

Seeing the Unseen:

Data Governance, Trust, and Fraud Protection

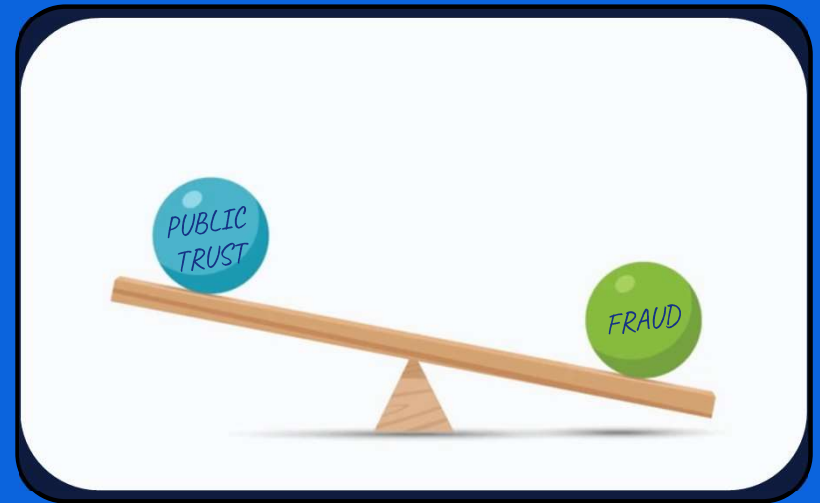
Kyle Rozanitis, Senior Solutions Architect

October 2025



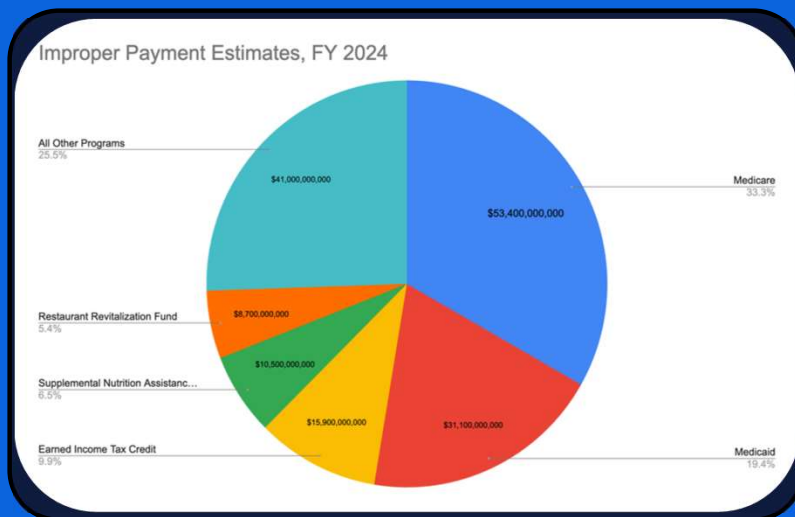
Fraud detection and public trust must move together.

Agencies can't protect taxpayer dollars without data, and they can't protect trust without governance.



Fraud drains billions and erodes trust.

According to GAO's government-wide estimates based on data from fiscal years 2018-2022, the federal government loses between \$233bn to \$521bn annually to fraud.



Data sourced from the U.S. Government Accountability Office (GAO).



Fraud is evolving faster than agencies can respond.

- ✗ Criminals exploit **automation and AI**
- ✗ Data volumes are **overwhelming**
- ✗ Manual detection is **too slow**

The Reality on the Ground

v1

Doesn't scale with the problem

Fraud detection often depends on exporting data into Excel. Analysts spend hours massaging rows and columns instead of investigating. It works in small bursts, but when datasets balloon into millions of rows, Excel buckles.

Searching for only the familiar

Analysts massage the data to find repeat patterns — multiple claims from the same IP, rare routing numbers, disposable email domains. But they only uncover what they already know to look for.

Blind to the unknown

Novel or anomalous fraud behaviors slip through the cracks. Without automation and machine learning, agencies are always one step behind.



The Reality on the Ground v2

Doesn't scale with the problem

Fraud detection often depends on exporting data into Excel. Analysts spend hours massaging rows and columns instead of investigating. It works in small bursts, but when datasets balloon into millions of rows, Excel buckles.

Searching for only the familiar

Analysts massage the data to find repeat patterns — multiple claims from the same IP, rare routing numbers, disposable email domains. But they only uncover what they already know to look for.

Blind to the unknown

Novel or anomalous fraud behaviors slip through the cracks. Without automation and machine learning, agencies are always one step behind.



Data powers fraud detection—but raises new risks.

The same citizen data that reveals suspicious patterns can undermine trust if misused.

Identity Data

Transaction Records

Network & Behavioral Data

The Governance Triangle

Collect only what's necessary

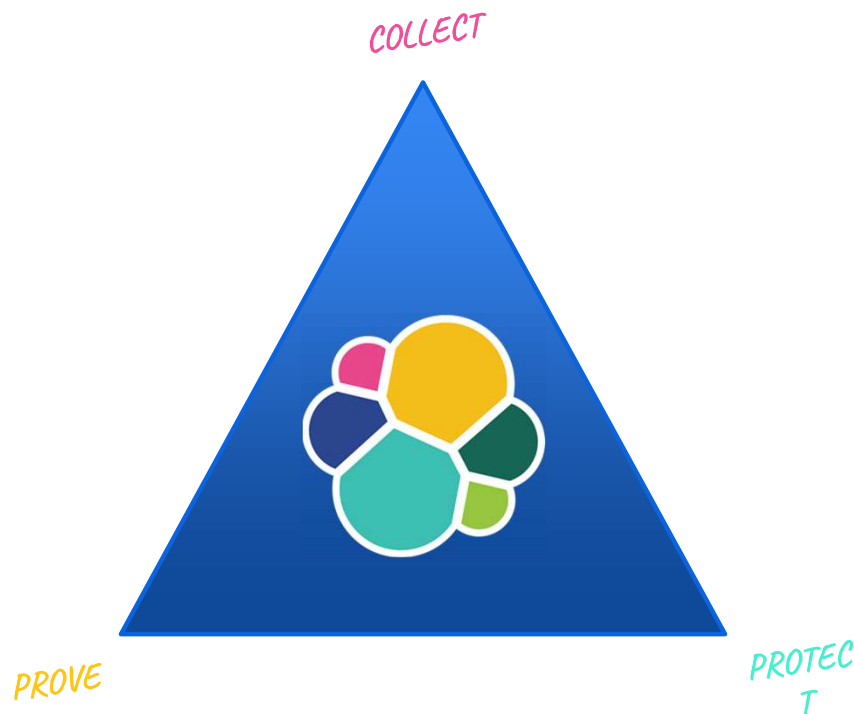
Good governance starts with restraint. Agencies should minimize the data they collect, focusing only on what's needed for fraud detection and compliance — nothing more.

Protect with layered security

Safeguards must be built in at every level: SSO, SAML, RBAC, and fine-grained controls like document-level and field-based access. Data should be encrypted in transit and at rest, and the systems, servers, and endpoints where that data lives must be hardened against attack.

Prove accountability and compliance

Trust requires visibility. Audit logging shows who accessed what, when, and how. Regular checks against frameworks like NIST CSF and state data policies demonstrate compliance. Transparency with citizens builds confidence that their information is being handled responsibly.



So what does this look like in practice?

Fraud Detection with Elastic

Spot the Fraud Hidden within the Noise

It's Monday morning at the Department of Labor.

A new wave of unemployment claims has arrived overnight — just like every day.

Most are legitimate. Some aren't.

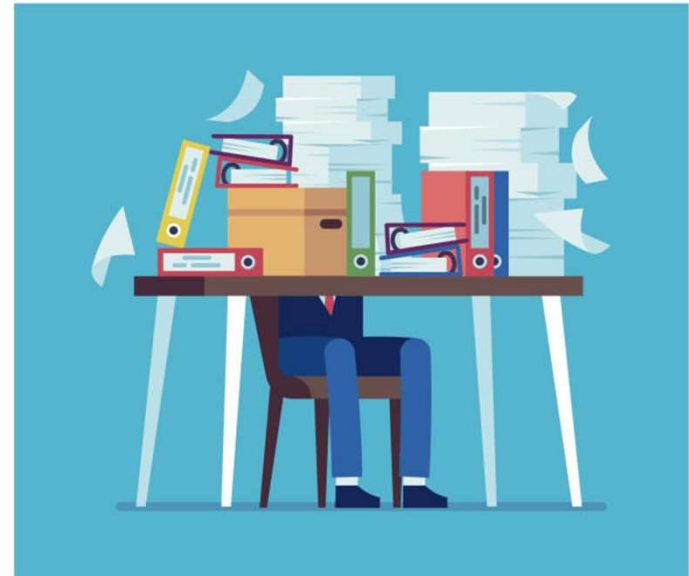
Among the thousands of claims, fraud is hiding:


- ✗ Multiple emails from the same IP
- ✗ Identities linked to the same bank account
- ✗ Suspicious spikes in ZIP codes

Fraud analysts are doing their best, but the tools are limited:

- ✗ Static reports and spreadsheets that don't scale
- ✗ Manual investigations that take hours
- ✗ Hunting only for known patterns

Meanwhile, fraudulent claims are getting approved and payouts are going out the door.

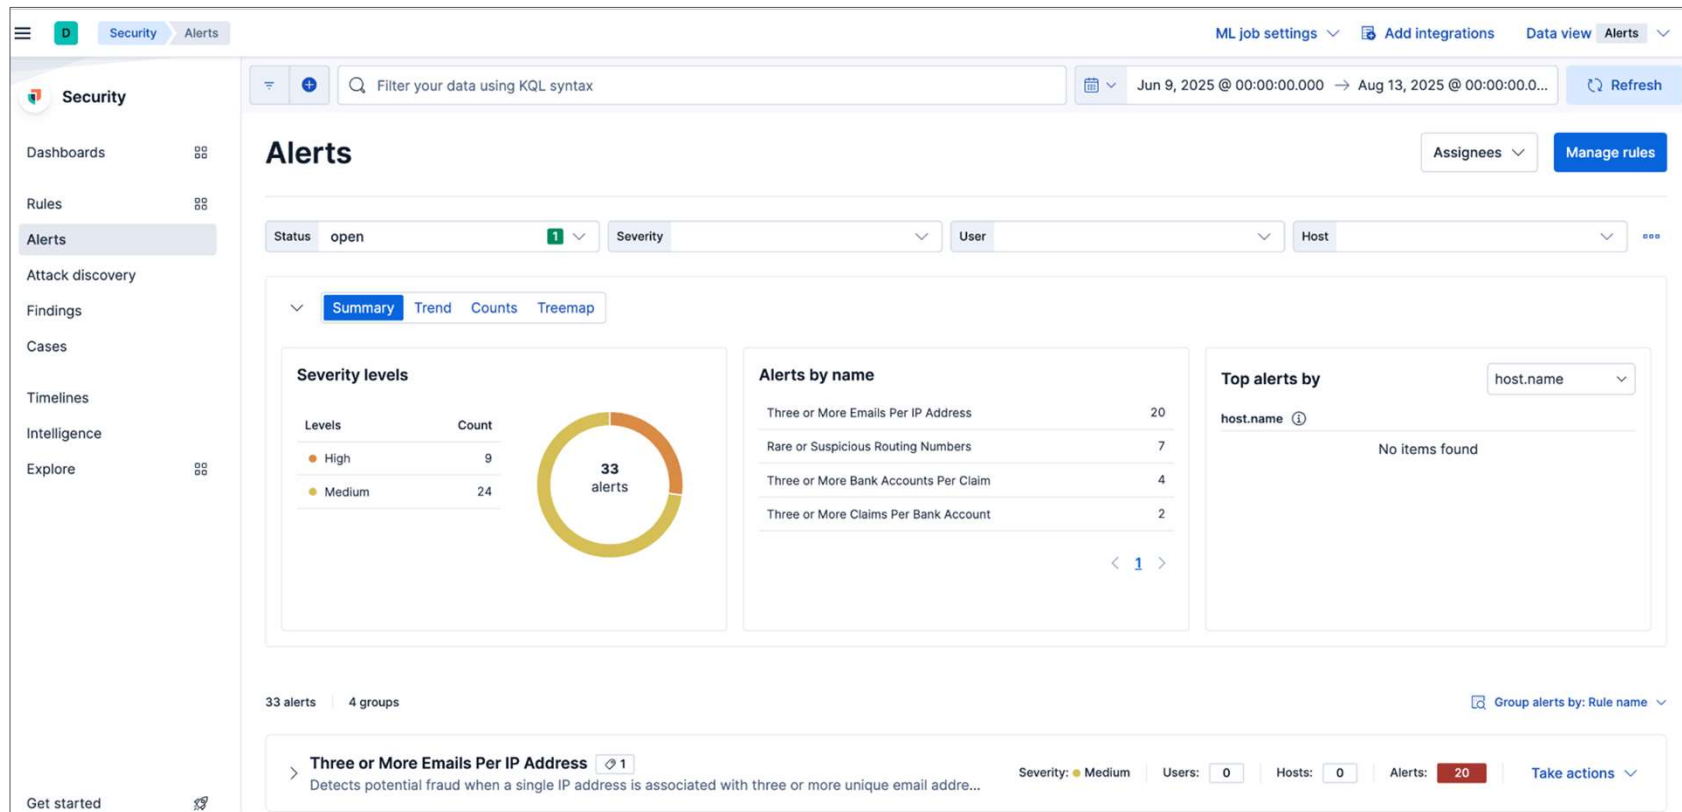


The background is a solid blue color. It features several white, 3D-style geometric shapes that look like architectural moldings or frames. These shapes are positioned in the corners and along the edges, creating a modern, minimalist aesthetic. One shape is in the top-left, another in the top-right, one in the bottom-left, and one in the bottom-right. They are all white with blue outlines and shadows, giving them a three-dimensional appearance.

Scene 1: A Suspicious Spike

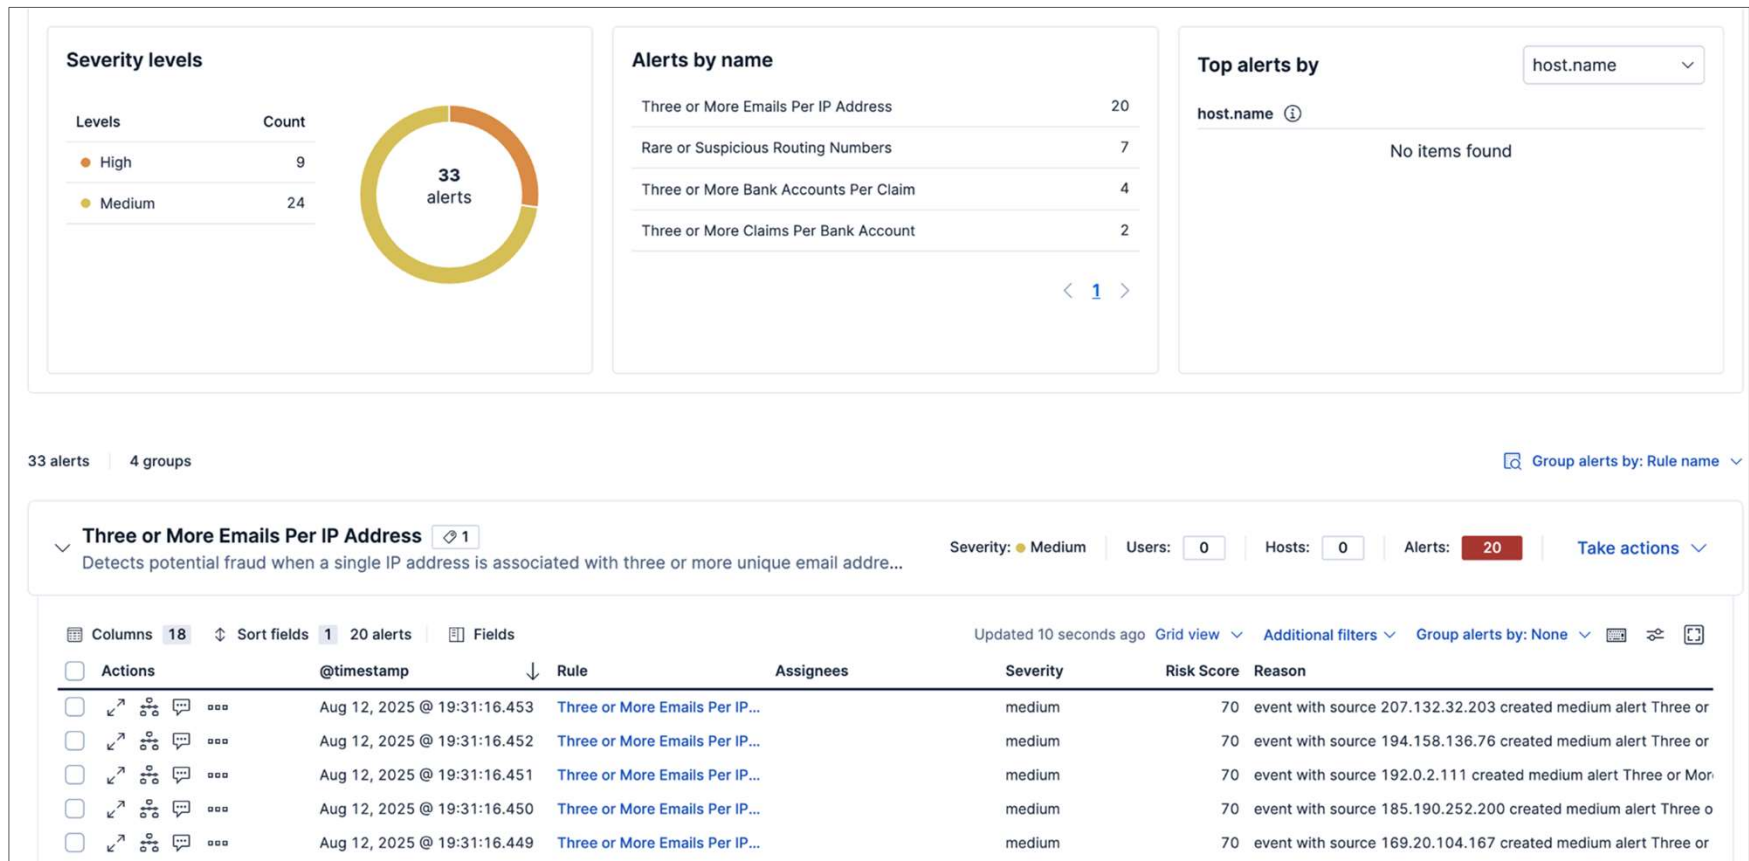
Fraud Detection with Elastic

Scene 1: A Suspicious Spike



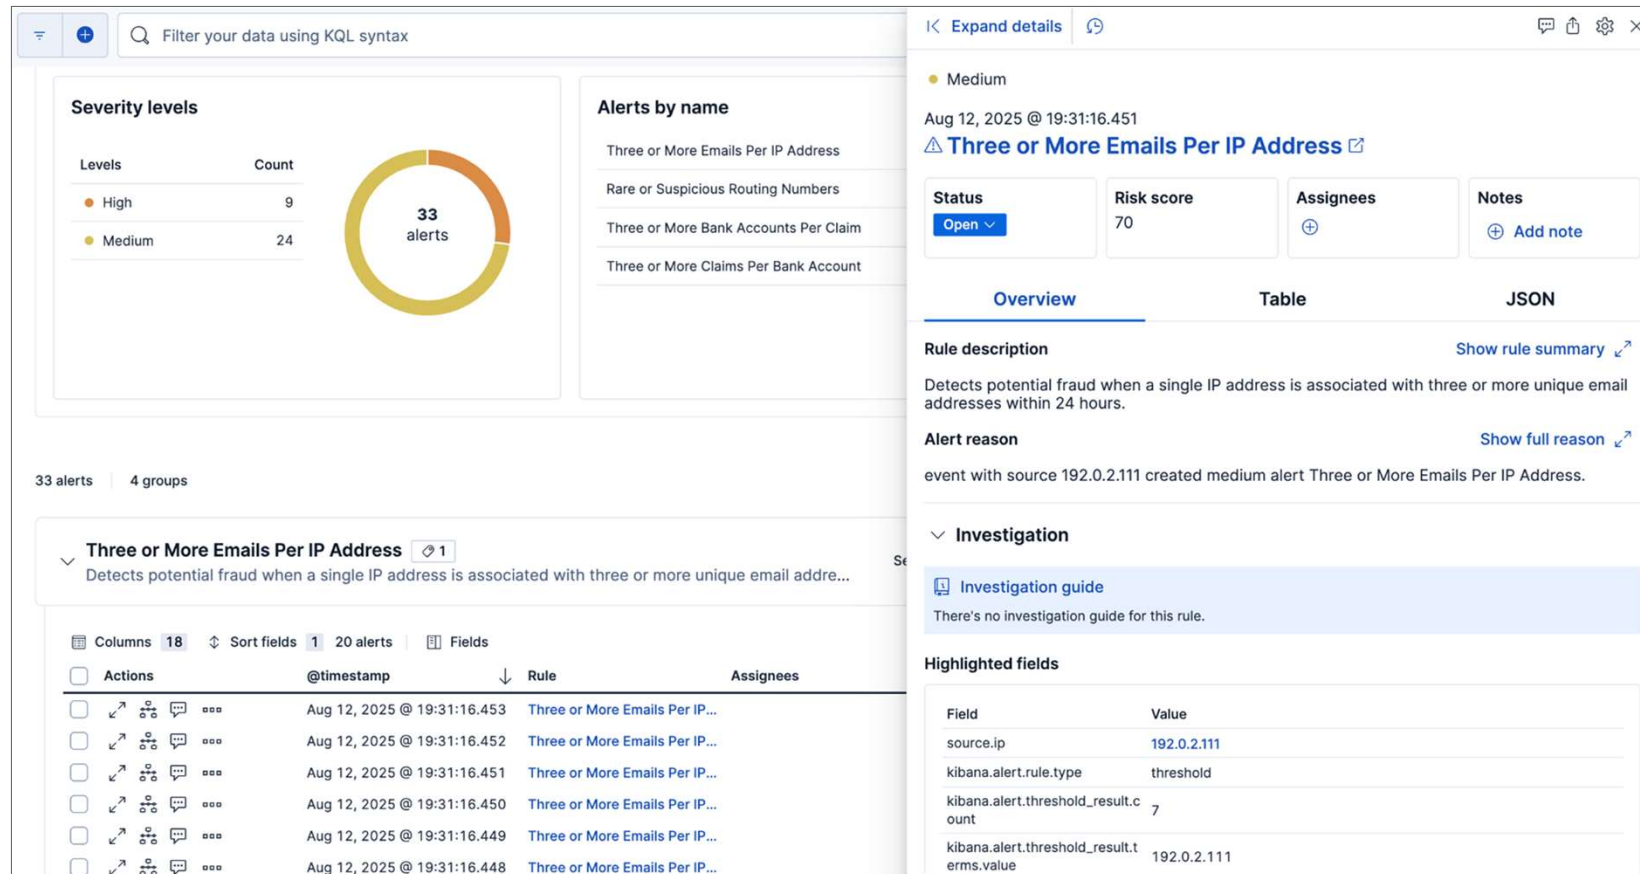
Fraud Detection with Elastic

Scene 1: A Suspicious Spike



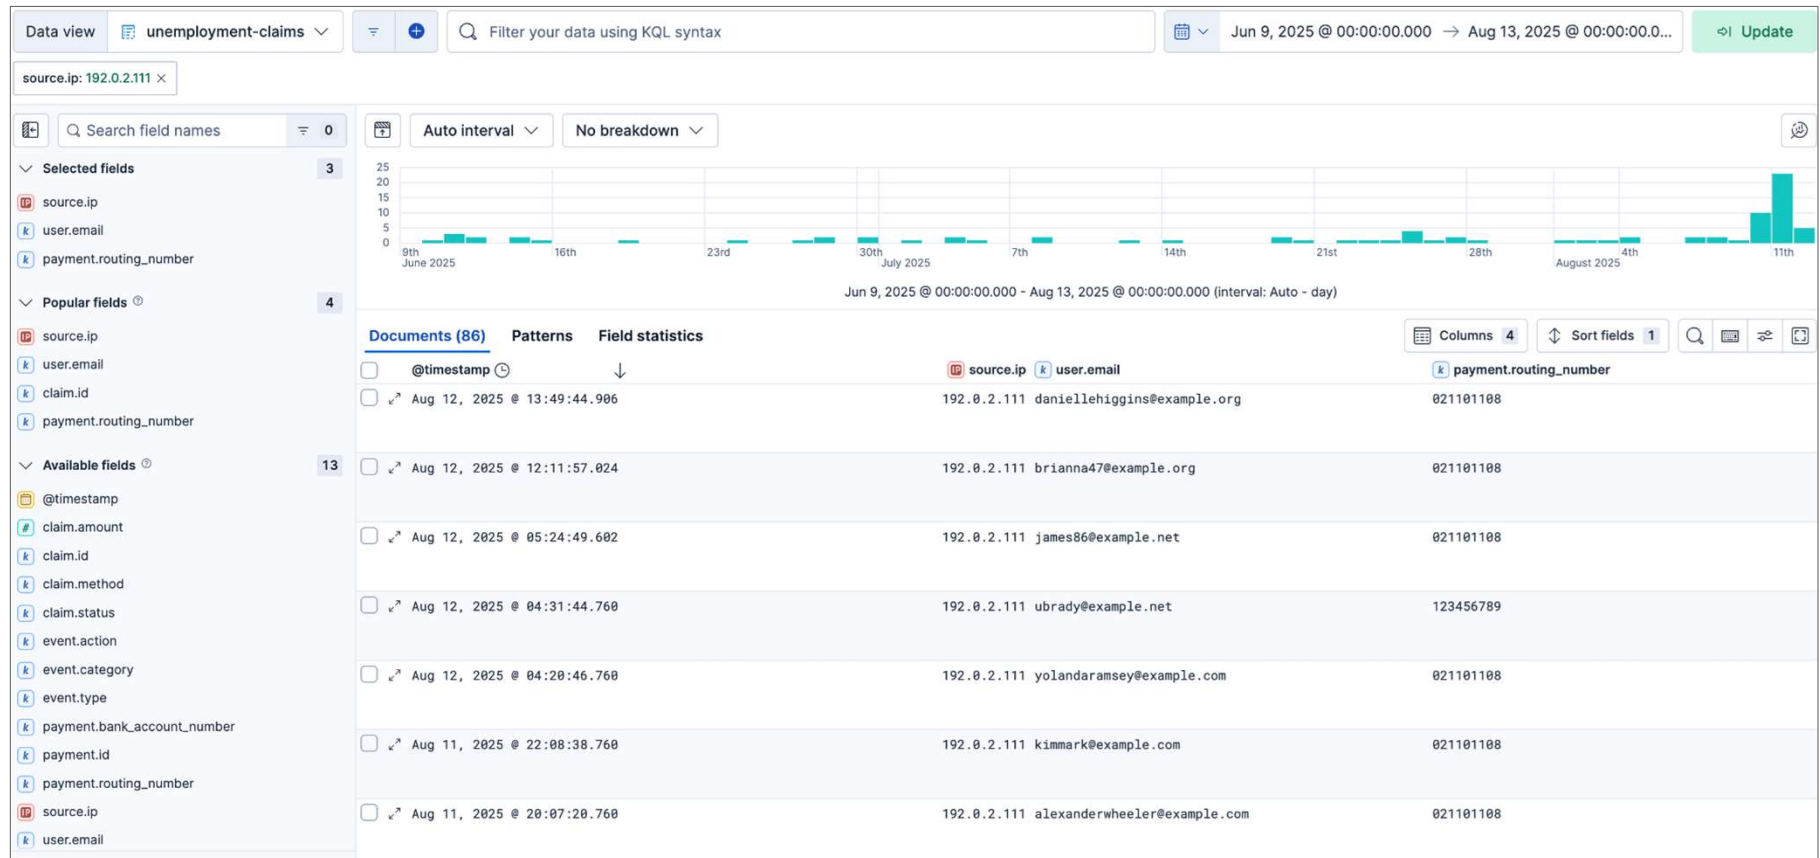
Fraud Detection with Elastic

Scene 1: A Suspicious Spike



Fraud Detection with Elastic

Scene 1: A Suspicious Spike



Fraud Detection with Elastic

Scene 1: A Suspicious Spike

ES|QL help

Jun 9, 2025 @ 00:00:00.000 → Aug 13, 2025 @ 00:00:00.0...

Run

1 FROM unemployment-claims

2 | WHERE source.ip: "192.0.2.111"

3 | KEEP @timestamp, source.ip, user.email

3 lines @timestamp found LIMIT 1000 rows

Submit feedback

Show recent queries

Search field names

0

Selected fields

3

@timestamp

source.ip

user.email

Available fields

3

@timestamp

source.ip

user.email

86 results

Columns 3

Sort fields 1

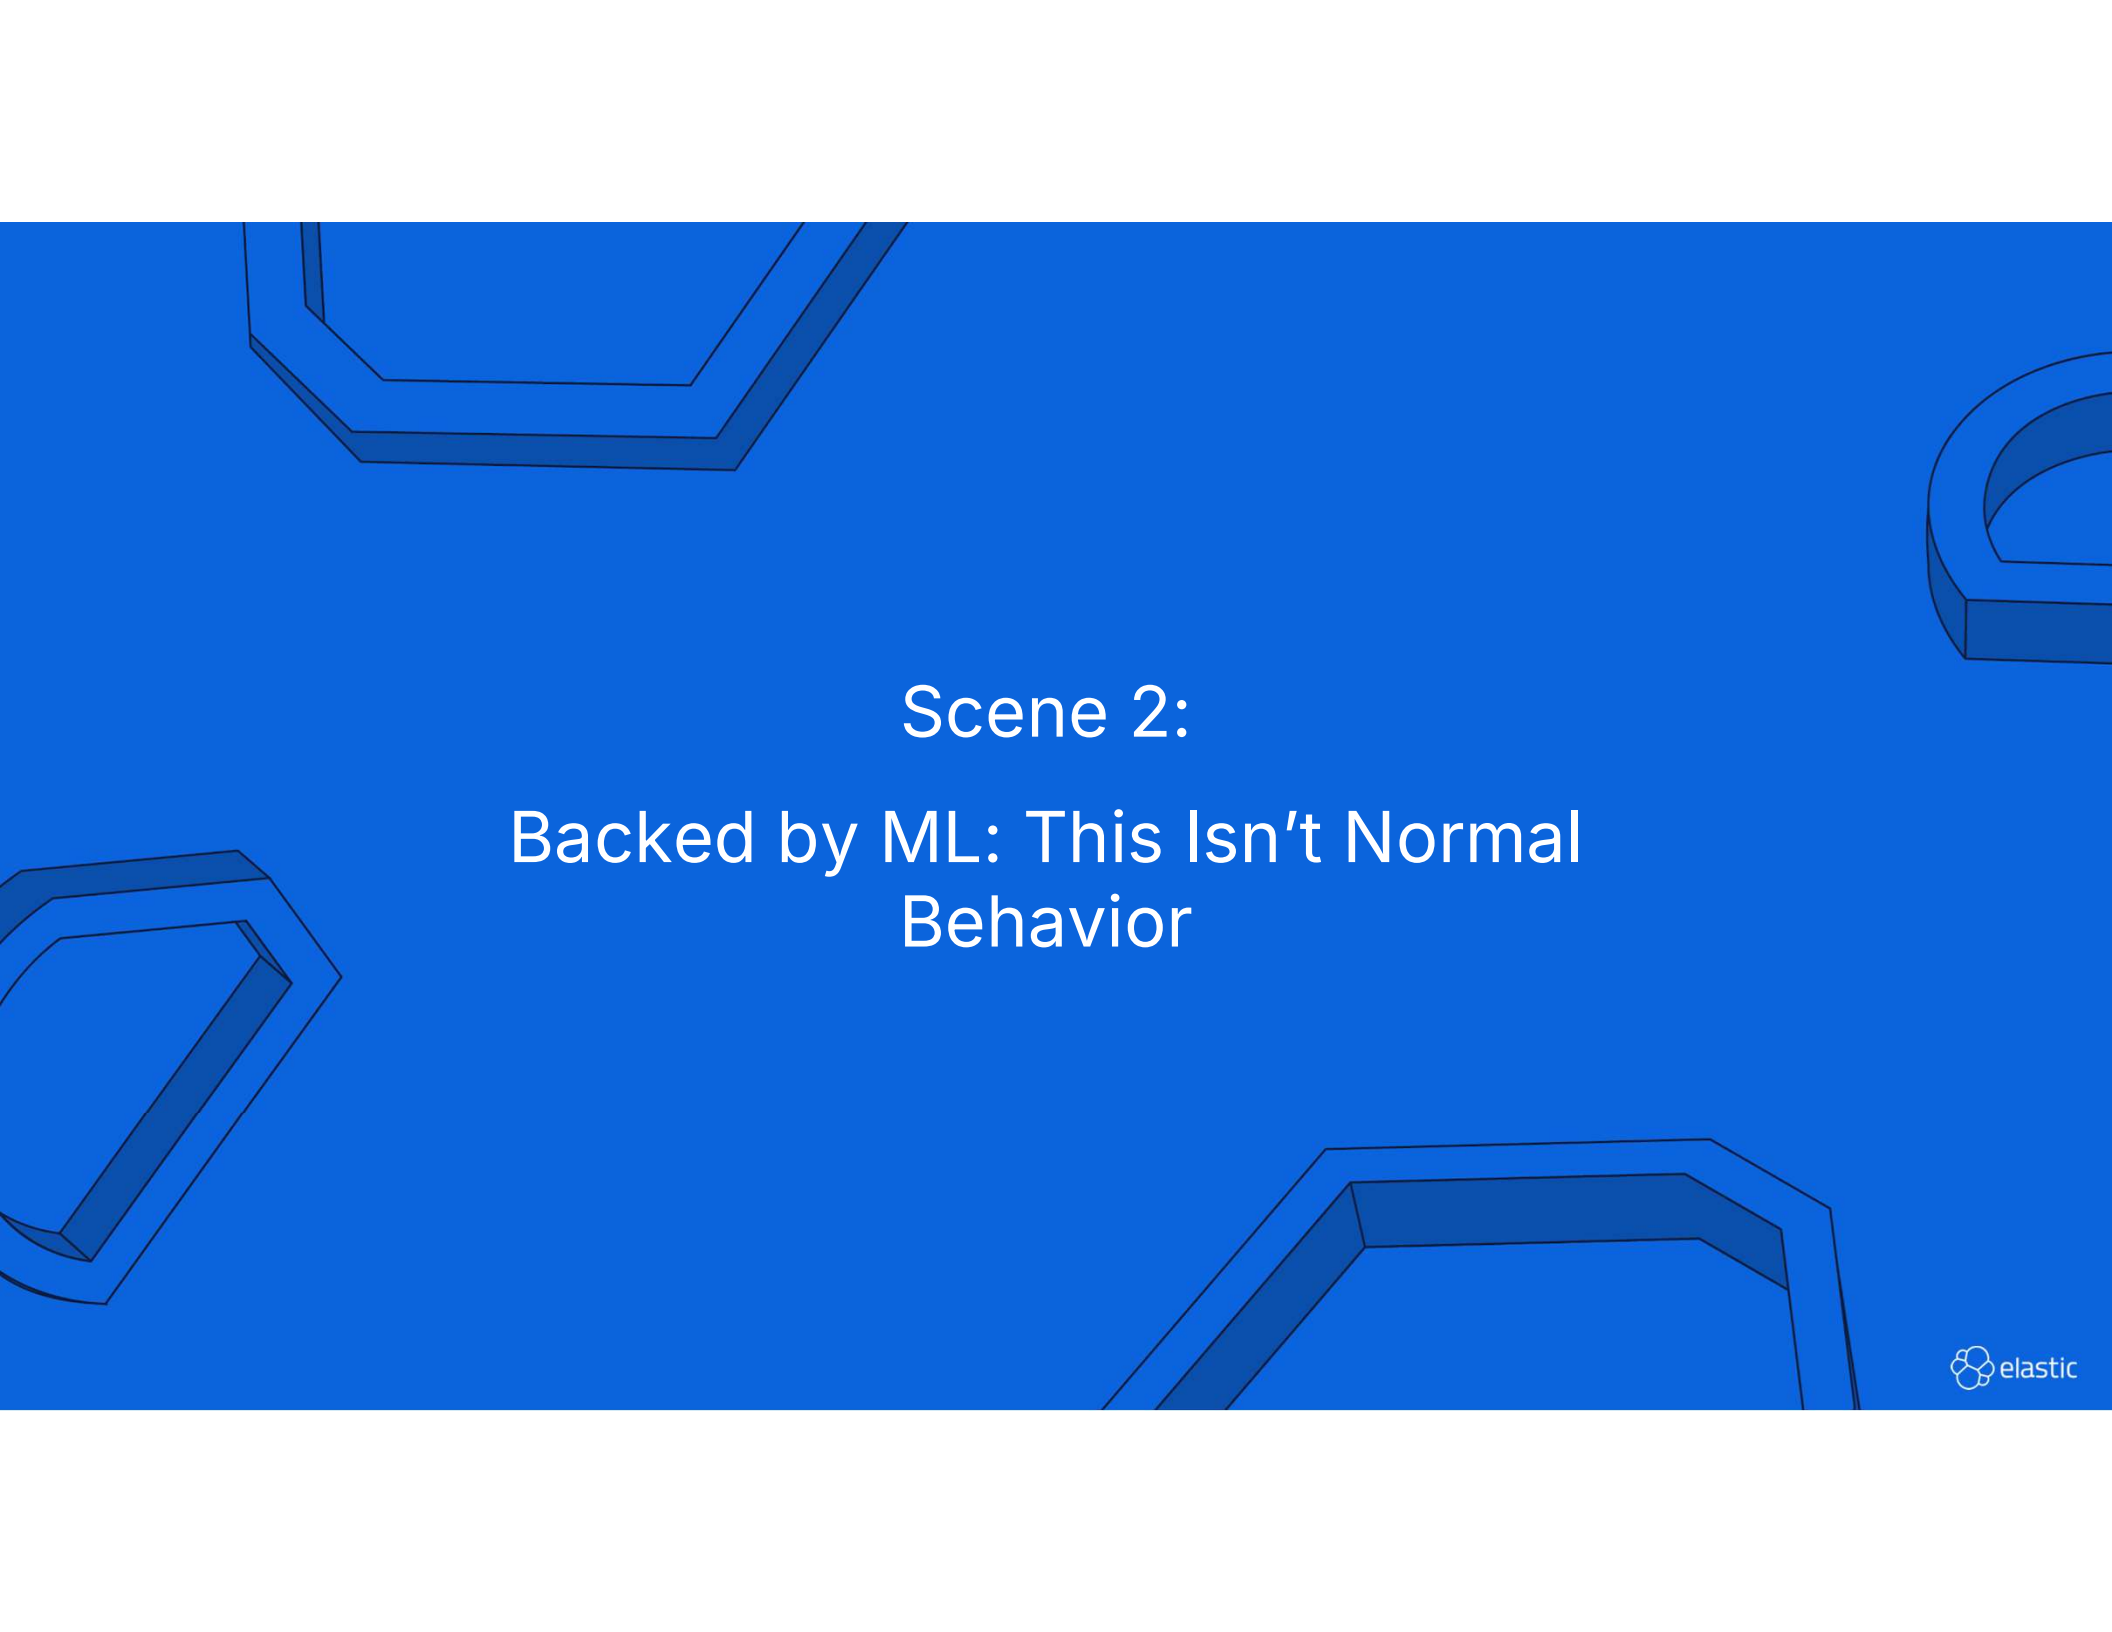
Q

+

+

+

@timestamp	source.ip	user.email
Aug 12, 2025 @ 13:49:44.986	192.0.2.111	daniellehiggins@example.org
Aug 12, 2025 @ 12:11:57.824	192.0.2.111	brianna47@example.org
Aug 12, 2025 @ 05:24:49.682	192.0.2.111	james86@example.net
Aug 12, 2025 @ 04:31:44.760	192.0.2.111	ubradu@example.net
Aug 12, 2025 @ 04:28:46.760	192.0.2.111	yolandaramsey@example.com
Aug 11, 2025 @ 22:08:38.760	192.0.2.111	kimmark@example.com
Aug 11, 2025 @ 20:07:20.760	192.0.2.111	alexanderwheeler@example.com
Aug 11, 2025 @ 19:29:46.760	192.0.2.111	edwincole@example.net
Aug 11, 2025 @ 19:21:51.760	192.0.2.111	mathew46@example.net

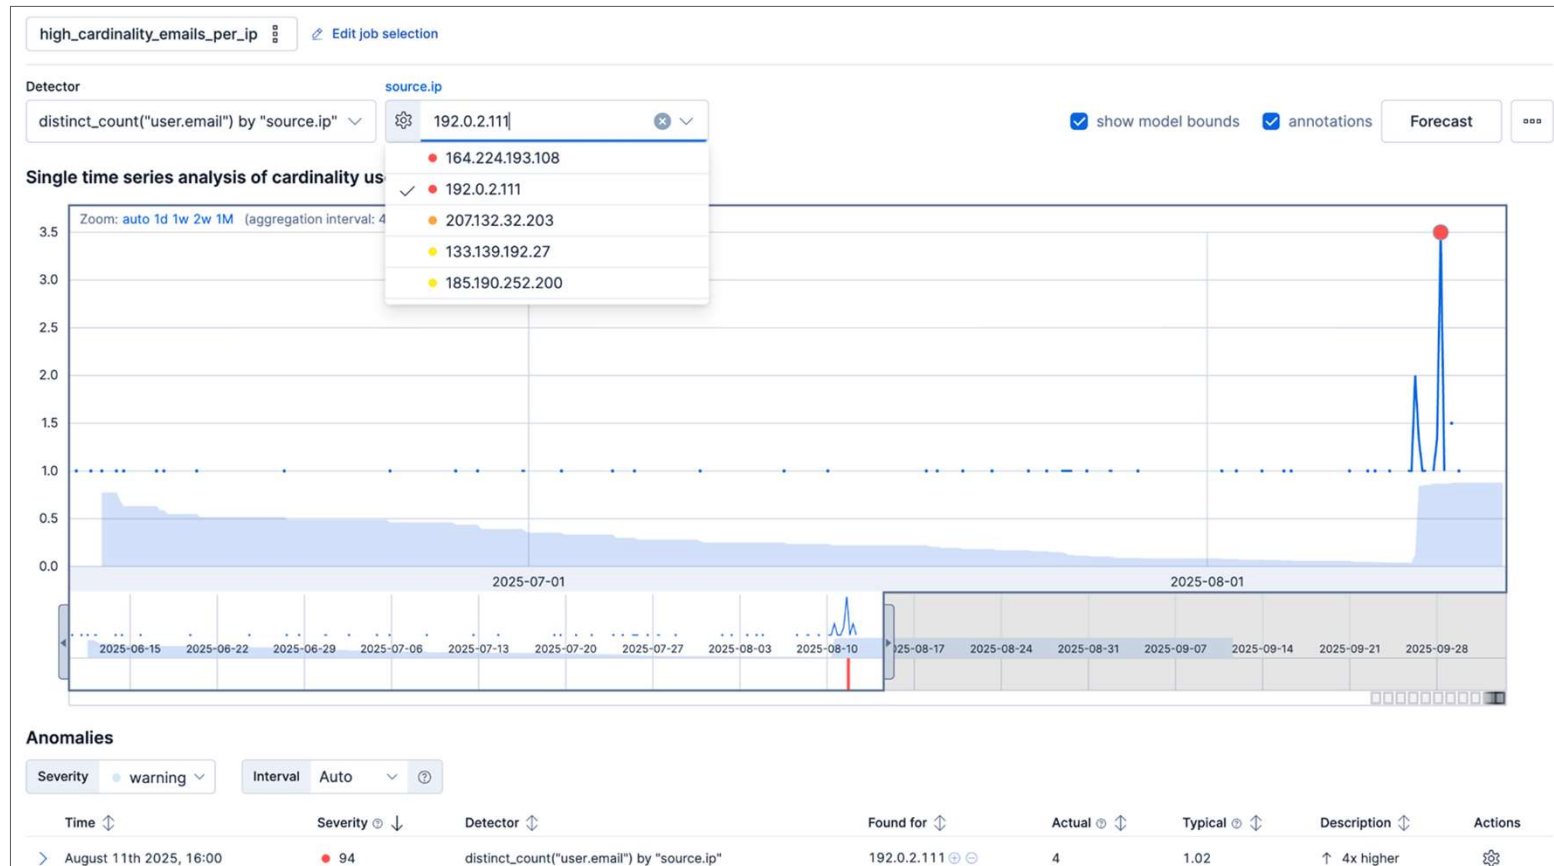
The background is a solid blue color. It features several white, three-dimensional geometric shapes that look like stylized architectural elements or frames. These shapes are positioned in the corners and along the edges, creating a sense of depth and structure. The shapes include octagons, rectangles, and curved forms, all rendered with white outlines and some internal shading to give them a 3D appearance.

Scene 2:

Backed by ML: This Isn't Normal Behavior

Fraud Detection with Elastic

Scene 2: Backed by ML: This Isn't Normal Behavior




Fraud Detection with Elastic

Scene 2: Backed by ML: This Isn't Normal Behavior

Select data view or saved Discover session

✕

Types ▼ [+ Create a data view](#)

Type	Title
	Unemployment Claims

< [1](#) >

Fraud Detection with Elastic

Scene 2: Backed by ML: This Isn't Normal Behavior

Create a job from the data view Unemployment Claims

Use a wizard



Single metric

Detect anomalies in a single time series.



Multi-metric

Detect anomalies with one or more metrics and optionally split the analysis.



Population

Detect unusual activity in a population. Recommended for high cardinality data.



Advanced

Use the full range of options to create a job for more advanced use cases.



Categorization

Group log messages into categories and detect anomalies within them.



Rare

Detect rare values in time series data.

Learn more about your data

If you're not sure what type of job to create, first explore the fields and metrics in your data.



Data Visualizer

Learn more about the characteristics of your data and identify the fields for analysis with machine learning.

Fraud Detection with Elastic

Scene 2: Backed by ML: This Isn't Normal Behavior

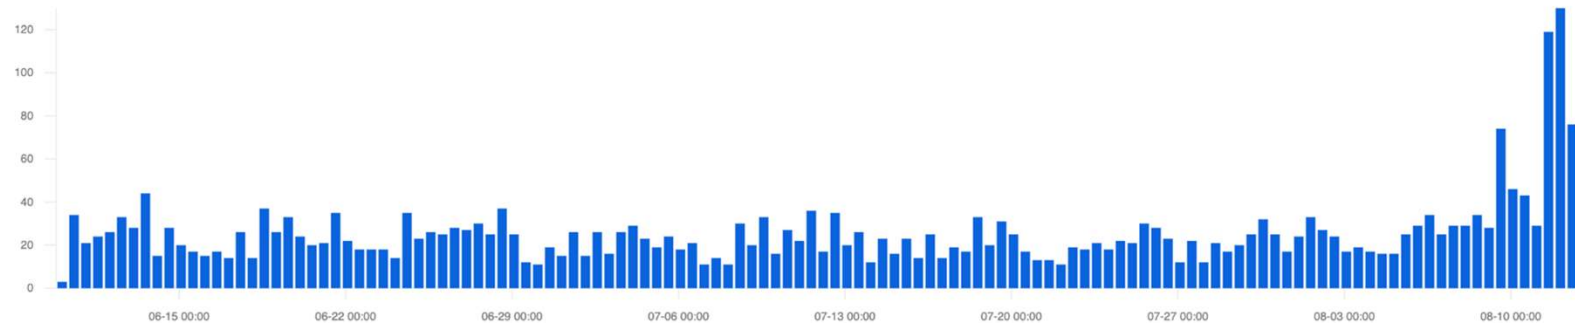
Create job: Multi-metric

Using data view Unemployment Claims (unemployment-*)



Time range

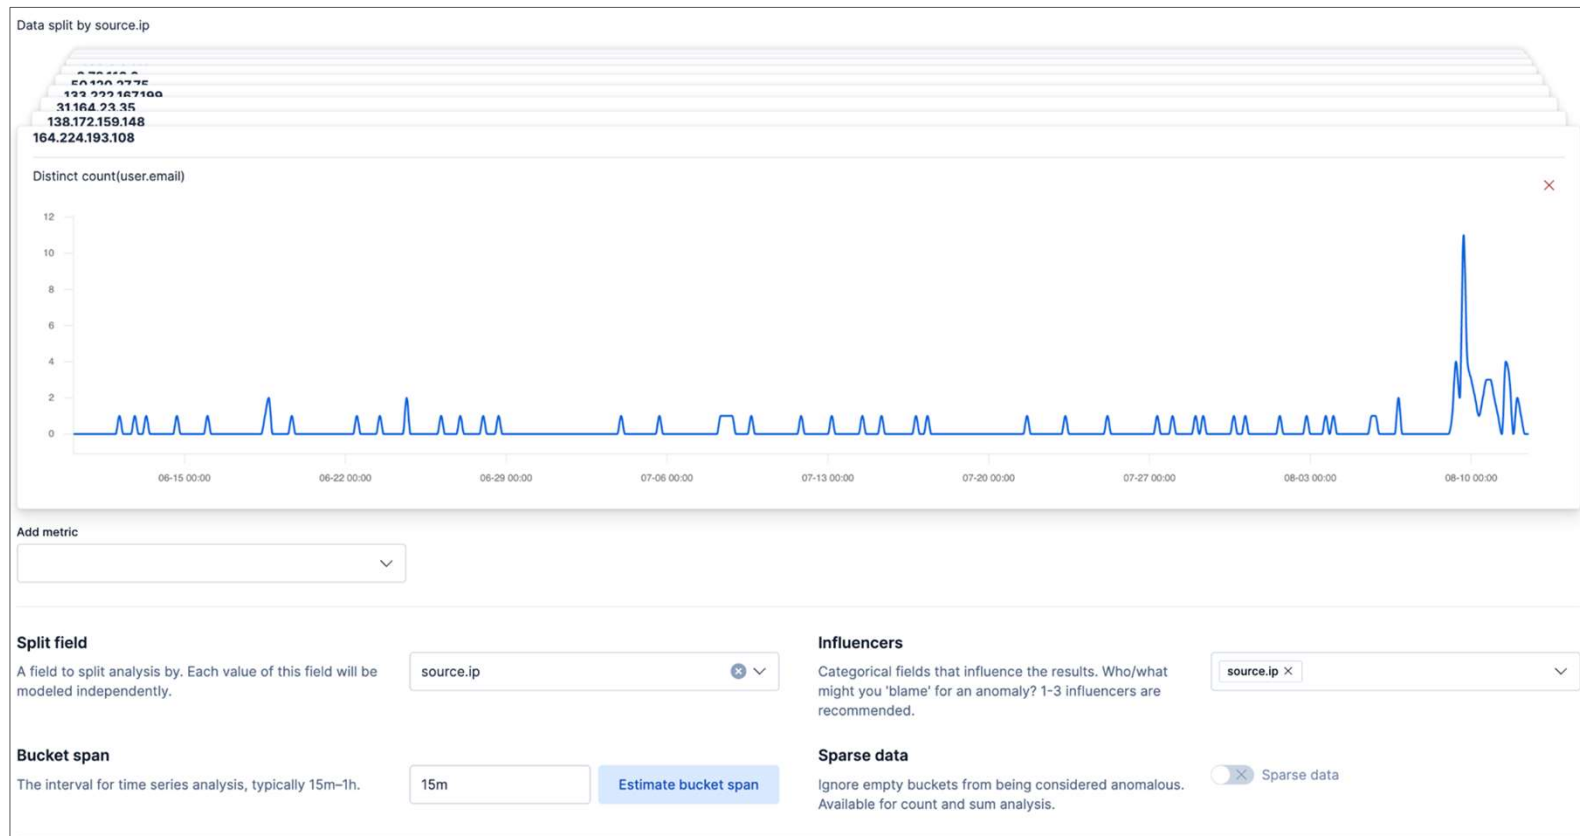
Jun 10, 2025 @ 05:20:38.760 → Aug 12, 2025 @ 15:31:13.528 [Use full data](#)



[Next >](#)

Fraud Detection with Elastic

Scene 2: Backed by ML: This Isn't Normal Behavior



Fraud Detection with Elastic

Scene 2: Backed by ML: This Isn't Normal Behavior

Create job: Multi-metric

Using data view Unemployment Claims (unemployment-*)

✓

✓

3

4

5

Time range

Choose fields

Job details

Validation

Summary

Job details

Job ID

A unique identifier for the job. Spaces and the characters / ? , " < > | * are not allowed

count-emails-per-ip

Job description

Optional descriptive text.

Groups

Optional grouping for jobs. New groups can be created or picked from the list of existing groups.

fraud_detection × × ▾

> Additional settings

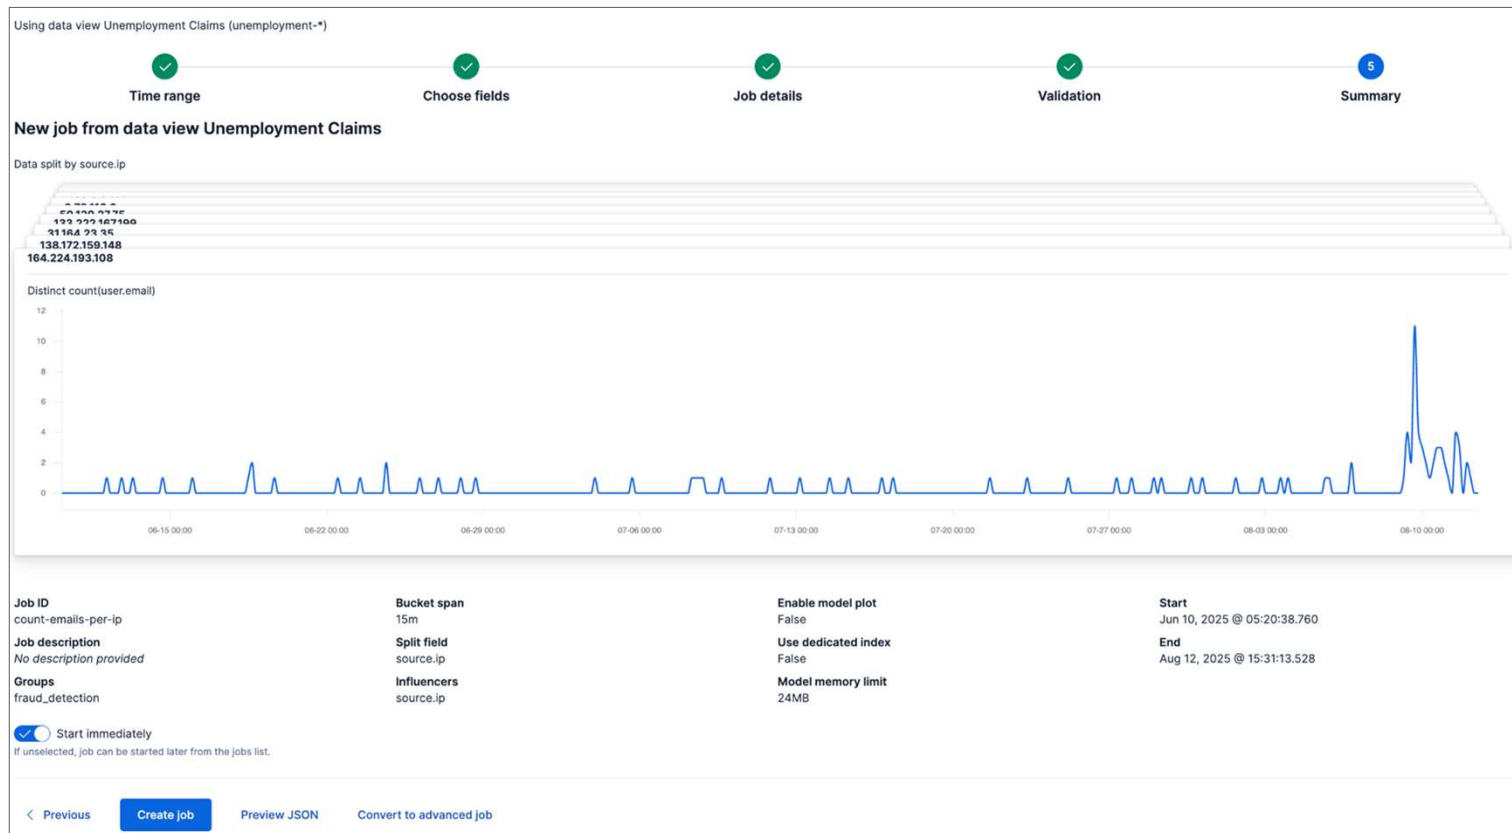
> Advanced

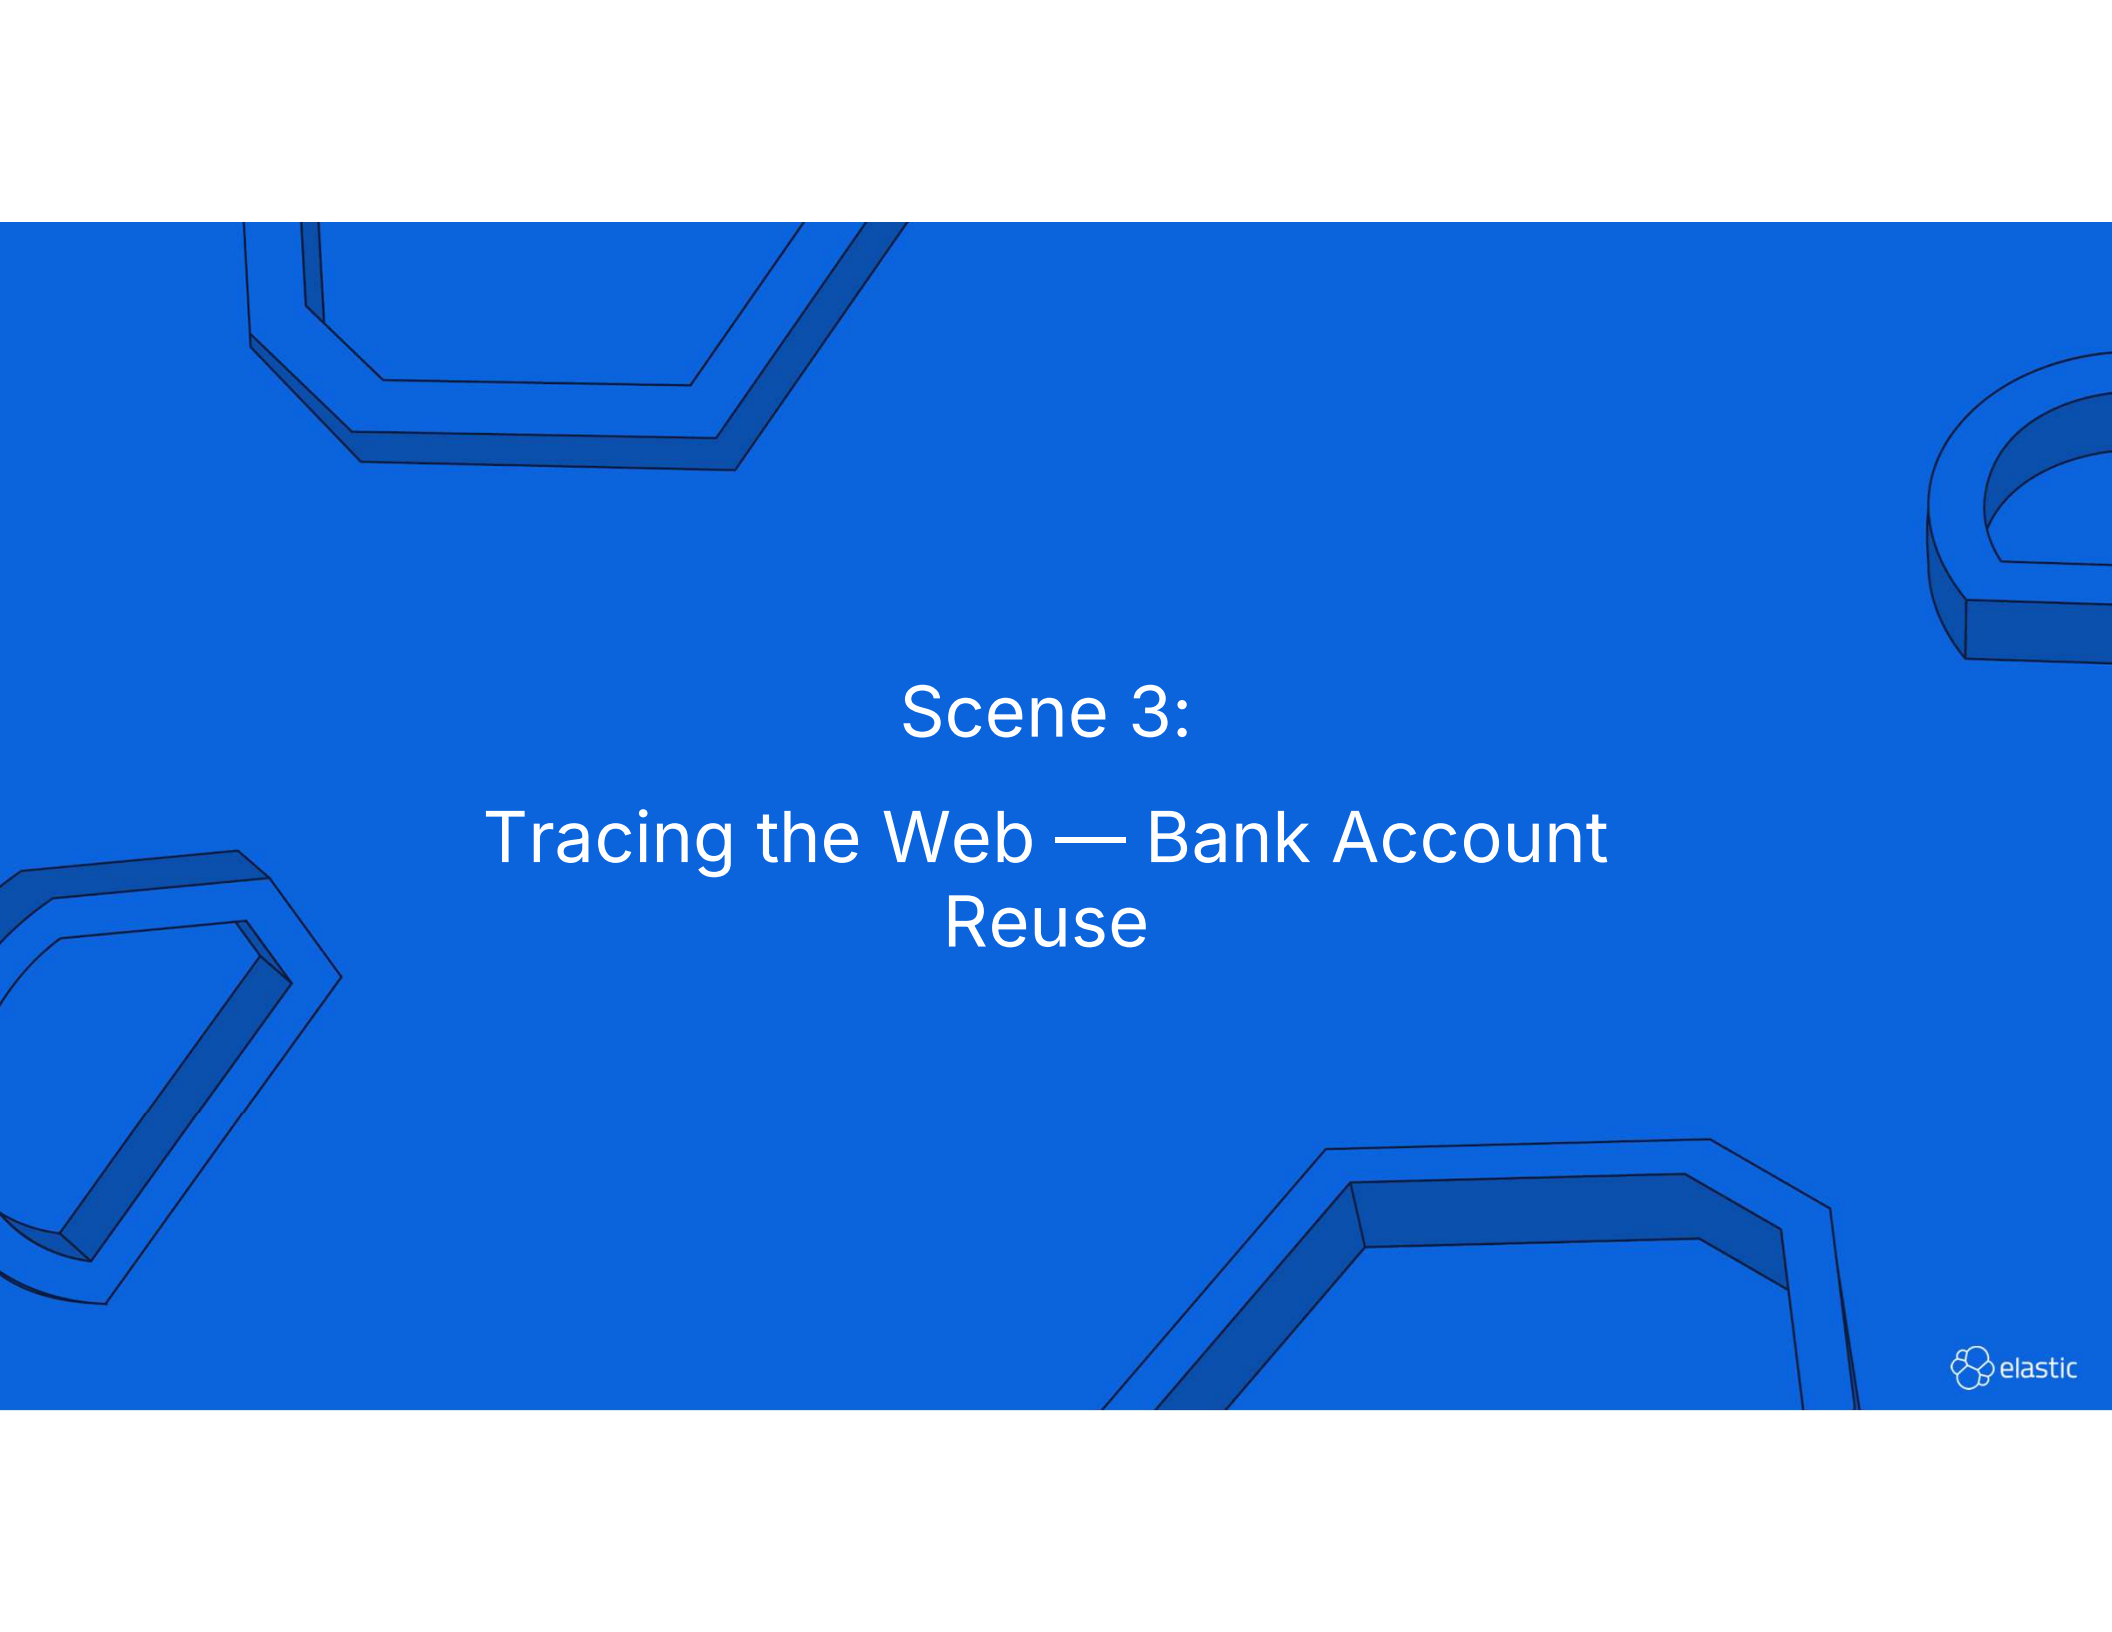
< Previous

Next >

Fraud Detection with Elastic

Scene 2: Backed by ML: This Isn't Normal Behavior



The background is a solid blue color. It features several white, 3D-style geometric shapes that look like stylized window frames or architectural elements. These shapes are positioned in the corners and along the edges, creating a modern, abstract aesthetic. The shapes include octagons, rectangles, and a semi-circle, all rendered with white outlines and some internal shading to give them a three-dimensional appearance.

Scene 3:

Tracing the Web — Bank Account Reuse

Fraud Detection with Elastic

Scene 3: Tracing the Web — Bank Account Reuse

33 alerts 4 groups Group alerts by: Rule name

> **Three or More Emails Per IP Address** 1
Detects potential fraud when a single IP address is associated with three or more unique email addresses within 24 hours. Severity: Medium Users: 0 Hosts: 0 Alerts: 20 [Take actions](#)

> **Rare or Suspicious Routing Numbers** 1
Detects unemployment claims made to known suspicious or rare routing numbers, indicating possible fraud or money laundering. Severity: High Users: 0 Hosts: 0 Alerts: 7 [Take actions](#)

✓ **Three or More Bank Accounts Per Claim** 1
Detects potential fraud when a single claim is associated with three or more unique bank accounts within 24 hours. Severity: Medium Users: 0 Hosts: 0 Alerts: 4 [Take actions](#)

Columns 18 Sort fields 1 4 alerts Fields Updated now Grid view Additional filters Group alerts by: None

Actions	@timestamp	Rule	Assignees	Severity	Risk Score	Reason	host.name
<input type="checkbox"/>	Aug 12, 2025 @ 19:09:08.642	Three or More Bank Accou...		medium	70	event created medium alert Three or More Bank Accounts Per Claim.	—
<input type="checkbox"/>	Aug 12, 2025 @ 19:09:08.641	Three or More Bank Accou...		medium	70	event created medium alert Three or More Bank Accounts Per Claim.	—
<input type="checkbox"/>	Aug 12, 2025 @ 10:54:00.767	Three or More Bank Accou...		medium	70	event created medium alert Three or More Bank Accounts Per Claim.	—
<input type="checkbox"/>	Aug 12, 2025 @ 10:54:00.766	Three or More Bank Accou...		medium	70	event created medium alert Three or More Bank Accounts Per Claim.	—

> **Three or More Claims Per Bank Account** 1
Detects potential fraud when a single bank account is associated with three or more unique claims within 24 hours. Severity: High Users: 0 Hosts: 0 Alerts: 2 [Take actions](#)

Fraud Detection with Elastic

Scene 3: Tracing the Web — Bank Account Reuse

The screenshot displays the Elastic Security console interface. On the left, a summary card shows 33 alerts, with 24 at Medium severity. Below this, a list of rules is shown, including 'Three or More Bank Accounts Per Claim', 'Three or More Claims Per Bank Account', 'Three or More Emails Per IP Address', 'Rare or Suspicious Routing Numbers', and 'Three or More Bank Accounts Per Claim'. The 'Three or More Claims Per Bank Account' rule is selected, showing its description: 'Detects potential fraud when a single bank account is associated with three or more unique claims within 24 hours.' Below the rule list, a table shows two alerts. The first alert, dated Aug 12, 2025 @ 19:31:21.898, is titled 'Three or More Claims Per Bank Account' and has a high severity. The second alert, dated Aug 11, 2025 @ 16:54:08.869, is also titled 'Three or More Claims Per Bank Account' and has a high severity. The right panel shows the details for the selected alert, including its status (Open), risk score (90), and assignees. It also provides an overview of the rule, its description, and the alert reason. The 'Investigation' section shows an investigation guide with the text 'There's no investigation guide for this rule.' The 'Highlighted fields' section lists fields such as 'payment.bank_account_number', 'kibana.alert.rule.type', 'kibana.alert.threshold_result.count', 'kibana.alert.threshold_result.term.s.value', and 'kibana.alert.threshold_result.cardinality.field'. At the bottom right, there is a 'Take action' button.

Medium 24 33 alerts

Three or More Bank Accounts Per Claim

Three or More Claims Per Bank Account

33 alerts 4 groups

> **Three or More Emails Per IP Address** 1

Detects potential fraud when a single IP address is associated with three or more unique email addresses within 24 hours.

> **Rare or Suspicious Routing Numbers** 1

Detects unemployment claims made to known suspicious or rare routing numbers, indicating possible fraud or money laundering.

> **Three or More Bank Accounts Per Claim** 1

Detects potential fraud when a single claim is associated with three or more unique bank accounts within 24 hours.

> **Three or More Claims Per Bank Account** 1

Detects potential fraud when a single bank account is associated with three or more unique claims within 24 hours.

Columns 18 Sort fields 1 2 alerts Fields

Actions	@timestamp	Rule	Assignees	Severity
<input type="checkbox"/>	Aug 12, 2025 @ 19:31:21.898	Three or More Claims Per B...		high
<input type="checkbox"/>	Aug 11, 2025 @ 16:54:08.869	Three or More Claims Per B...		high

Rows per page: 25

Untitled timeline 186 Unsaved

High

Aug 12, 2025 @ 19:31:21.898

Three or More Claims Per Bank Account

Status **Open**

Risk score 90

Assignees

Notes Add note

Overview Table JSON

Rule description [Show rule summary](#)

Detects potential fraud when a single bank account is associated with three or more unique claims within 24 hours.

Alert reason [Show full reason](#)

event created high alert Three or More Claims Per Bank Account.

Investigation

[Investigation guide](#)

There's no investigation guide for this rule.

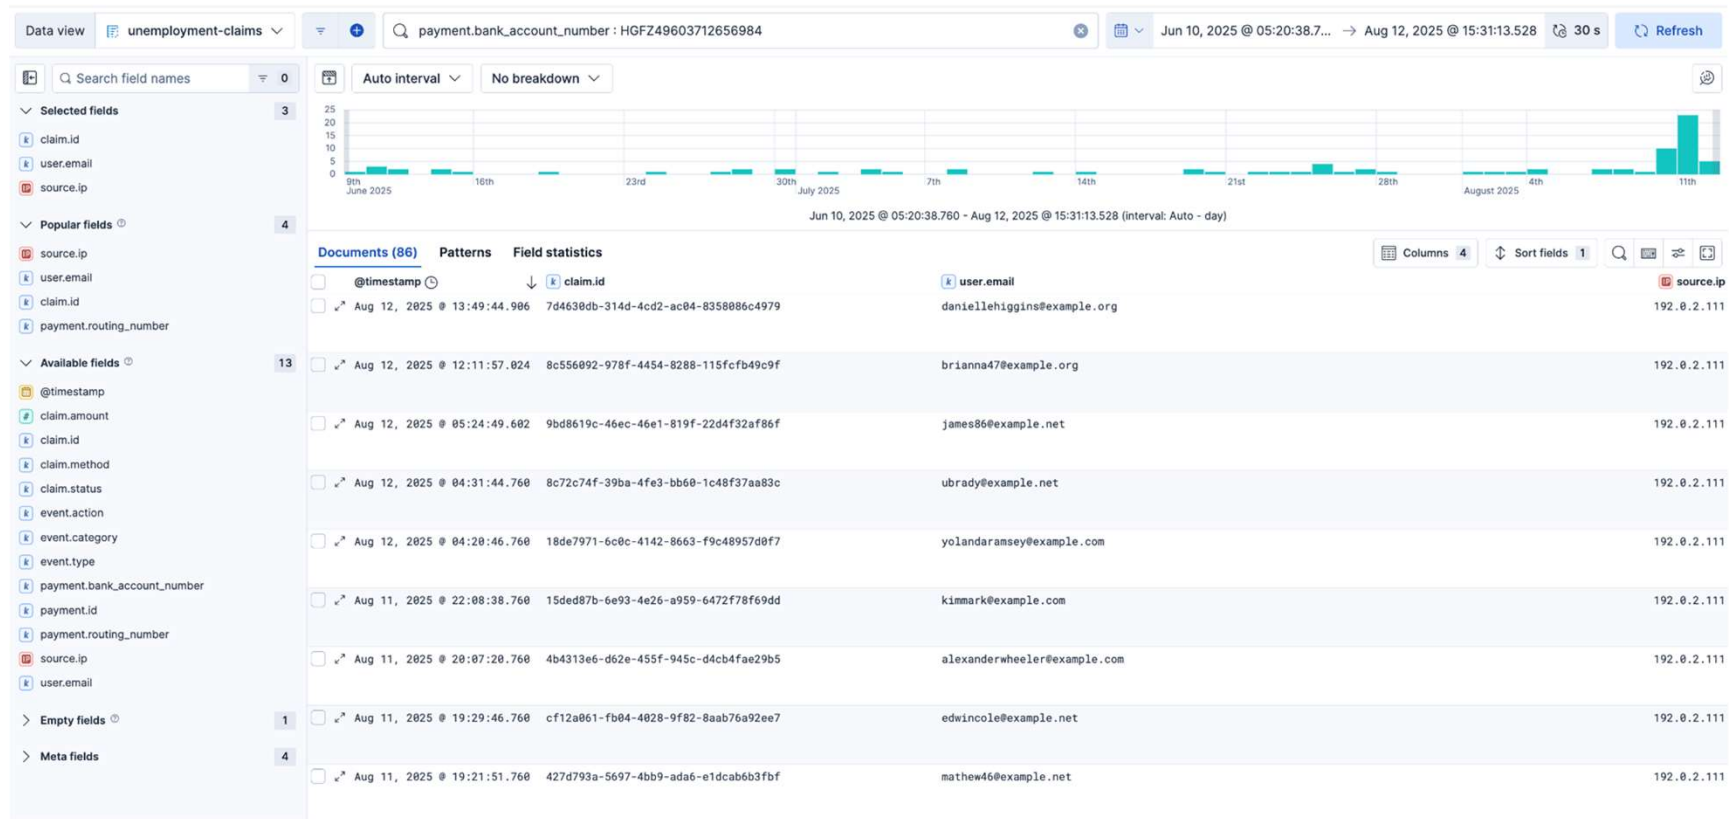
Highlighted fields

Field	Value
payment.bank_account_number	HGFZ49603712656984
kibana.alert.rule.type	threshold
kibana.alert.threshold_result.count	7
kibana.alert.threshold_result.term.s.value	HGFZ49603712656984
kibana.alert.threshold_result.cardinality.field	claim.id

Take action

Fraud Detection with Elastic

Scene 3: Tracing the Web — Bank Account Reuse



Fraud Detection with Elastic

Scene 3: Tracing the Web — Bank Account Reuse

New chat

ES|QL Query for Unemployment Claims with Specific Bank Account Number Elastic Managed LLM

You are now using the Elastic Managed LLM connector
Elastic AI Assistant and other AI features are powered by an LLM. The Elastic Managed LLM connector is used by default (additional costs incur) when no custom connectors are available. You can configure a custom connector if you prefer.

You at: 10/1/2025, 11:59:56 AM

Write me an ES|QL query that searches the `unemployment-claims` index and retrieves the "claim.id", "user.email", and "source.ip" where the "payment.bank_account_number" is "HGFZ49603712656984"?

Assistant at: 10/1/2025, 12:00:14 PM

```
FROM unemployment-claims
| WHERE payment.bank_account_number == "HGFZ49603712656984"
| KEEP claim.id, user.email, source.ip
```

You at: 10/1/2025, 12:01:07 PM

This shows me records with duplicate email addresses. Modify the ES|QL query to show distinct `user.email` and `source.ip` values for records with the specified bank account number.

Assistant at: 10/1/2025, 12:01:33 PM

```
FROM unemployment-claims
| WHERE payment.bank_account_number == "HGFZ49603712656984"
| STATS emails = VALUES(user.email), ips = VALUES(source.ip)
```

Ask me anything from 'summarize this alert' to 'help me build a query...'

Workflow suggestions

Query conversion

Agent integration advice

+ Add quick prompt...

New chat

ES|QL Query for Unemployment Claims with Specific Bank Account Number Elastic Managed LLM

You at: 10/1/2025, 12:01:07 PM

This shows me records with duplicate email addresses. Modify the ES|QL query to show distinct `user.email` and `source.ip` values for records with the specified bank account number.

Assistant at: 10/1/2025, 12:01:33 PM

```
FROM unemployment-claims
| WHERE payment.bank_account_number == "HGFZ49603712656984"
| STATS emails = VALUES(user.email), ips = VALUES(source.ip)
```

This query will:

1. Filter records with the specified bank account number
2. Use the STATS command to collect all unique values of `user.email` and `source.ip`
3. Return these values as multi-valued fields named "emails" and "ips"

If you want the values to be sorted, you can use `MV_SORT` like this:

```
FROM unemployment-claims
| WHERE payment.bank_account_number == "HGFZ49603712656984"
| STATS emails = MV_SORT(VALUES(user.email)), ips = MV_SORT(VALUES(source.ip))
```

Regenerate

Ask me anything from 'summarize this alert' to 'help me build a query...'

Workflow suggestions

Query conversion

Agent integration advice

+ Add quick prompt...

Fraud Detection with Elastic

Scene 3: Tracing the Web — Bank Account Reuse

ES|QL help

Jun 10, 2025 @ 05:20:38.7... → Aug 12, 2025 @ 15:31:13.528 30 s Refresh

```
1 FROM unemployment-claims
2 | WHERE payment.bank_account_number == "HGFZ49603712656984"
3 | STATS distinct_emails = VALUES(user.email), distinct_ips = VALUES(source.ip)
```

3 lines @timestamp found LIMIT 1000 rows 1 warning

Submit feedback Show recent queries

Search field names 0 1 result

Columns 2 Sort fields

Selected fields 2

- distinct_emails
- distinct_ips

Available fields 2

- distinct_emails
- distinct_ips

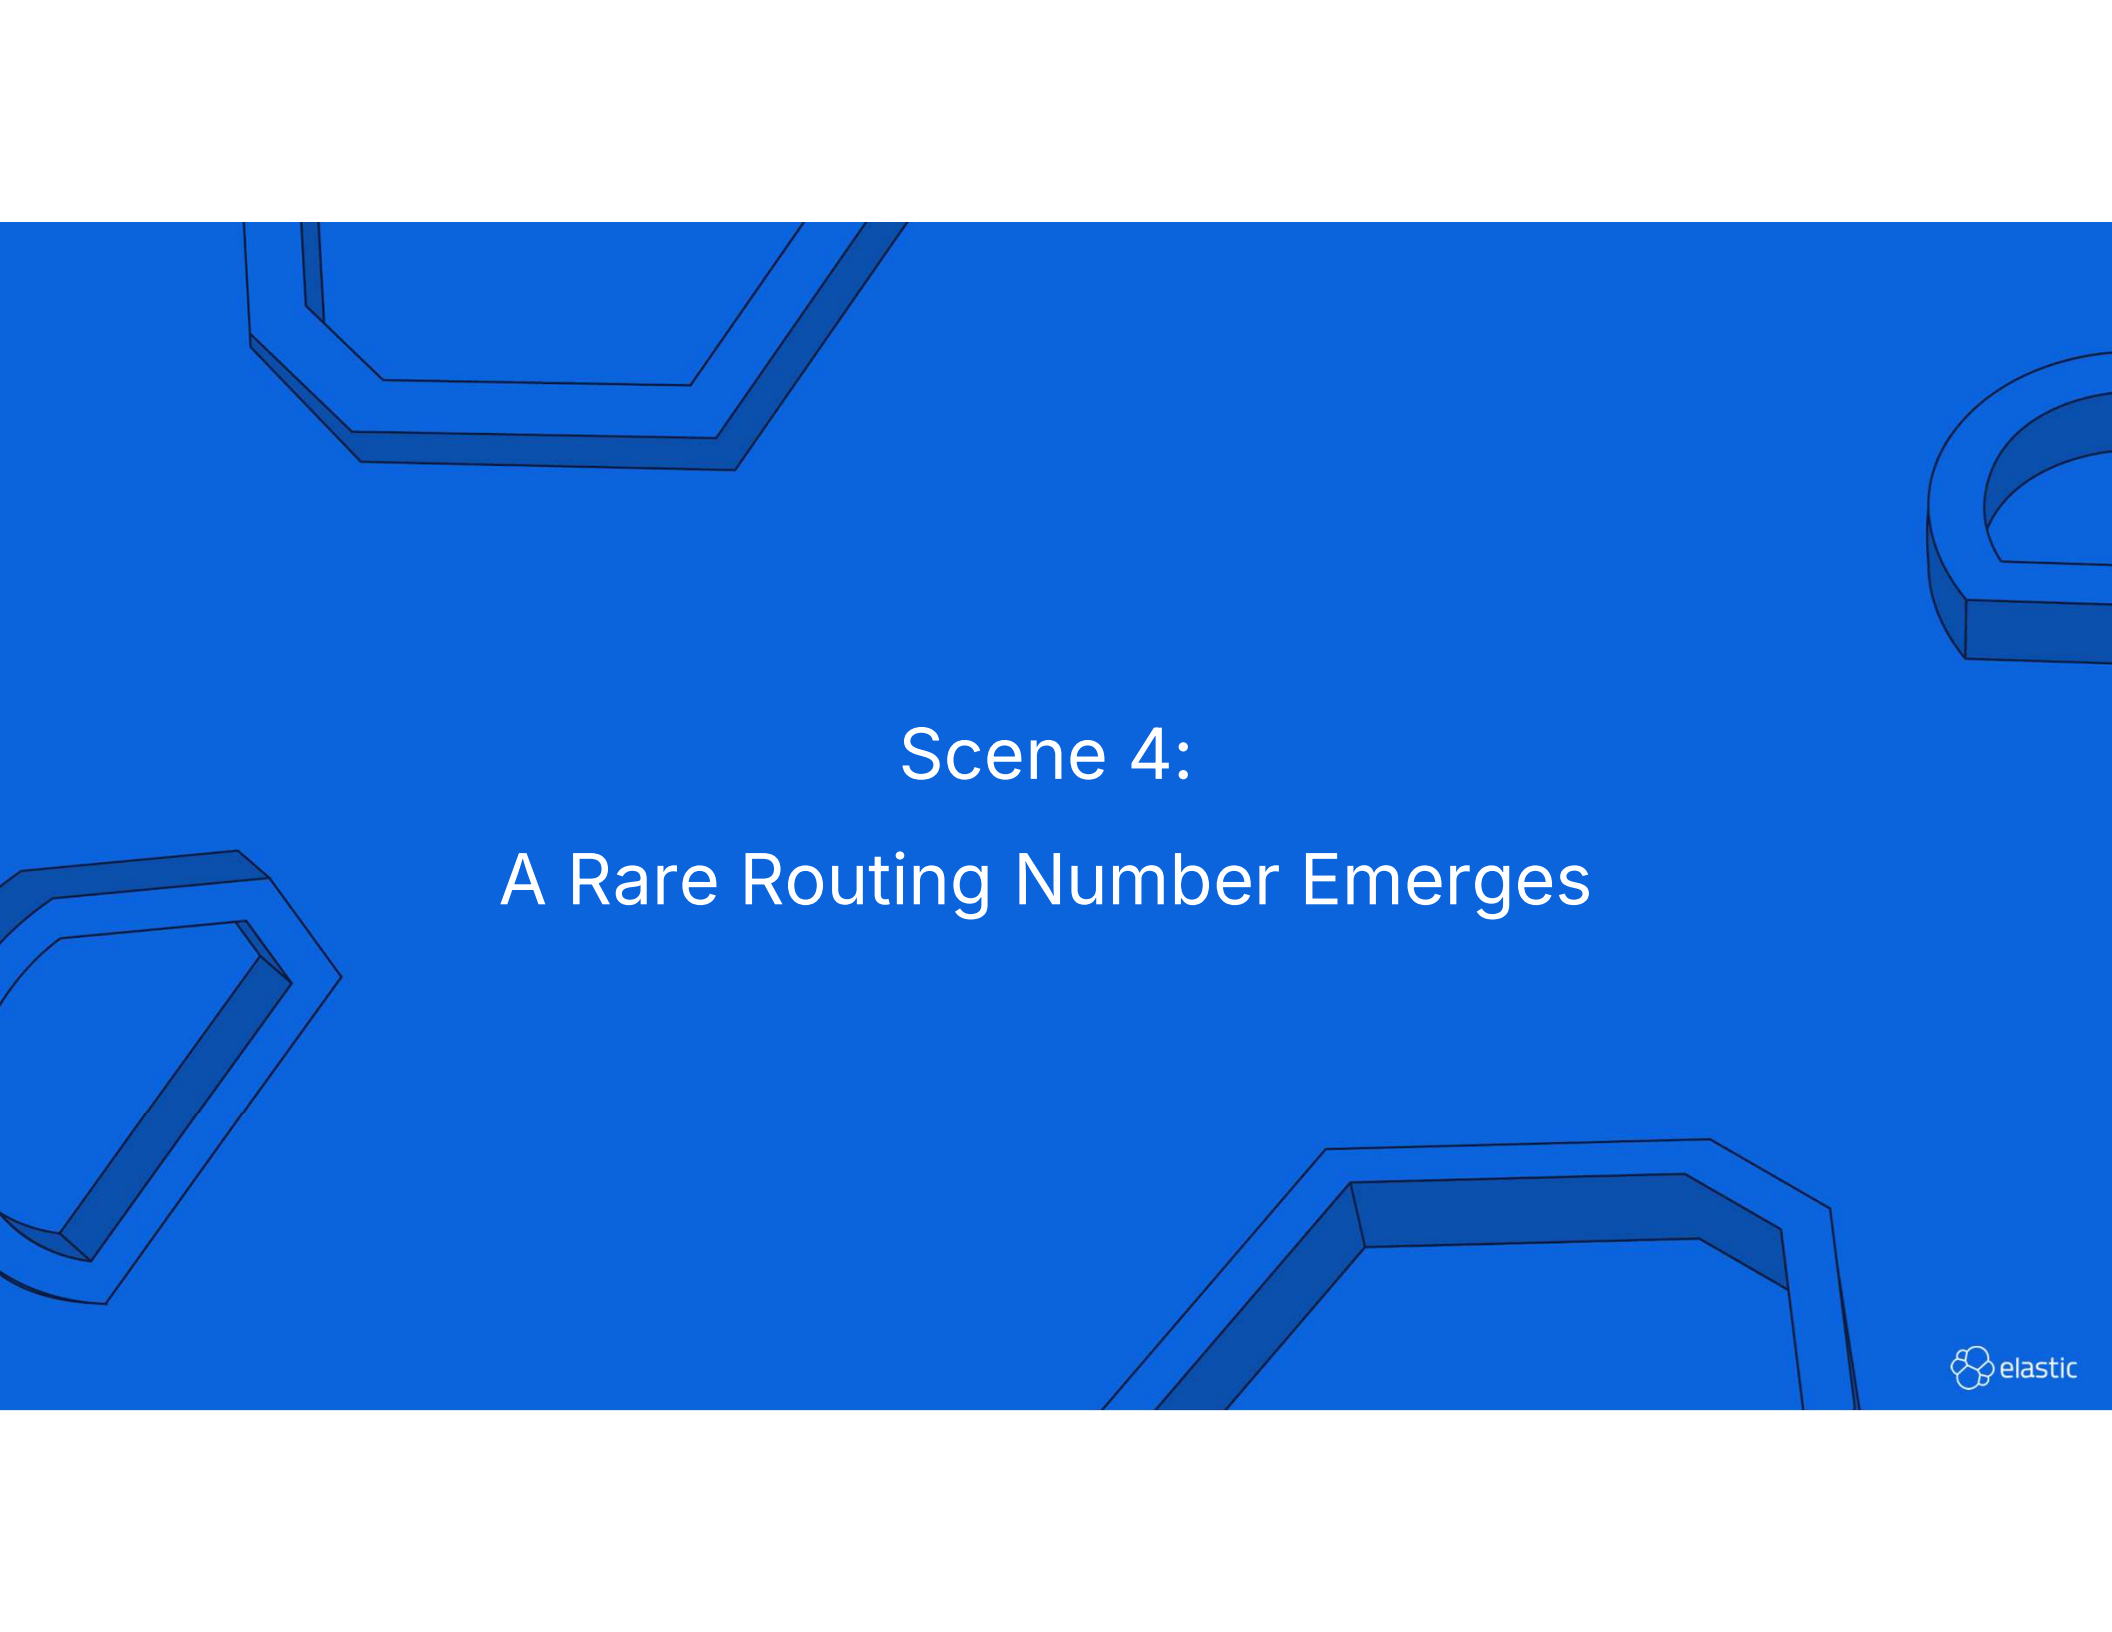
distinct_emails

[daniellehiggins@example.org, brianna47@example.org, james86@example.net, williamhurley@example.com, qcannon@example.com, pennywilliams@example.org, markmorris@example.org, riverastephanie@example.net, katherine71@example.net, willie48@example.org, matthewgates@example.org, coxalison@example.com, kennethmelton@example.com, jongross@example.com, parsonsmichael@example.net, cdavis@example.org, gilesdavid@example.org, xbrown@example.com, ehopkins@example.net, matthew46@example.net, stephentucker@example.com, asuarez@example.net, taylorbarbara@example.org, christopherpowell@example.com, hmoyer@example.com, stevenstimothy@example.com, deborah36@example.net, danadennis@example.org, thomasmartinez@example.net, manderson@example.com, benjamin43@example.com, douglaspatton@example.net, emily60@example.net, brichmond@example.net, cheryl76@example.com, edwincole@example.net, alexanderwheeler@example.com, kimmark@example.com, yolandaramsey@example.com, ubrady@example.net]

Filter for Filter out Copy value

distinct_ips

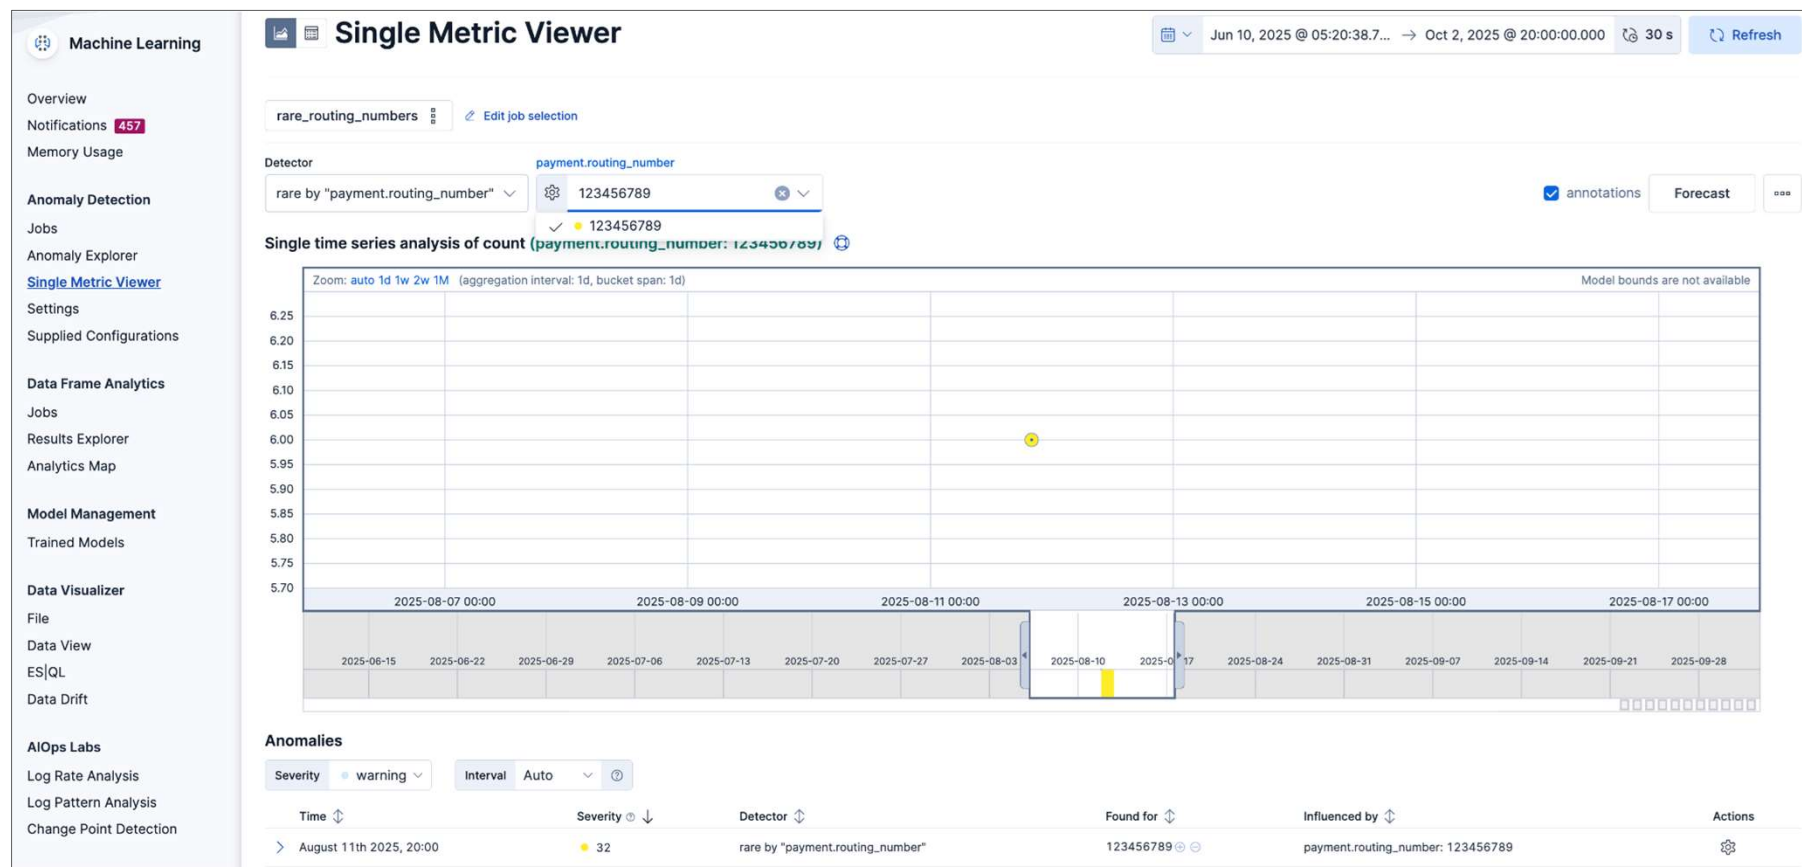
192.0.2.111

The background is a solid blue color. It features several white, three-dimensional geometric shapes that look like architectural moldings or frames. These shapes are positioned in the corners and along the edges, creating a sense of depth and structure. The shapes include octagons, rectangles with beveled edges, and a semi-circular arch-like structure.

Scene 4: A Rare Routing Number Emerges


Fraud Detection with Elastic

Scene 4: A Rare Routing Number Emerges



Fraud Detection with Elastic

Scene 4: A Rare Routing Number Emerges

 Security

Dashboards

Rules

Alerts

Attack discovery

Findings

Cases


Timelines

Intelligence

Explore


Rules

Management




Detection rules (SIEM)

Create and manage detection rules for threat detection and monitoring.




Benchmarks

View benchmark rules for Cloud Security Posture management.



Shared exception lists


Create and manage shared exception lists to prevent the creation of unwanted alerts.



SIEM Rule Migrations


Our generative AI powered SIEM migration tool automates some of the most time consuming migrations tasks and processed.

Discover



MITRE ATT&CK® Coverage

Review and maintain your protections MITRE ATT&CK® coverage.



Fraud Detection with Elastic

Scene 4: A Rare Routing Number Emerges

Security

Dashboards

Rules

Alerts

Attack discovery

Findings

Cases

Timelines

Intelligence

Explore

Rules

[Add Elastic rules](#) 1495 [Manage value lists](#) [Import rules](#) [Create new rule](#)

Get more value out of Elastic prebuilt rules!

Learn how to customize prebuilt rules and update them to get the latest improvements.

[Read the blog](#)

Installed Rules 4 Rule Monitoring 4

Tags 1 Last response 3 Elastic rules (0) Custom rules (4) Enabled rules Disabled rules

Showing 1-4 of 4 rules | Selected 0 rules [Select all 4 rules](#) [Bulk actions](#) [Refresh](#) Updated 10 seconds ago [On](#)

<input type="checkbox"/> Rule	Risk s...	Severity	Last run	Last response	Last updated	Notify	Enabled
<input type="checkbox"/> Three or More Bank Accounts Per Claim	1 70	Medium	3 minutes ago	Succeeded	Aug 12, 2025 @ 19:31:...		
<input type="checkbox"/> Three or More Claims Per Bank Account	1 90	High	3 minutes ago	Succeeded	Aug 12, 2025 @ 19:31:...		
<input type="checkbox"/> Three or More Emails Per IP Address	1 70	Medium	3 minutes ago	Succeeded	Aug 12, 2025 @ 19:31:...		
<input type="checkbox"/> Rare or Suspicious Routing Numbers	1 80	High	2 minutes ago	Succeeded	Aug 12, 2025 @ 19:31:...		


Rows per page: 20

Fraud Detection with Elastic


Scene 4: A Rare Routing Number Emerges

1 Define rule


Rule type

**Custom query**
Use KQL or Lucene to detect issues across indices.


Select

**Machine Learning**
Select ML job to detect anomalous activity.


✓ Selected

**Threshold**
Aggregate query results to detect when number of matches exceeds threshold.


Select

**Event Correlation**
Use Event Query Language (EQL) to match events, generate sequences, and stack data


Select

**Indicator Match**
Use indicators from intelligence sources to detect matching events and alerts.

Select

**New Terms**
Find documents with values appearing for the first time.

Select

**ES|QL**
Use Elasticsearch Query Language (ES|QL) to find events and aggregate search results.

Select

Machine Learning job


rare_routing_numbers ×

Create custom job

We've provided a few common jobs to get you started. To add your own custom jobs, assign a group of 'security' to those jobs in the [Machine Learning](#) application to make them appear here.

Anomaly score threshold

25



Fraud Detection with Elastic

Scene 4: A Rare Routing Number Emerges

2

About rule

Name

alert_on_rare_routing_number

Description

This detection rule is triggered when an anomaly is discovered by our ML job rare_routing_numbers

Default severity

Select a severity level for all alerts generated by this rule.

Critical

☐ Severity override

Use source event values to override the default severity.

Default risk score

Select a risk score for all alerts generated by this rule.

0255075100

99

☐ Risk score override

Use a source event value to override the default risk score.

Tags

Optional

fraud_detection ×

Type one or more custom identifying tags for this rule. Press enter after each tag to begin a new one.

> Advanced settings

Continue

Fraud Detection with Elastic

Scene 4: A Rare Routing Number Emerges

3

Schedule rule

Runs every

5

Minutes

Rules run periodically and detect alerts within the specified time frame.

Additional look-back time

1

Minutes

Adds time to the look-back period to prevent missed alerts.

Continue

Fraud Detection with Elastic


Scene 4: A Rare Routing Number Emerges


4


Rule actions


Actions


Select a connector type



Cases



D3 Security



Email



IBM Resilient



Index



Jira



Microsoft Teams



Opsgenie



PagerDuty



Server log



ServiceNow ITOM



ServiceNow ITSM



ServiceNow SecOps



Slack



Swimlane


TheHive


Tines



Torq



Webhook


xMatters

Response Actions

Response actions are run on each rule execution.


Osquery


Elastic Defend

Create rule without enabling it

Create & enable rule

The background is a solid blue color. It features several white, three-dimensional geometric shapes that look like stylized architectural elements or frames. These shapes are positioned in the corners and along the edges, creating a sense of depth and structure. The shapes include octagons, rectangles, and curved forms, all rendered with white outlines and some internal shading to give them a 3D appearance.

Scene 5: A Call from Medicare

Fraud Detection with Elastic

Scene 5: A Call from Medicare

Data Visualizer



Upload data from a file

Upload your file, analyze its data, and optionally import the data into an index. The data can also be automatically vectorized using semantic text.

The following file formats are supported up to 60 MB:

- PDF
- Microsoft Office Word, Excel, and PowerPoint
- Plain text and Rich Text Format
- Open Document Format

The following file formats are supported up to 500 MB:

- Delimited text files, such as CSV and TSV
- Newline-delimited JSON
- Log files with a common format for the timestamp

Upload data



Select or drag and drop a file

Fraud Detection with Elastic

Scene 5: A Call from Medicare

medicare-claims.csv [Select a different file](#)

File contents

First 501 lines

```
1  medicare_beneficiary_id,medicare_claim_id,medicare_claim_line,@timestamp,medicare_procedure_code,medicare_modifier1,medicare_modifier2,medicare_diagnosis_code,medicare_provider_npi,medicare_rendering_npi,
   medicare_claim_type,medicare_place_of_service,medicare_units,medicare_claim_amount,medicare_status,source.ip
2  B100308,C500308,2,2025-08-08T21:35:42Z,99213,,I25.10,9556315838,8746968888,Carrier (Part B),22,1,206.68,Paid,198.51.100.235
3  B100013,C500013,2,2025-08-07T13:26:17Z,96372,59,,E11.9,1496088892,,Outpatient (Part B),22,4,213.73,Denied,203.0.113.200
4  B100414,C500414,1,2025-08-10T17:22:20Z,90658,,M54.5,4882535309,6046273051,Carrier (Part B),02,1,287.93,Paid,203.0.113.34
5  B100032,C500032,1,2025-04-22T16:38:54Z,93000,RT,,Z20.828,7256602822,3380205275,Carrier (Part B),11,4,160.81,Paid,203.0.113.82
6  B100460,C500460,3,2025-08-07T15:02:48Z,90658,,R53.1,6942421836,7157126938,Inpatient (Part A),22,1,61.49,Adjusted,203.0.113.224
7  B100155,C500155,3,2025-04-11T02:23:26Z,99212,,N39.0,3611789641,4414211530,Carrier (Part B),11,2,141.96,Denied,203.0.113.243
8  B100249,C500249,1,2025-08-10T20:18:30Z,36415,26,,R51.9,6789071802,2059546535,Outpatient (Part B),11,1,232.23,Paid,203.0.113.53
9  B100017,C500017,2,2025-08-08T00:14:10Z,90658,EN,C70.5,4713757500,6132640367,Carrier (Part B),11,2,204.87,Paid,203.0.113.106
```

Summary

Number of lines analyzed	501
Format	delimited
Delimiter	,
Has header row	true
Time field	@timestamp
Time format	ISO8601

[Import](#) [Override settings](#) [Analysis explanation](#)

File stats

All fields **16** of 16 total Number fields **6** of 6 total

Field name **16** ▼ Field type **4** ▼ ⓘ

>	Type ↕	Name ↑	Documents (%) ↕	Distinct values ↕	Distributions ⚙
>	🕒	@timestamp	500 (100%)	500	
>	🔑	medicare_beneficiary_id	500 (100%)	500	<div>██████████ top 10 of 500 categories</div>

Fraud Detection with Elastic

Scene 5: A Call from Medicare

File stats					
All fields 16 of 16 total		Number fields 6 of 6 total		Field name 16	Field type 4
Type	Name	Documents (%)	Distinct values	Distributions	
>	@timestamp	500 (100%)	500		
>	medicare_beneficiary_id	500 (100%)	500		
>	medicare_claim_amount	500 (100%)	483	min	median
				20	198.67
>	medicare_claim_id	500 (100%)	500		
>	medicare_claim_line	500 (100%)	3	min	median
				1	2
>	medicare_claim_type	500 (100%)	3		
>	medicare_diagnosis_code	500 (100%)	14		
>	medicare_modifier1	233 (46.6%)	6		
>	medicare_modifier2	51 (10.2%)	6		
>	medicare_place_of_service	500 (100%)	6	min	median
				2	11
>	medicare_procedure_code	500 (100%)	14		
>	medicare_provider_npi	500 (100%)	500	min	median
				1001976602	5326097620
>	medicare_rendering_npi	358 (71.6%)	358	min	median
				1004549232	5446608883
>	medicare_status	500 (100%)	3		
>	medicare_units	500 (100%)	4	min	median
				1	3
>	source_ip	500 (100%)	328		

Fraud Detection with Elastic

Scene 5: A Call from Medicare

medicare-claims.csv

Import data
[Simple](#) [Advanced](#)

Index name

☐ Create data view

Data view name

Automatically created fields
[Add additional field](#)

Index settings

```
1 {
2   "index.mode": "lookup"
3 }
```

Mappings

```
39 {
40   "medicare_rendering_npi": {
41     "type": "long"
42   },
43   "medicare_status": {
44     "type": "keyword"
45   },
46   "medicare_units": {
47     "type": "long"
48   },
49   "source_ip": {
50     "type": "ip"
51   }
52 }
```

Ingest pipeline

```
1 {
2   "description": "Ingest pipeline created by text
3   structure finder",
4   "processors": [
5     {
6       "csv": {
7         "field": "message",
8         "target_fields": [
9           "medicare_beneficiary_id",
10          "medicare_claim_id",
11          "medicare_claim_line",
12          "@timestamp",
13          "medicare_procedure_code",
14          "medicare_modifier1",
```

[Import](#)

[Back](#)

[Select a different file](#)

Fraud Detection with Elastic

Scene 5: A Call from Medicare

ES|QL help

Jun 10, 2025 @ 05:20:38.7... → Aug 12, 2025 @ 15:31:13.528 30 s Refresh

```
1 FROM "unemployment-claims"
2 | LOOKUP JOIN "medicare-claims" ON source.ip
3 | KEEP source.ip, claim.id, claim.amount, medicare_claim_id, medicare_claim_amount
4 | WHERE medicare_claim_id IS NOT NULL
5 | STATS distinct_ip_count = COUNT(source.ip) BY source.ip
6 | KEEP source.ip
7
```

7 lines @timestamp found LIMIT 1000 rows

Submit feedback Show recent queries

Search field names 0

1 result

Columns 1 Sort fields

Selected fields 1

source.ip

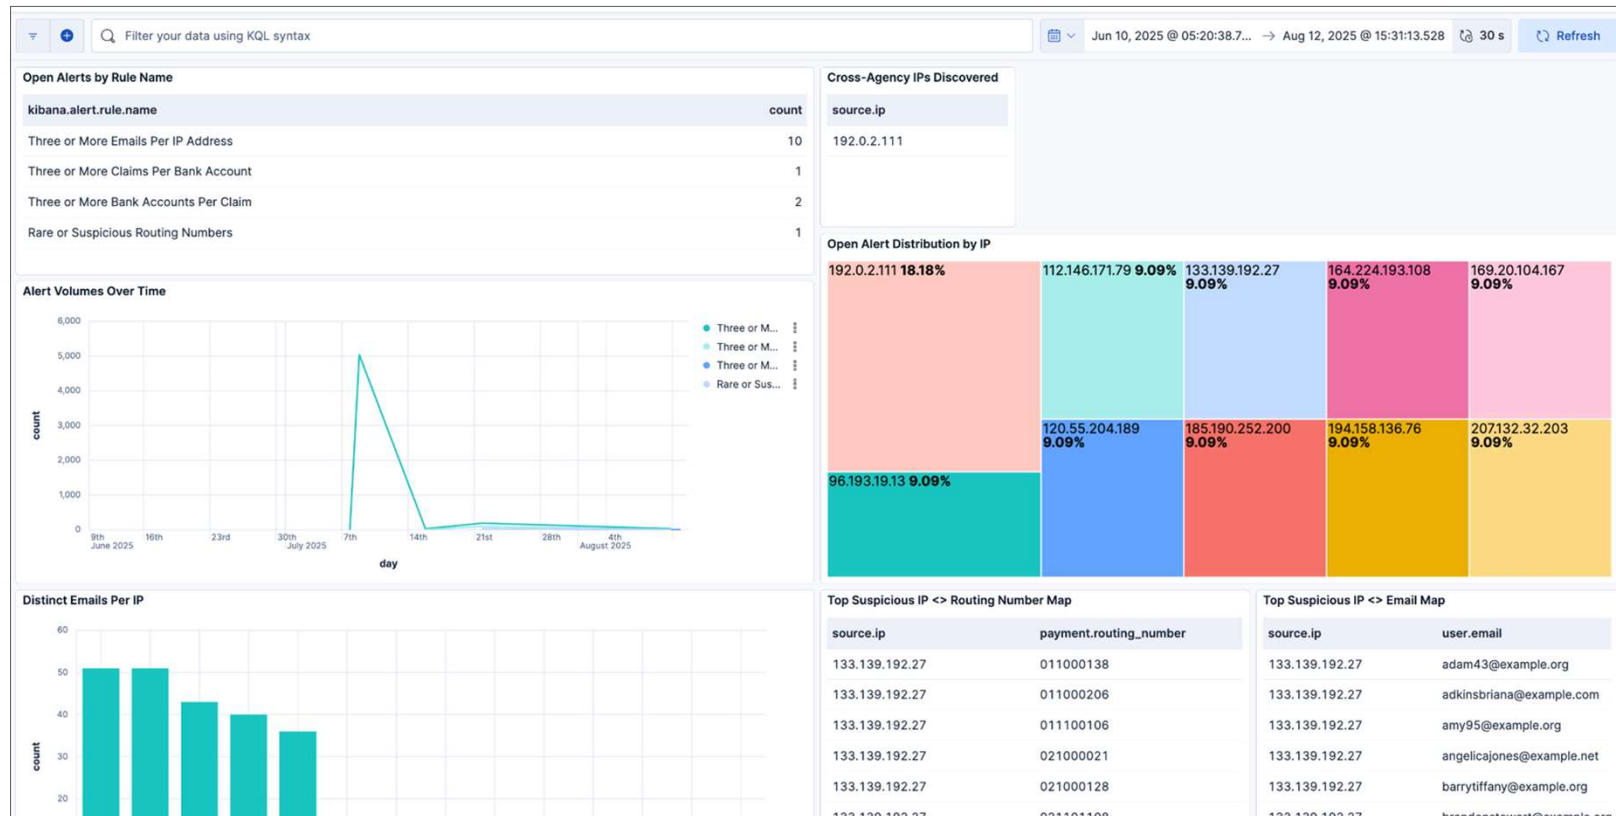
Available fields 1

source.ip

source.ip
192.0.2.111

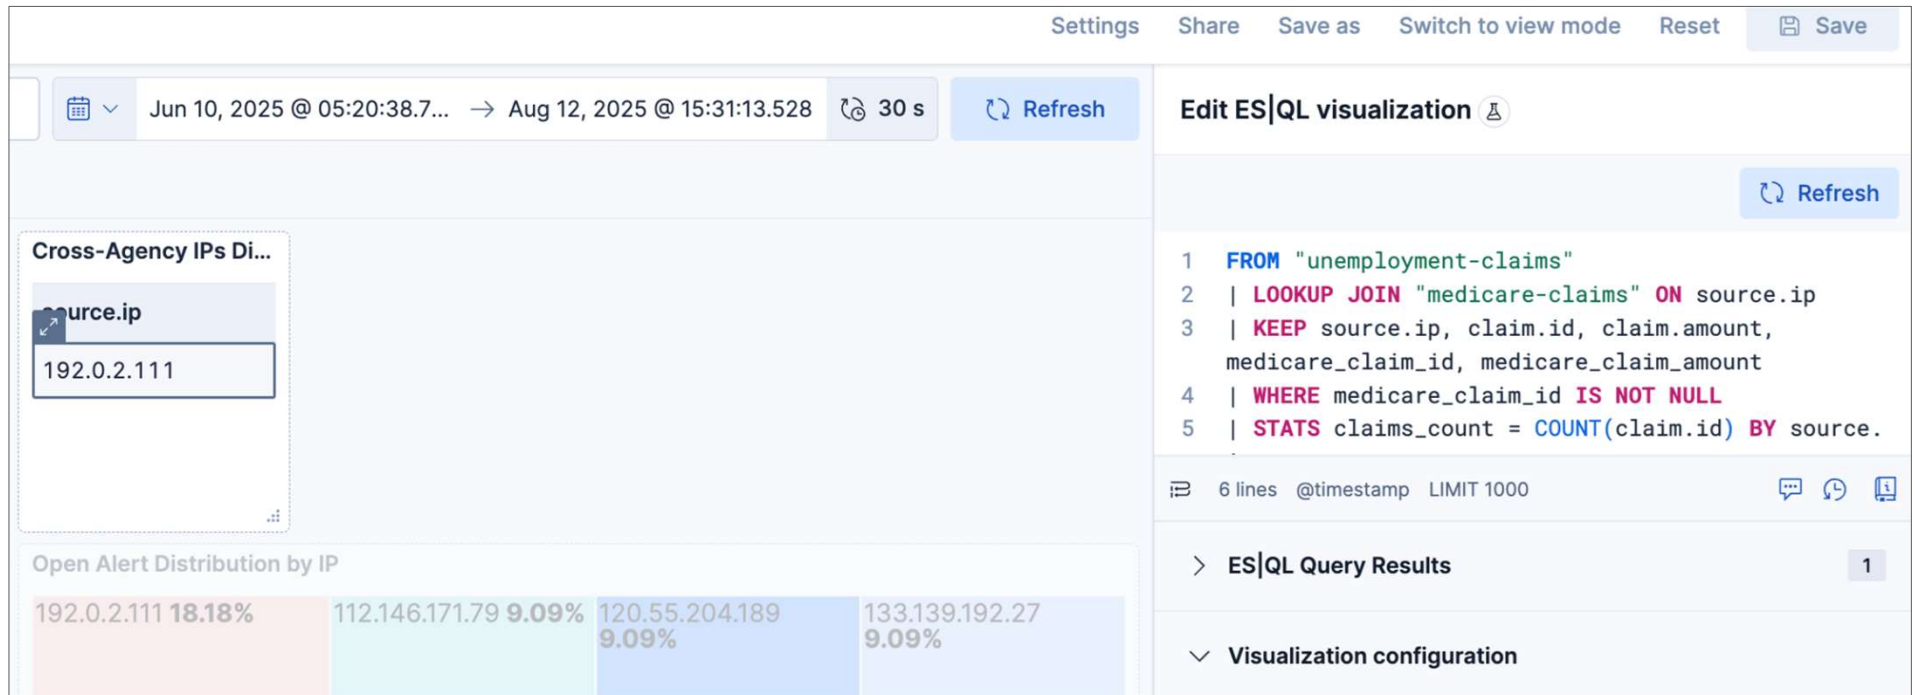
Fraud Detection with Elastic

Scene 6: Reporting Upstream



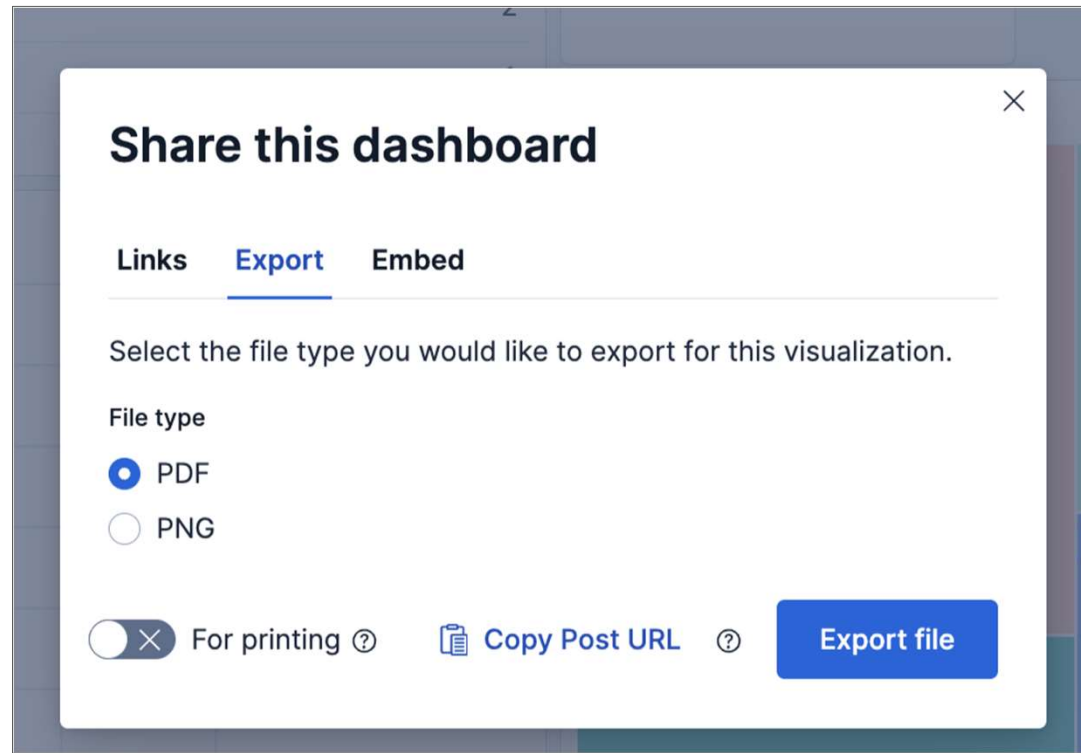
Fraud Detection with Elastic


Scene 6: Reporting Upstream



Fraud Detection with Elastic

Scene 6: Reporting Upstream



The background is a solid blue color. It features several white, 3D-style geometric shapes that look like architectural moldings or frames. These shapes are positioned in the corners and along the edges, creating a modern, abstract design.

Scene 7:

Attack Discovery for Fraud Detection? Can it really work?


Fraud Detection with Elastic

Scene 7: Attack Discovery for Fraud Detection? Can it really work?

The screenshot displays the Elastic Security console. On the left is a navigation sidebar with the following items: Security, Dashboards, Rules, Alerts, Attack discovery (highlighted), Findings, Cases, Timelines, Intelligence, and Explore. The main content area is titled 'Attack discovery' and includes a sub-header '1 discovery | 33 alerts | Generated: 2 days ago'. A top right corner contains the text 'KR-Personal-OpenAI' and a 'Generate' button. The primary alert is 'Coordinated Fraud Operation Detected', which includes an 'Attack chain' visualization (a sequence of red and white circles) and a count of 'Alerts: 33'. Below the alert title, a summary states: 'Fraud network involving 192.0.2.111 and 207.132.32.203'. A 'View in AI Assistant' link is provided. The alert details are organized into sections: 'Summary' (repeating the IP information), 'Details' (a paragraph describing a sophisticated fraud operation involving multiple tactics and IP addresses), and 'Attack Chain' (a horizontal timeline with stages: Reconnaissance, Resource Development, Initial Access, Execution, Persistence, Privilege Escalation, Defense Evasion, Credential Access, Discovery, Lateral Movement, and Collection). The 'Attack Chain' section also includes a 'View in AI Assistant' link and an 'Investigate in Timeline' button.

Fraud Detection with Elastic

Scene 7: Attack Discovery for Fraud Detection? Can it really work?

 **Security AI settings**

[Conversations](#) [Connectors](#) [System Prompts](#) [Quick Prompts](#) [Anonymization](#) [Knowledge Base](#)

Define privacy settings for event data sent to third-party LLM providers. You can choose which fields to include, and which to anonymize by replacing their values with random strings. Helpful defaults are provided below.

102
Allowed

3
Anonymized

102
Available

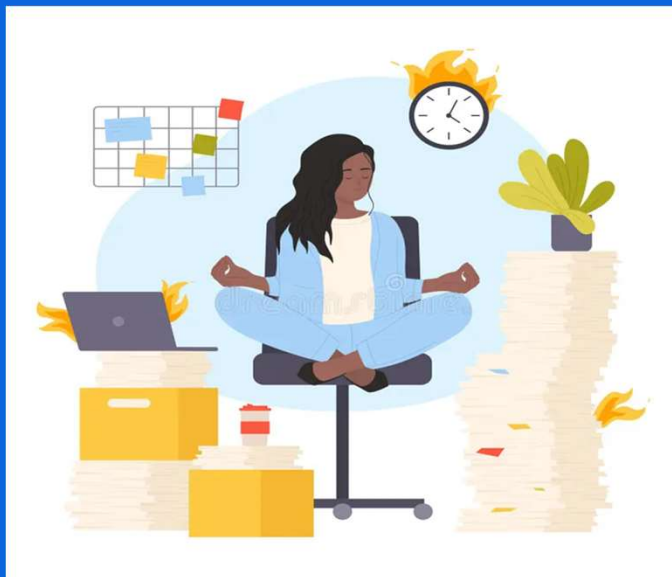
AllowedAnonymized

Selected 0 fields [Select all 102 fields](#) [Bulk actions](#)

<input type="checkbox"/>	Allowed <input type="checkbox"/>	Anonymized <input type="checkbox"/>	Field <input type="checkbox"/>
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	@timestamp
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	Ransomware.feature
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	Ransomware.files.data
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	Ransomware.files.entropy
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	Ransomware.files.extension
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	Ransomware.files.metrics
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	Ransomware.files.operation
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	Ransomware.files.path
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	Ransomware.files.score
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="radio"/> No	Ransomware.version

Rows per page: 10

[<](#) [1](#) [2](#) [3](#) [4](#) [5](#) ... [11](#) [>](#)



With Elastic for Fraud Detection

- 🎯 Detect fraud faster, before it becomes payout
- 🎯 Correlate patterns across IPs, emails, bank accounts, and locations
- 🎯 Scale discovery and analysis without scaling headcount
- 🎯 Provide explainable, defensible insights to leadership and investigators
- 🎯 Utilize AI to reduce the time and effort required to correlate alerts

Strengthen Fraud Prevention. Strengthen Trust.

Detect smartly

Use data and analytics to find fraud faster.

Govern responsibly

Protect citizen data with strong controls.

Build trust continuously

Prove accountability and transparency every step of the way.

Fraud Detection Value

-

Data Risk Exposure

+

Governance Maturity

=

Public Trust

What's the next step your agency can take to balance fraud prevention with citizen trust?

Thank You
Questions?



Kyle Rozanitis

kyle.rozanitis@elastic.co

<https://www.linkedin.com/in/kyle-rozanitis/>

Learn more at elastic.co