Beyond Detection:
AI's Potential For Supporting Threat Hunters

Robin Sommer

Corelight, Inc.

Co-Founder

`robin@corelight.com`

About me

PADERBORN
UNIVERSITY

SAARLAND
UNIVERSITY

TUM
TECHNISCHE
UNIVERSITÄT
MÜNCHEN

ICSI
INTERNATIONAL
COMPUTER SCIENCE
INSTITUTE

BERKELEY LAB
Lawrence Berkeley National Laboratory

corelight

An Open Source Network Security Monitoring Tool

Zeek (formerly Bro) is the world's leading platform for network security monitoring.

Flexible, open source, and powered by defenders.

ZEEK AND YE SHALL FIND

www.zeek.org

corelight

We transform network activity into evidence so that data-first defenders can stay ahead of ever-changing attacks.

FOUNDED	In 2013 in Berkeley, CA
LOCATIONS	San Francisco, CA (HQ); Columbus, OH; London, UK; Sydney, Australia; Dubai, UAE
PRODUCT	Open Network Detection and Response (NDR) platform for visibility, incident response, and threat hunting
CUSTOMERS	Fortune 500, critical infrastructure, national security, R&D
FUNDING	Series A/B/C/D (incl. Accel, Insight, Crowdstrike)
PEOPLE	~300

Remember SolarWinds in 2020?

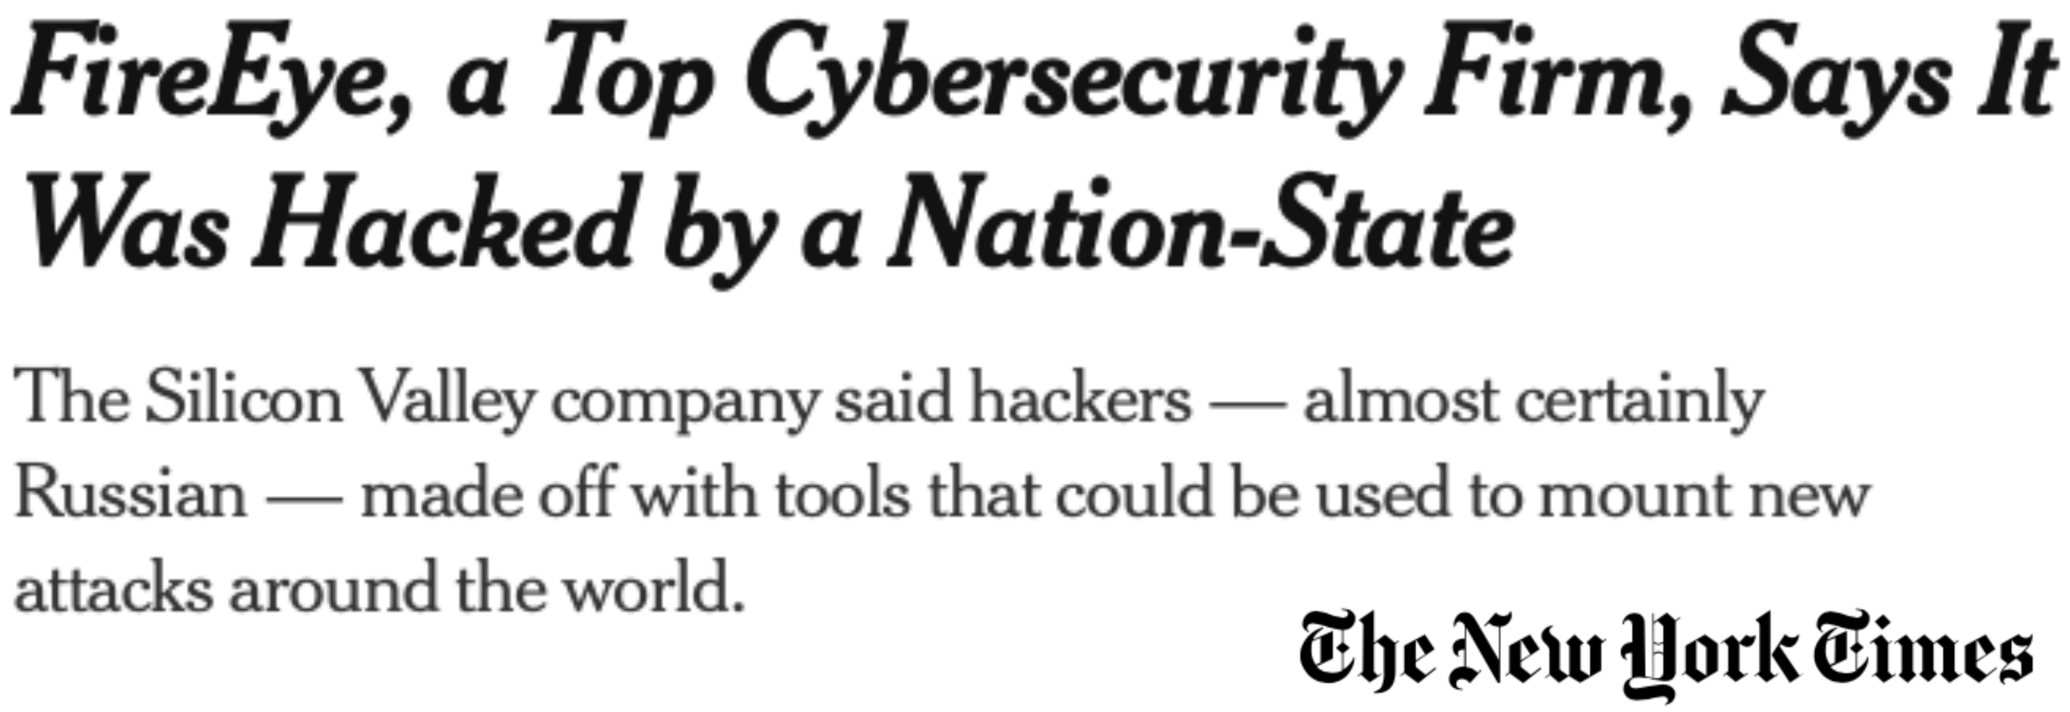

FireEye, a Top Cybersecurity Firm, Says It Was Hacked by a Nation-State

The Silicon Valley company said hackers — almost certainly Russian — made off with tools that could be used to mount new attacks around the world.

The New York Times

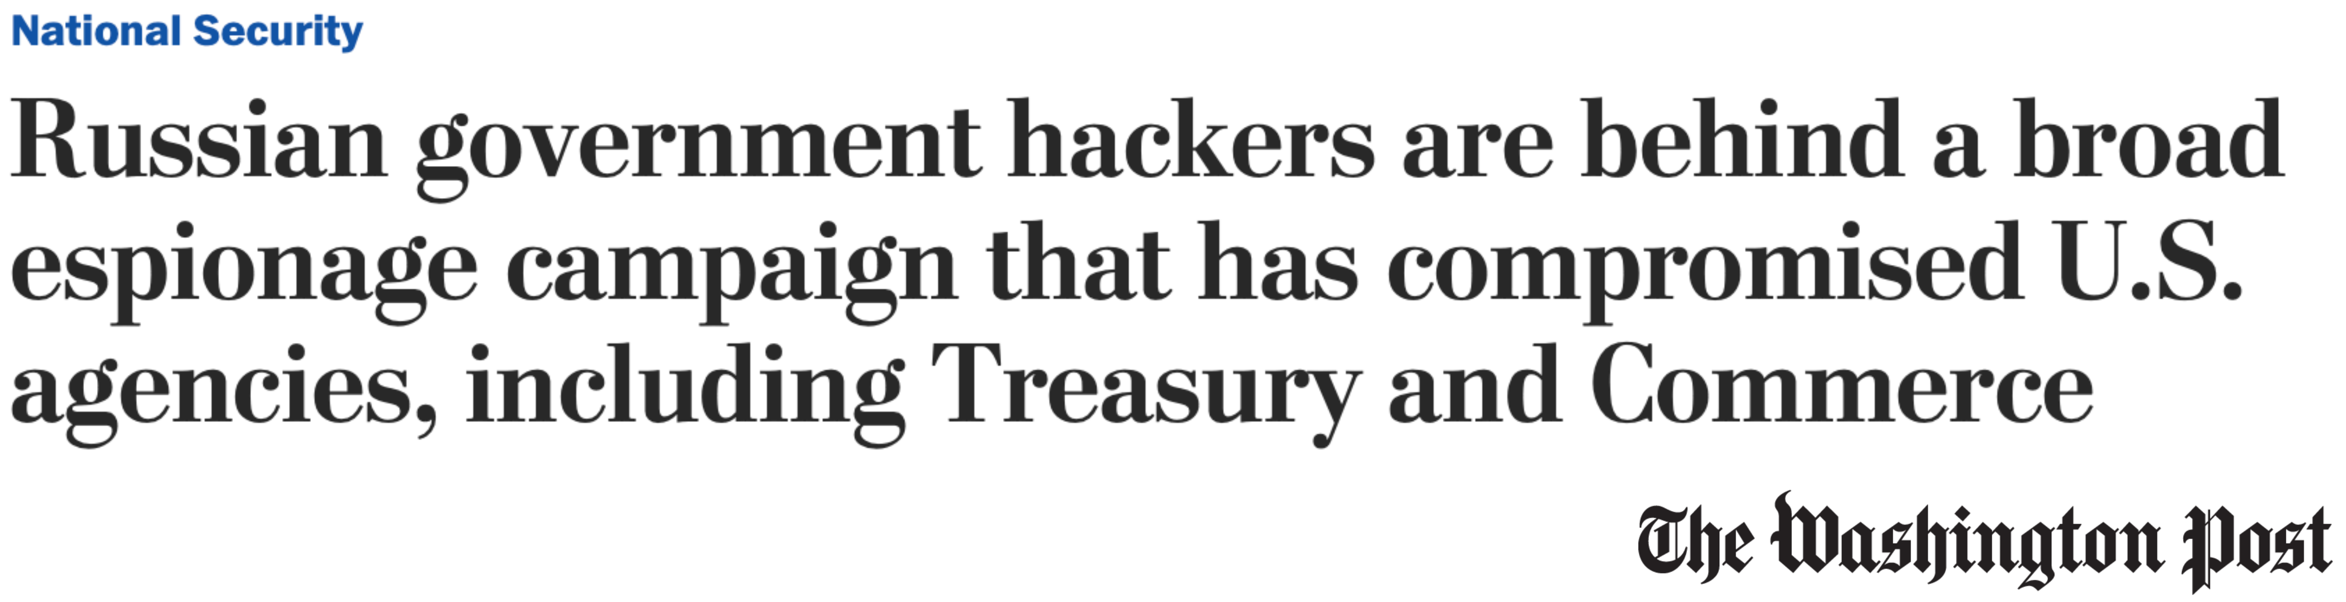

National Security

Russian government hackers are behind a broad espionage campaign that has compromised U.S. agencies, including Treasury and Commerce

The Washington Post

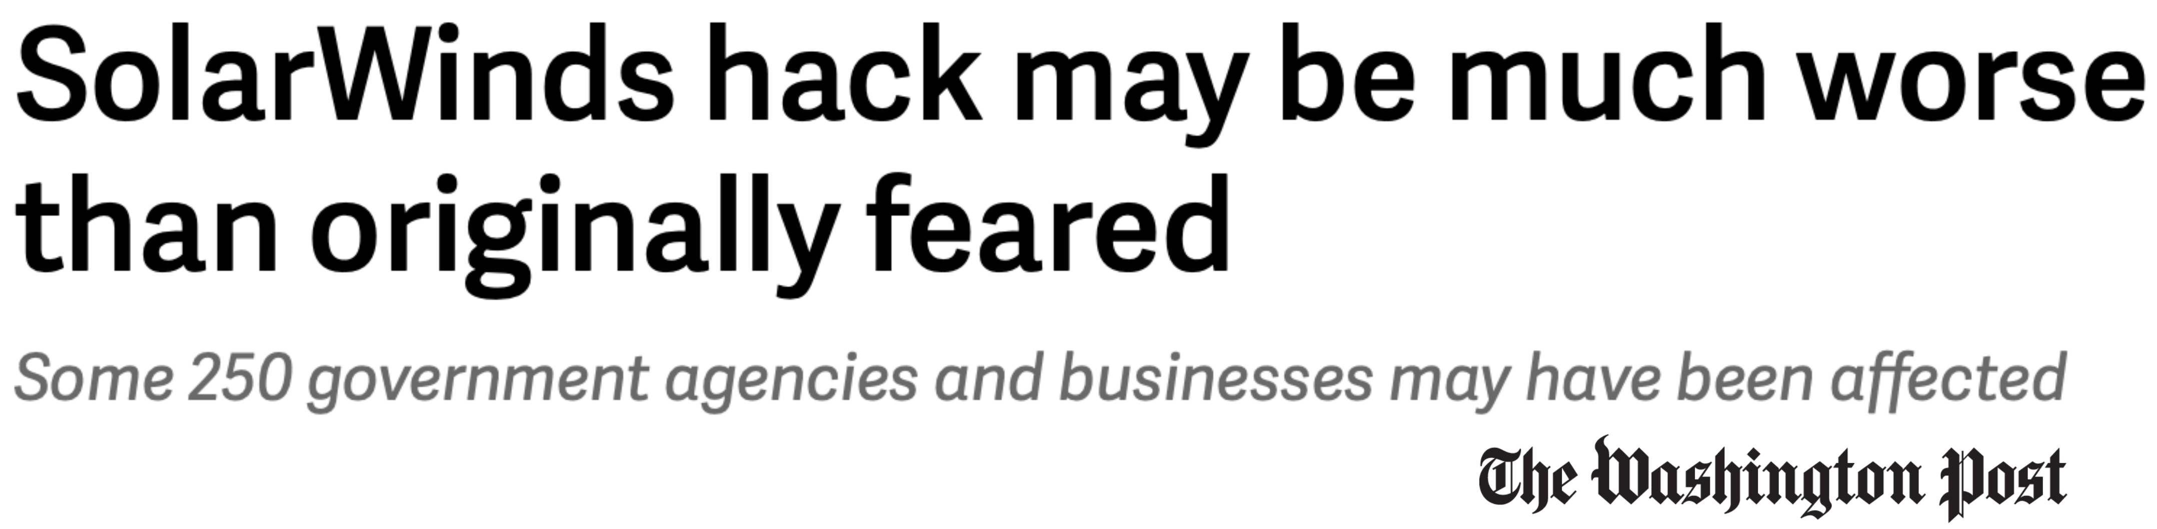

SolarWinds hack may be much worse than originally feared

Some 250 government agencies and businesses may have been affected

The Washington Post

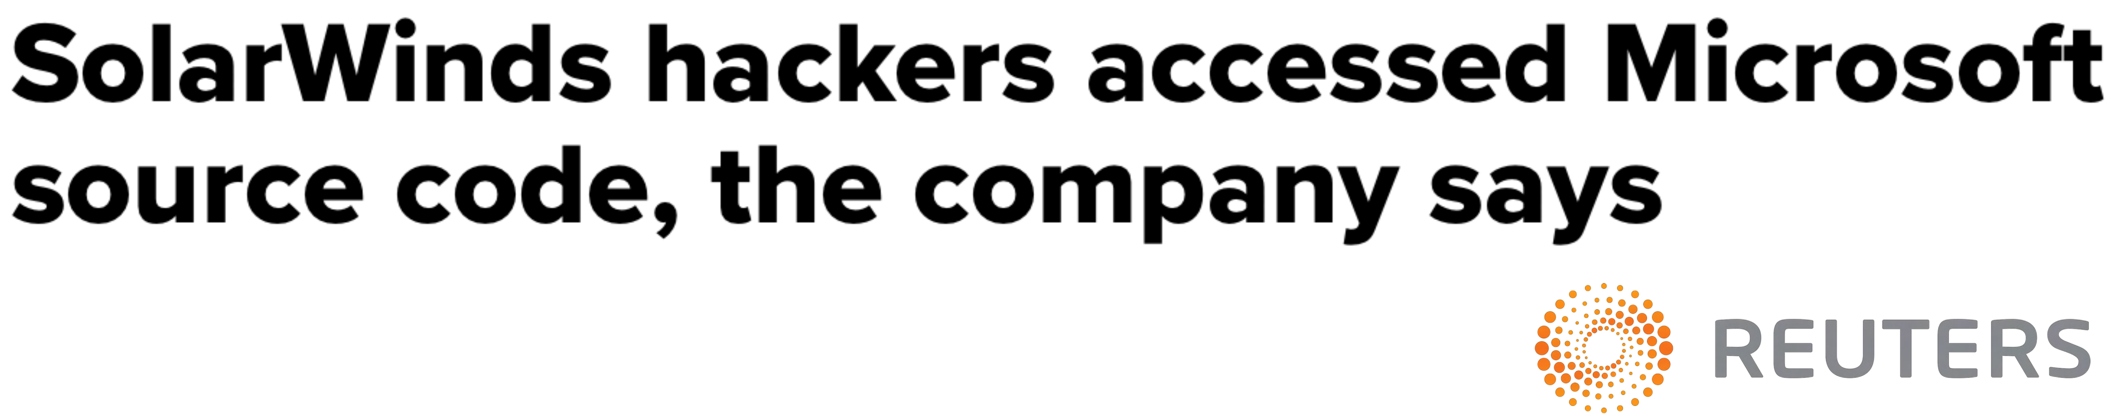

SolarWinds hackers accessed Microsoft source code, the company says

REUTERS

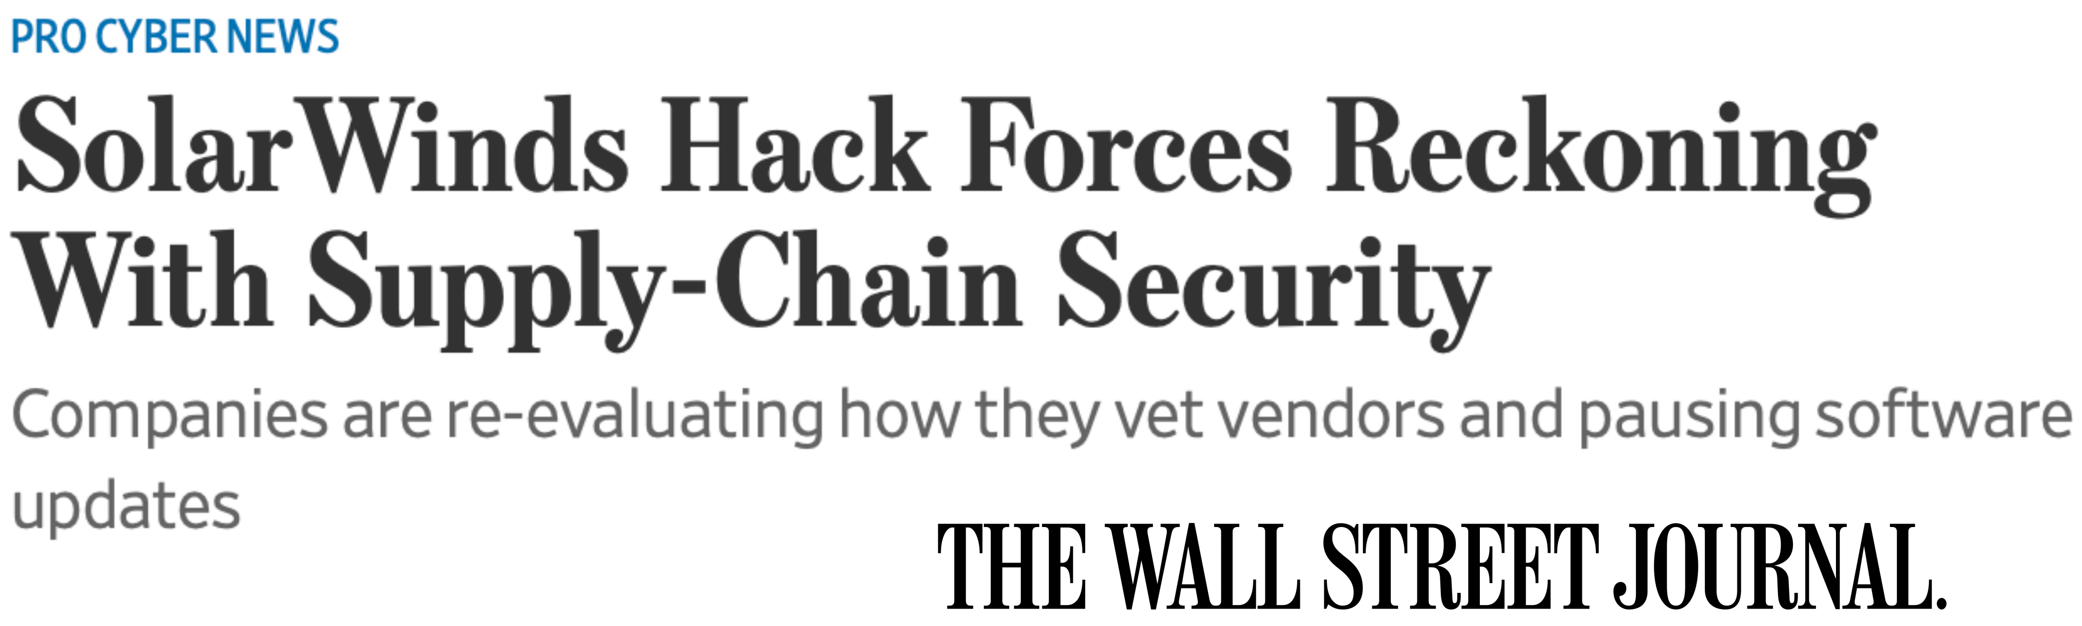

PRO CYBER NEWS

SolarWinds Hack Forces Reckoning With Supply-Chain Security

Companies are re-evaluating how they vet vendors and pausing software updates

THE WALL STREET JOURNAL.

Why wasn’t this detected earlier?

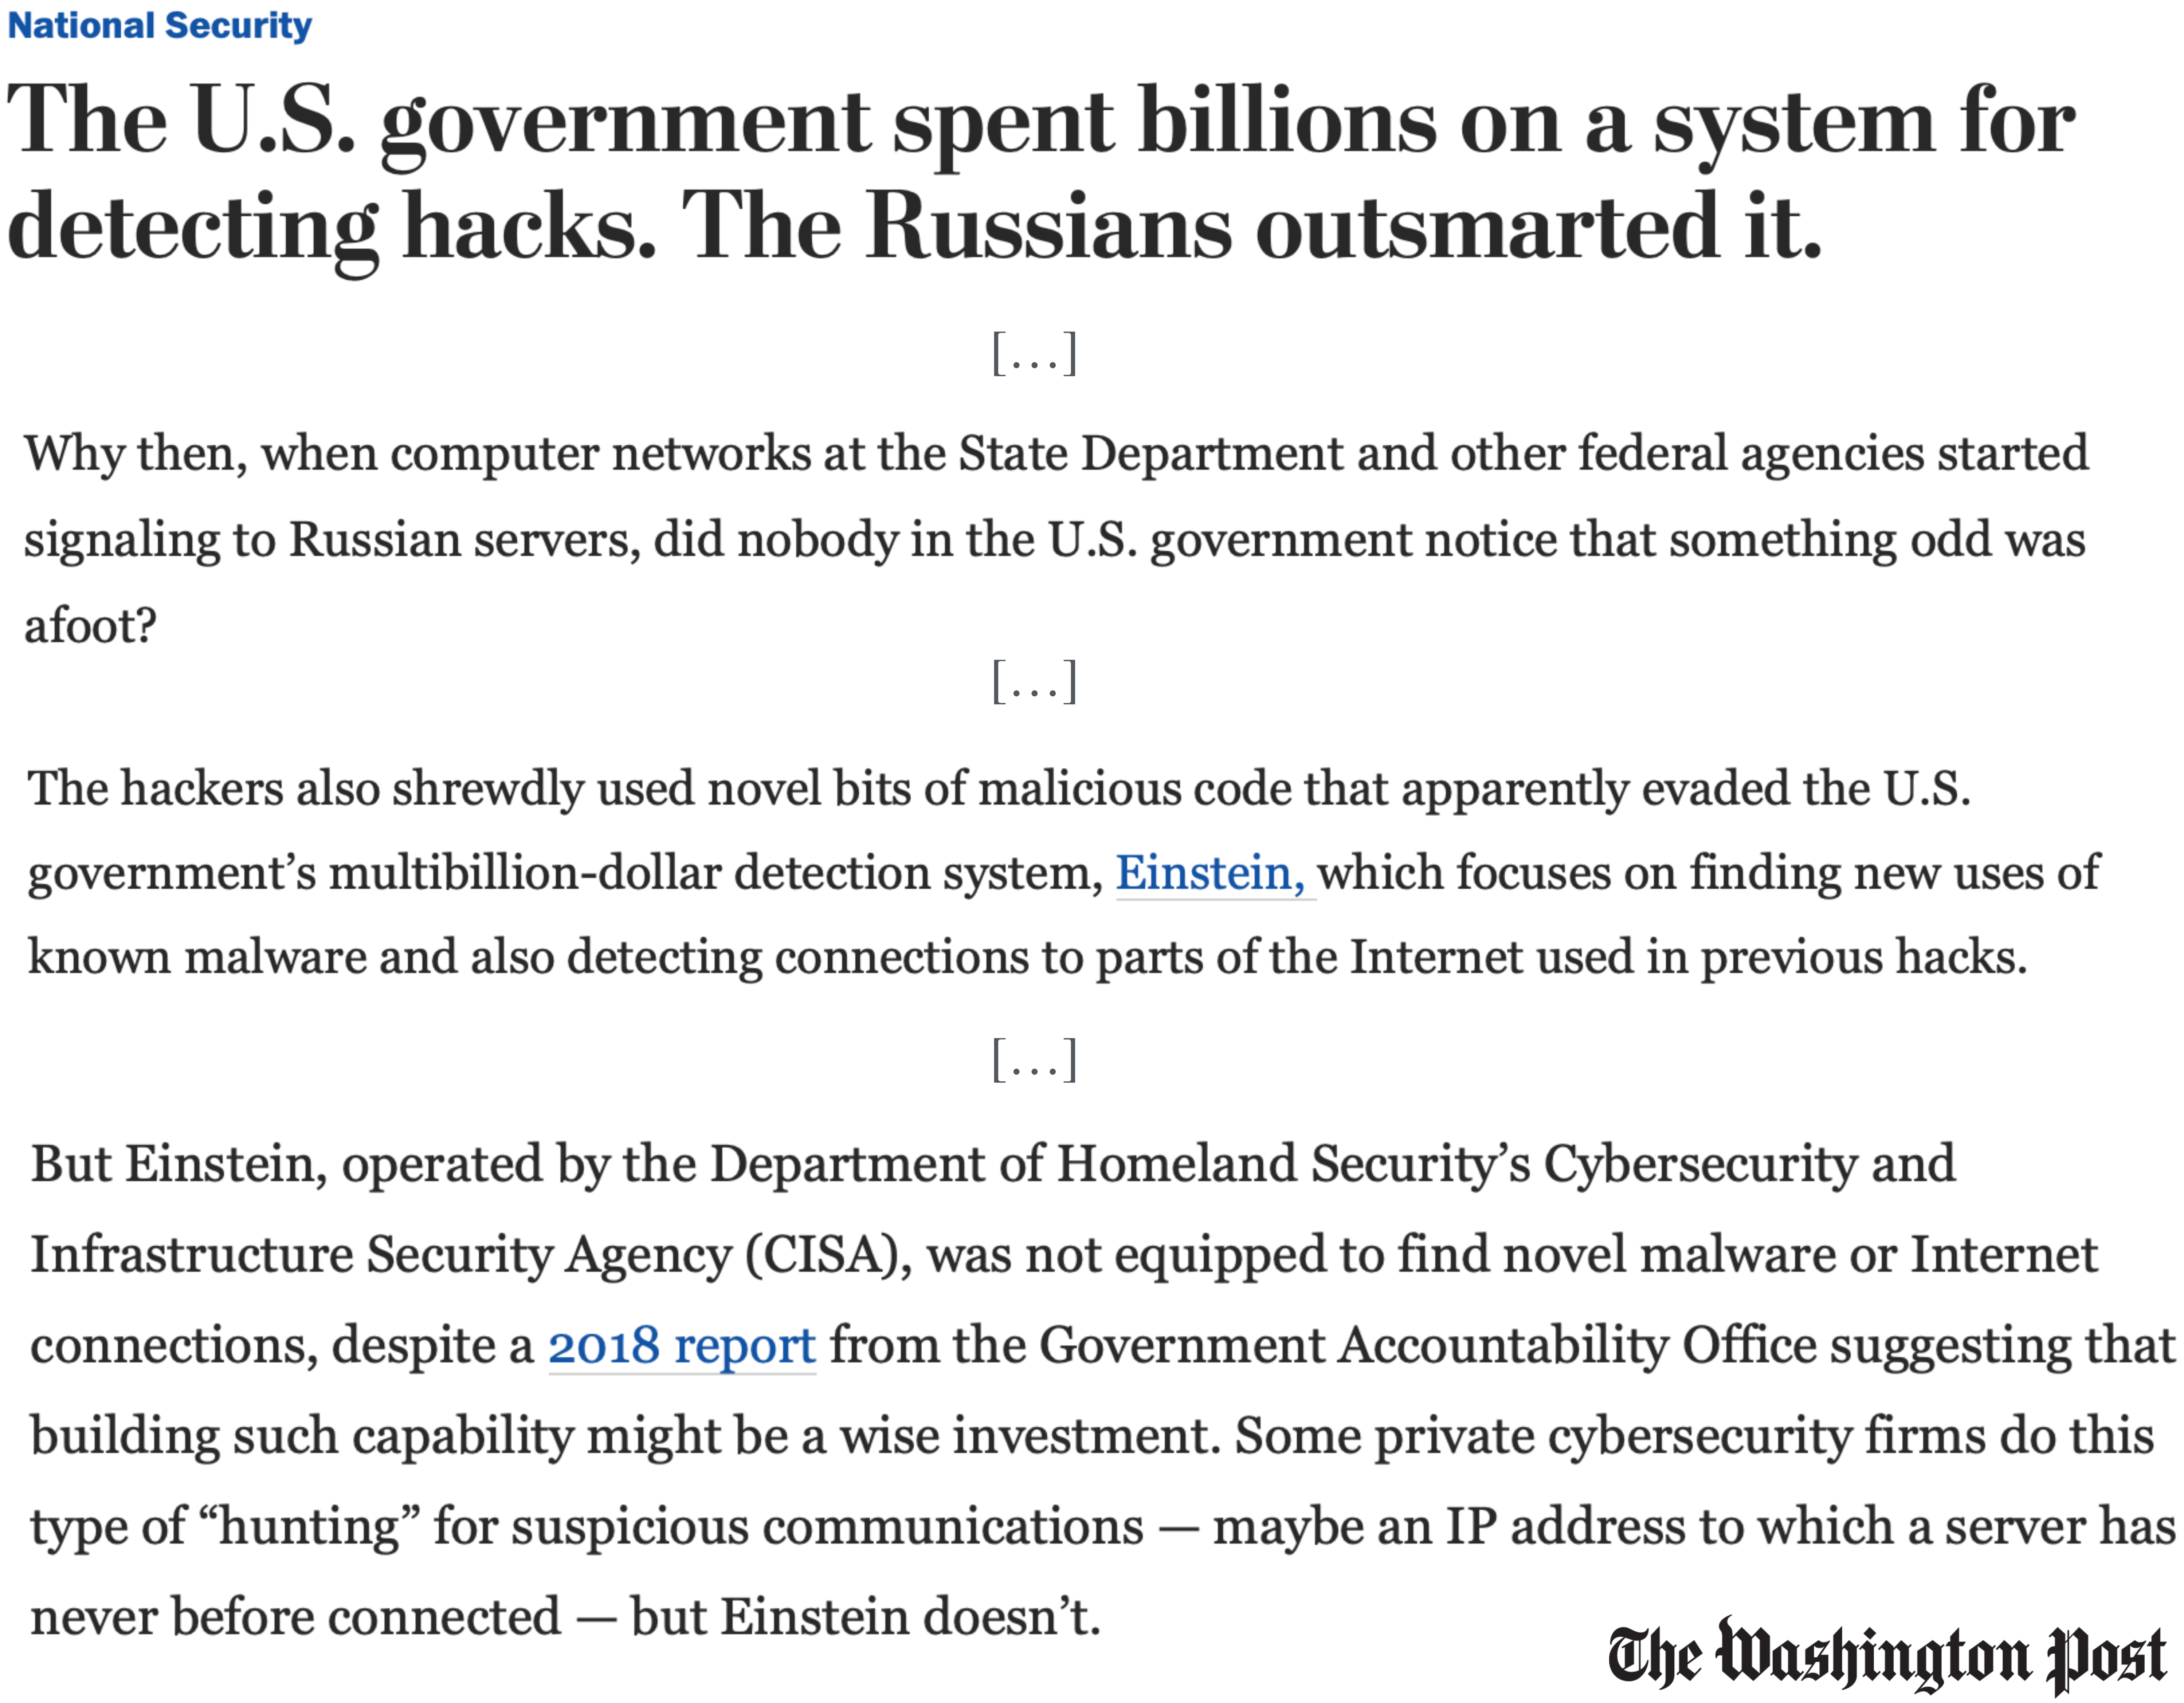

National Security

The U.S. government spent billions on a system for detecting hacks. The Russians outsmarted it.

[…]

Why then, when computer networks at the State Department and other federal agencies started signaling to Russian servers, did nobody in the U.S. government notice that something odd was afoot?

[…]

The hackers also shrewdly used novel bits of malicious code that apparently evaded the U.S. government’s multibillion-dollar detection system, Einstein, which focuses on finding new uses of known malware and also detecting connections to parts of the Internet used in previous hacks.

[…]

But Einstein, operated by the Department of Homeland Security’s Cybersecurity and Infrastructure Security Agency (CISA), was not equipped to find novel malware or Internet connections, despite a 2018 report from the Government Accountability Office suggesting that building such capability might be a wise investment. Some private cybersecurity firms do this type of “hunting” for suspicious communications — maybe an IP address to which a server has never before connected — but Einstein doesn’t.

The Washington Post

Why wasn't this detected earlier?

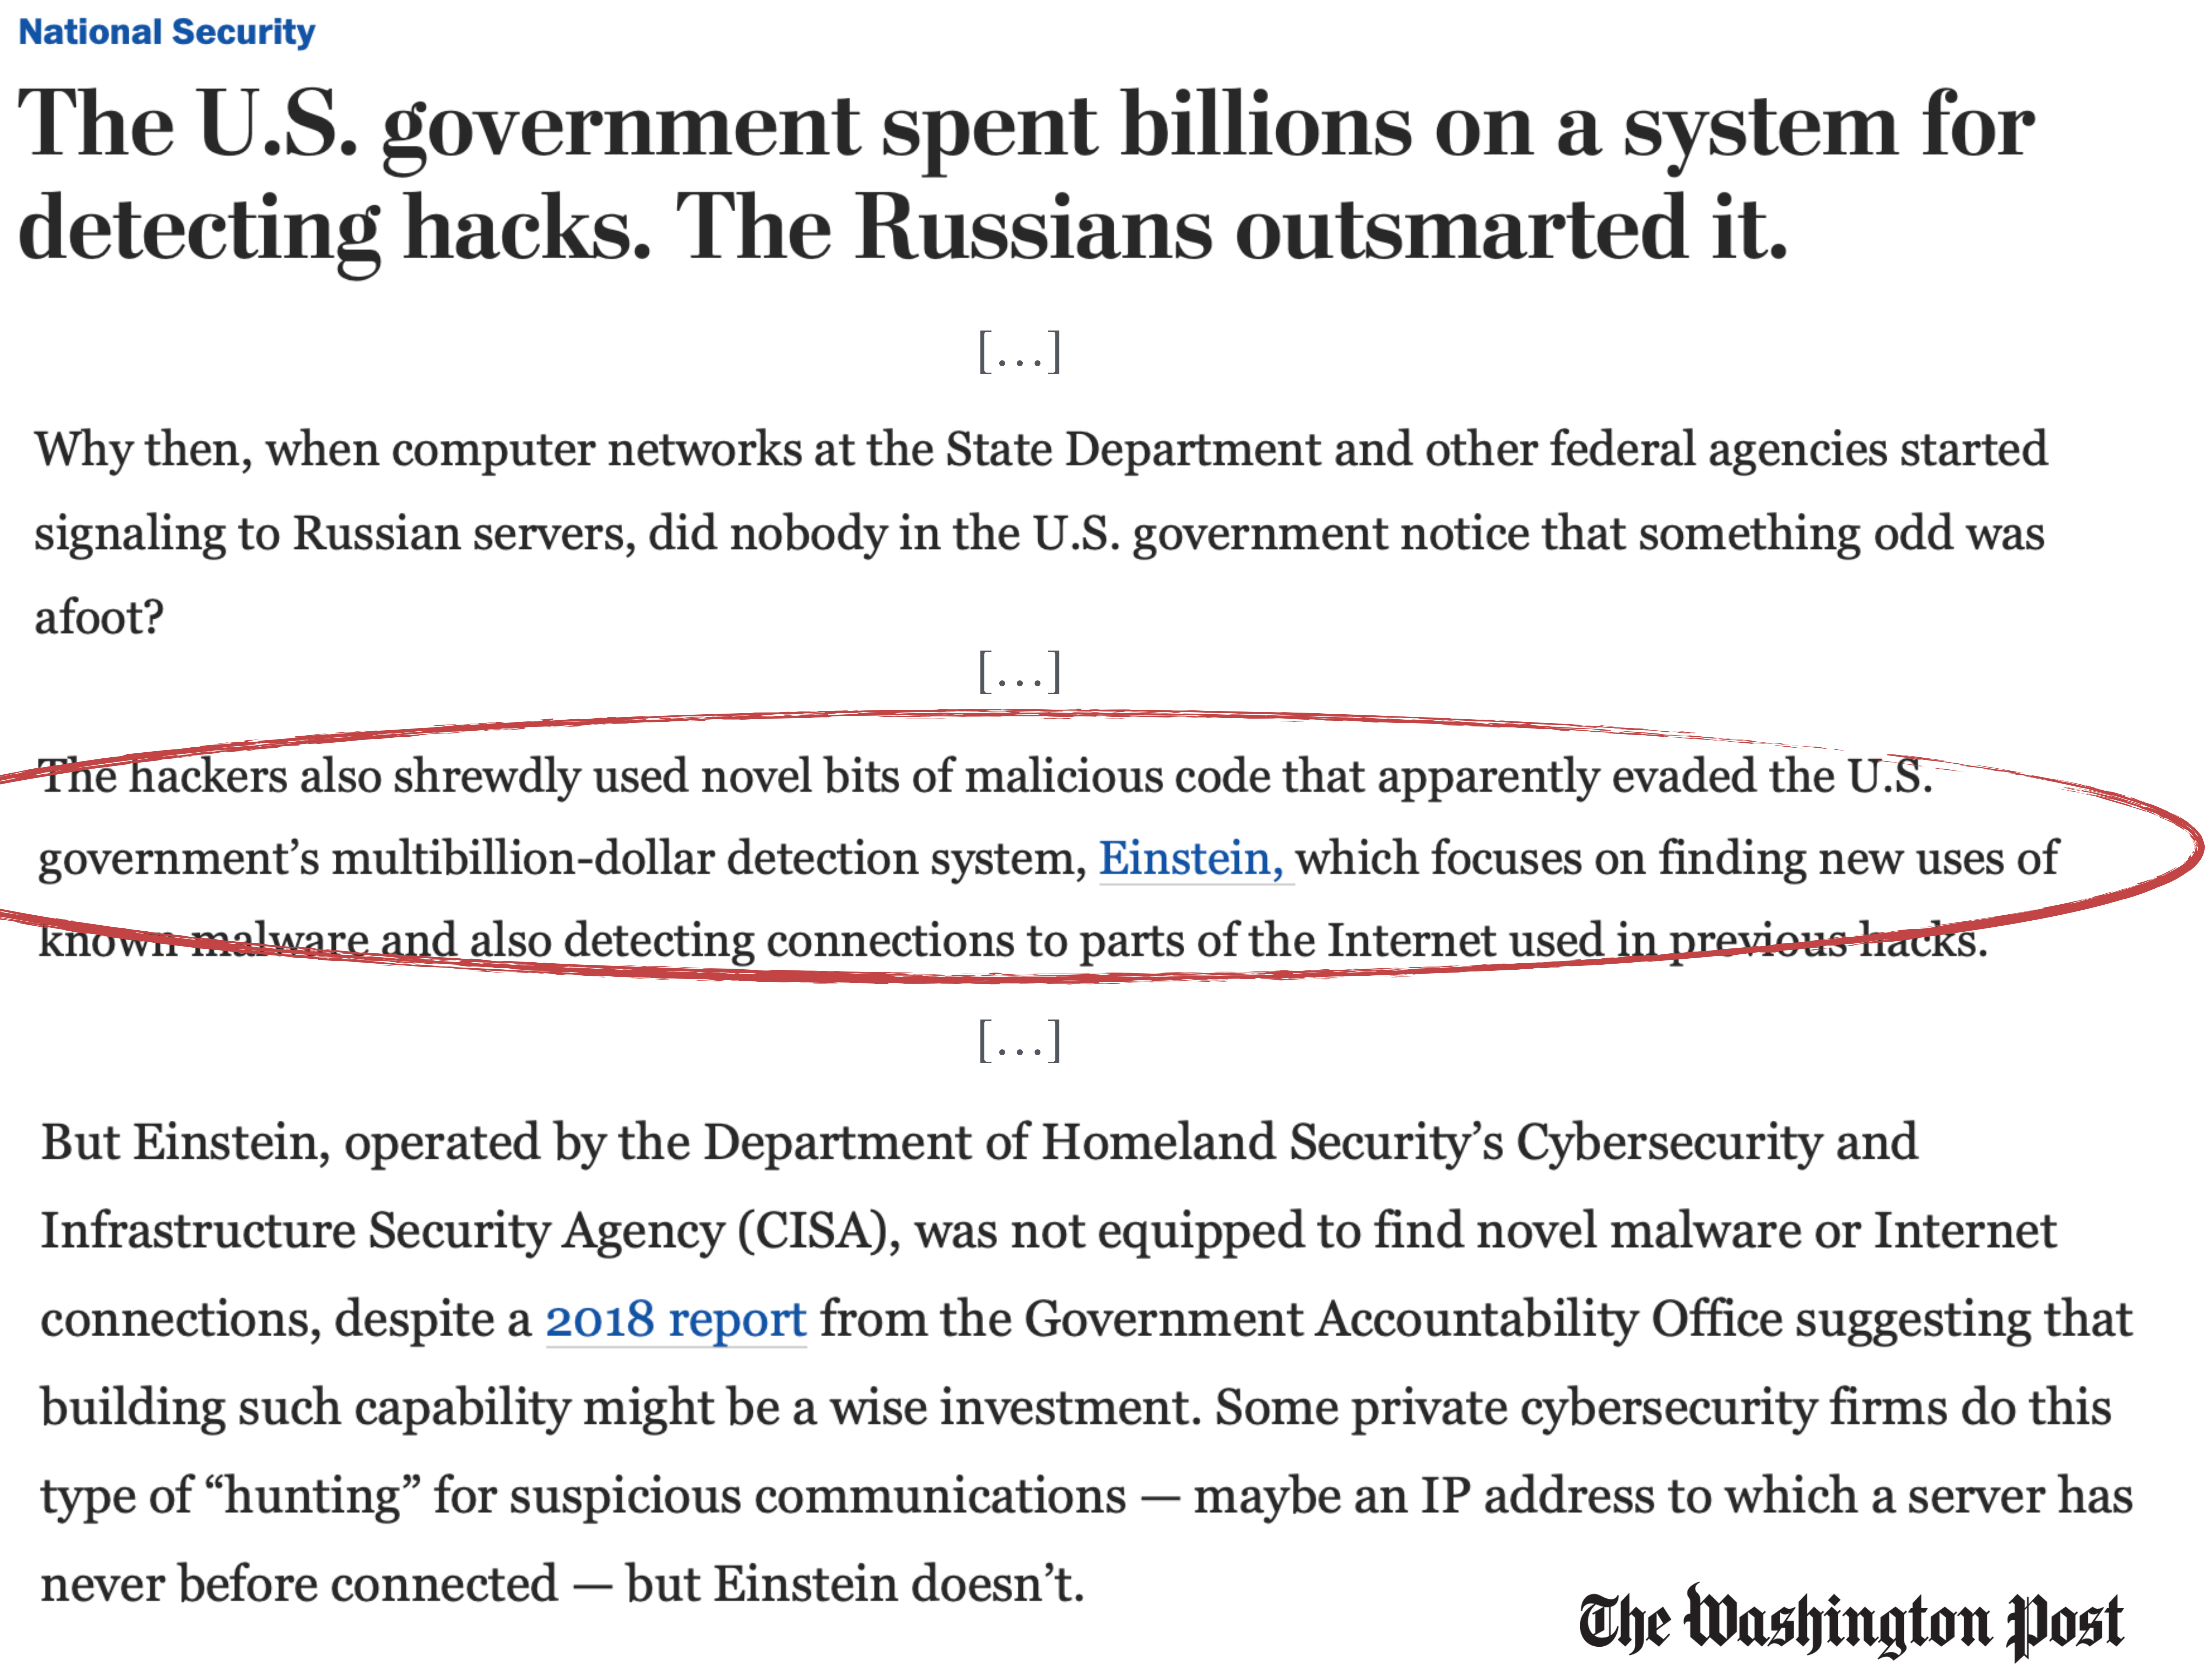

National Security

The U.S. government spent billions on a system for detecting hacks. The Russians outsmarted it.

[…]

Why then, when computer networks at the State Department and other federal agencies started signaling to Russian servers, did nobody in the U.S. government notice that something odd was afoot?

[…]

The hackers also shrewdly used novel bits of malicious code that apparently evaded the U.S. government's multibillion-dollar detection system, Einstein, which focuses on finding new uses of known malware and also detecting connections to parts of the Internet used in previous hacks.

[…]

But Einstein, operated by the Department of Homeland Security's Cybersecurity and Infrastructure Security Agency (CISA), was not equipped to find novel malware or Internet connections, despite a 2018 report from the Government Accountability Office suggesting that building such capability might be a wise investment. Some private cybersecurity firms do this type of "hunting" for suspicious communications — maybe an IP address to which a server has never before connected — but Einstein doesn't.

The Washington Post

Why wasn't this detected earlier?

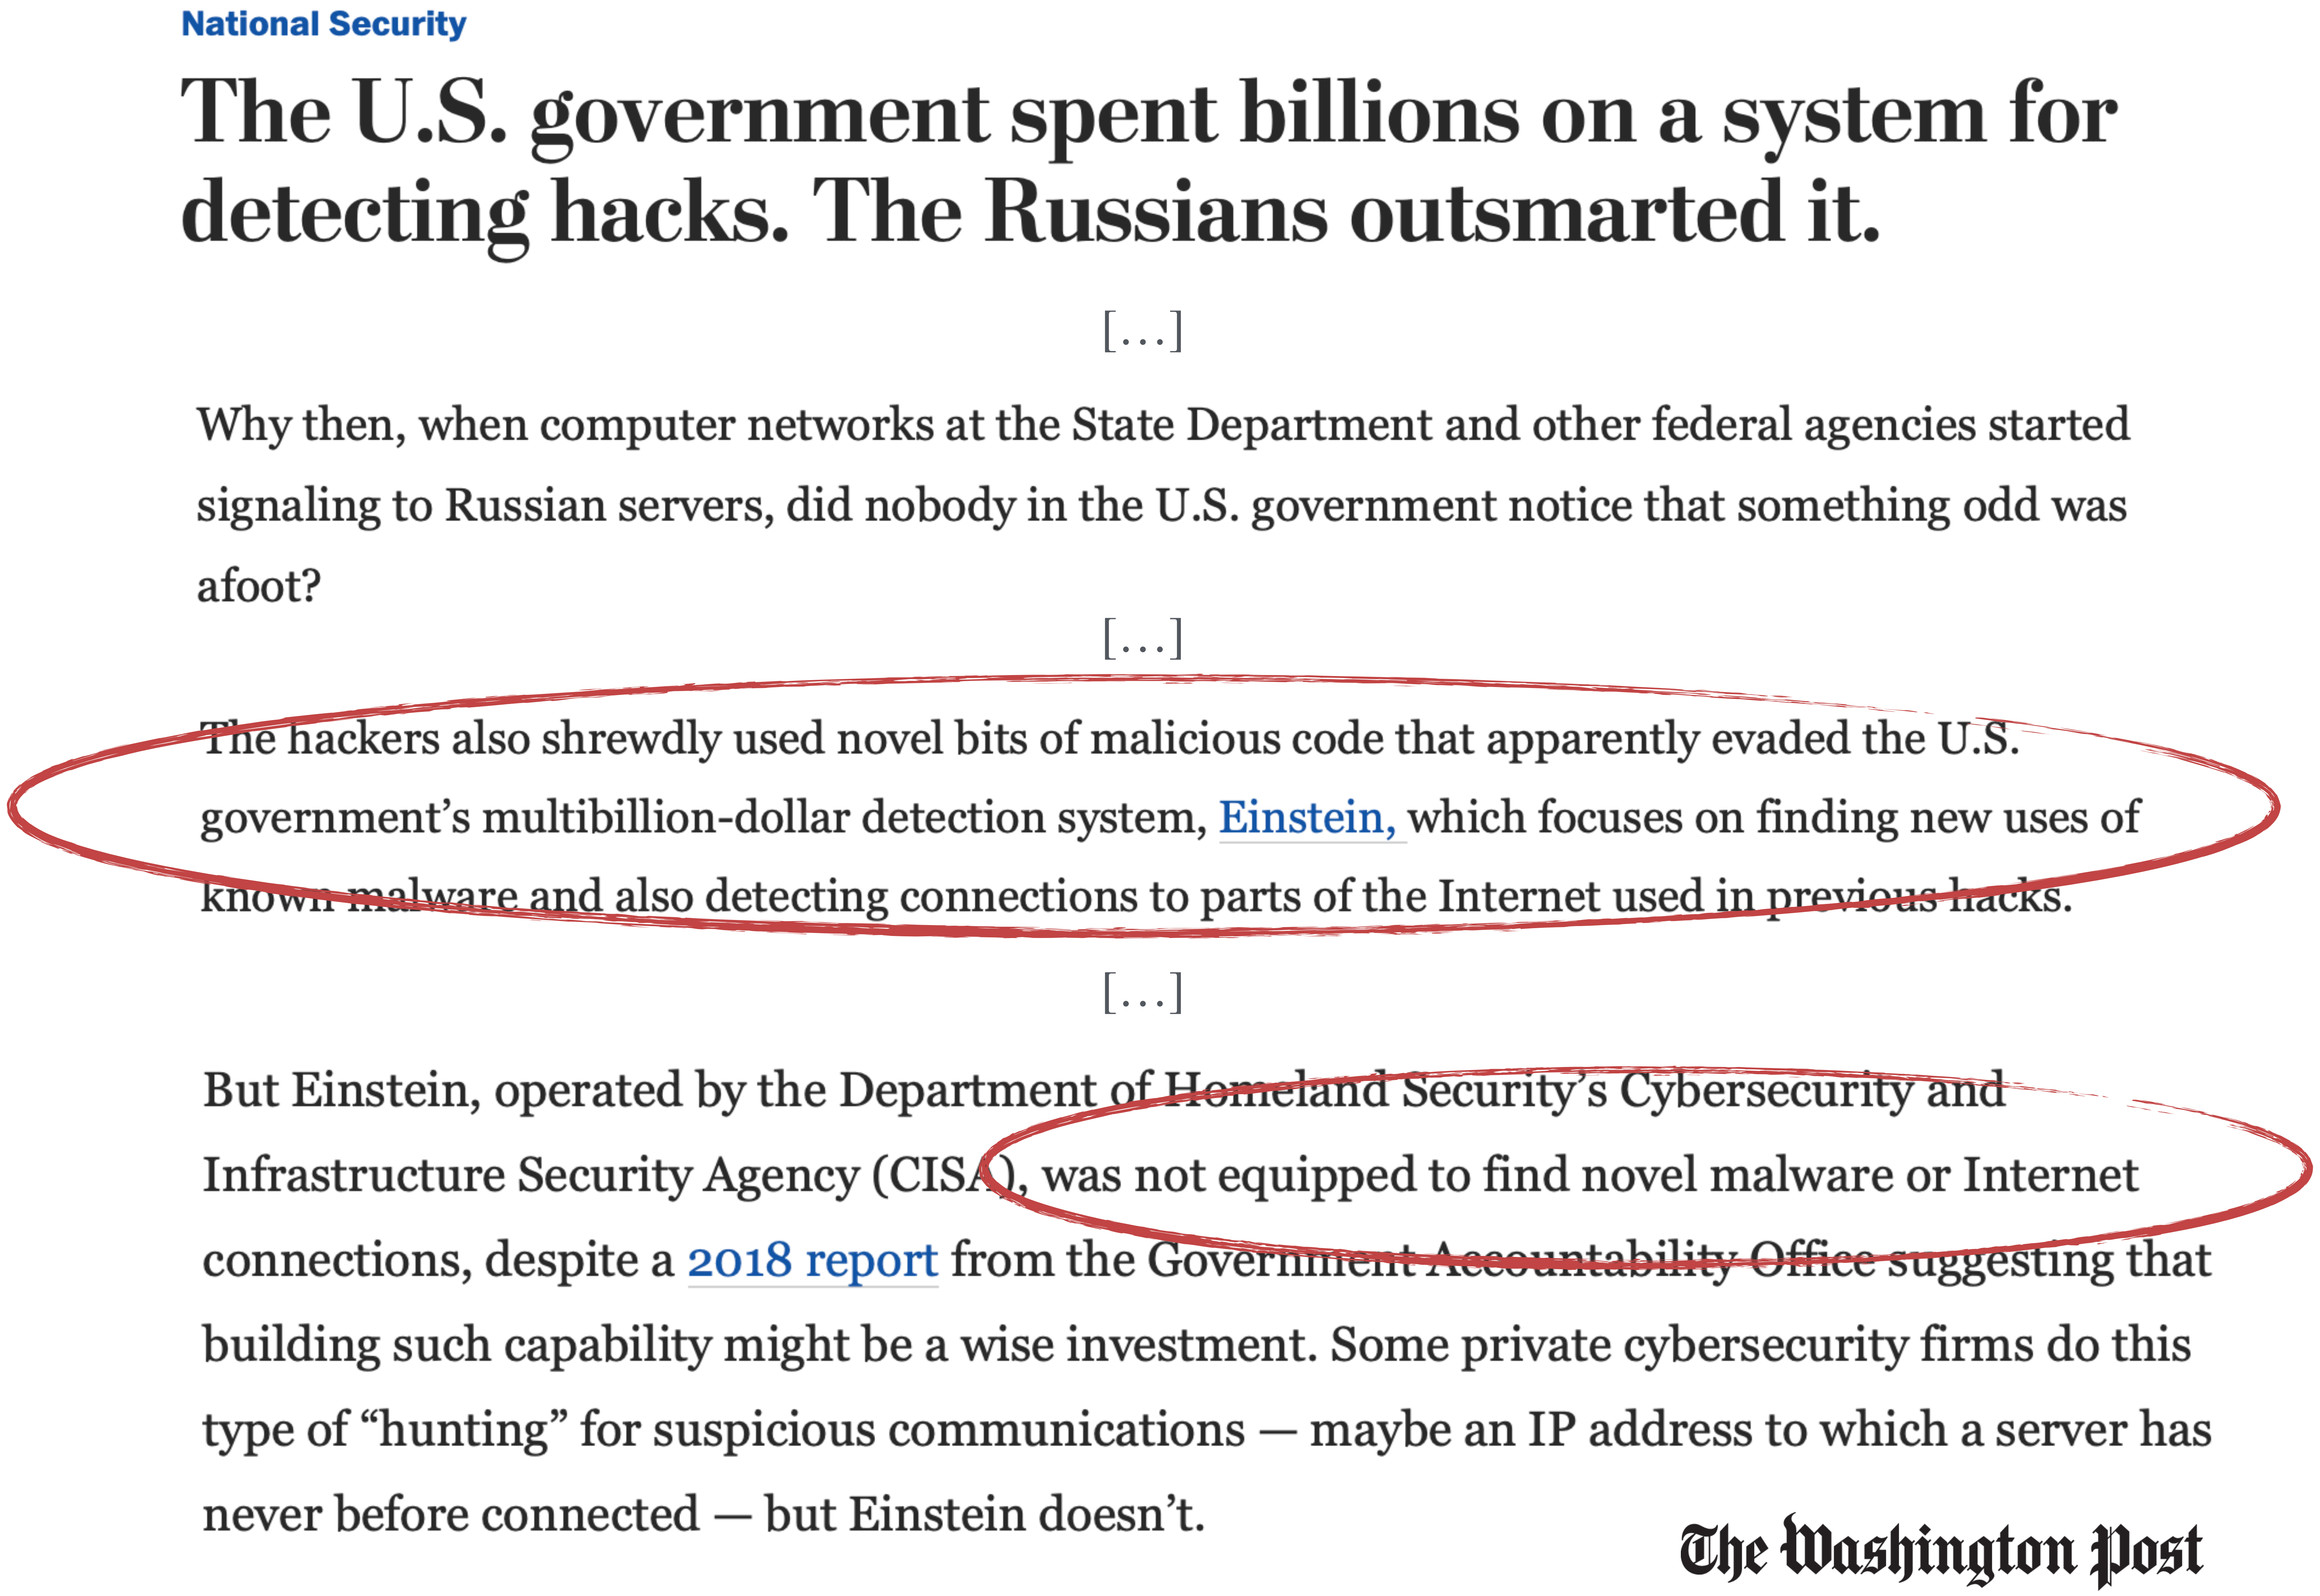

National Security

The U.S. government spent billions on a system for detecting hacks. The Russians outsmarted it.

[…]

Why then, when computer networks at the State Department and other federal agencies started signaling to Russian servers, did nobody in the U.S. government notice that something odd was afoot?

[…]

The hackers also shrewdly used novel bits of malicious code that apparently evaded the U.S. government's multibillion-dollar detection system, Einstein, which focuses on finding new uses of known malware and also detecting connections to parts of the Internet used in previous hacks.

[…]

But Einstein, operated by the Department of Homeland Security's Cybersecurity and Infrastructure Security Agency (CISA), was not equipped to find novel malware or Internet connections, despite a 2018 report from the Government Accountability Office suggesting that building such capability might be a wise investment. Some private cybersecurity firms do this type of "hunting" for suspicious communications — maybe an IP address to which a server has never before connected — but Einstein doesn't.

The Washington Post

Why wasn’t this detected earlier?

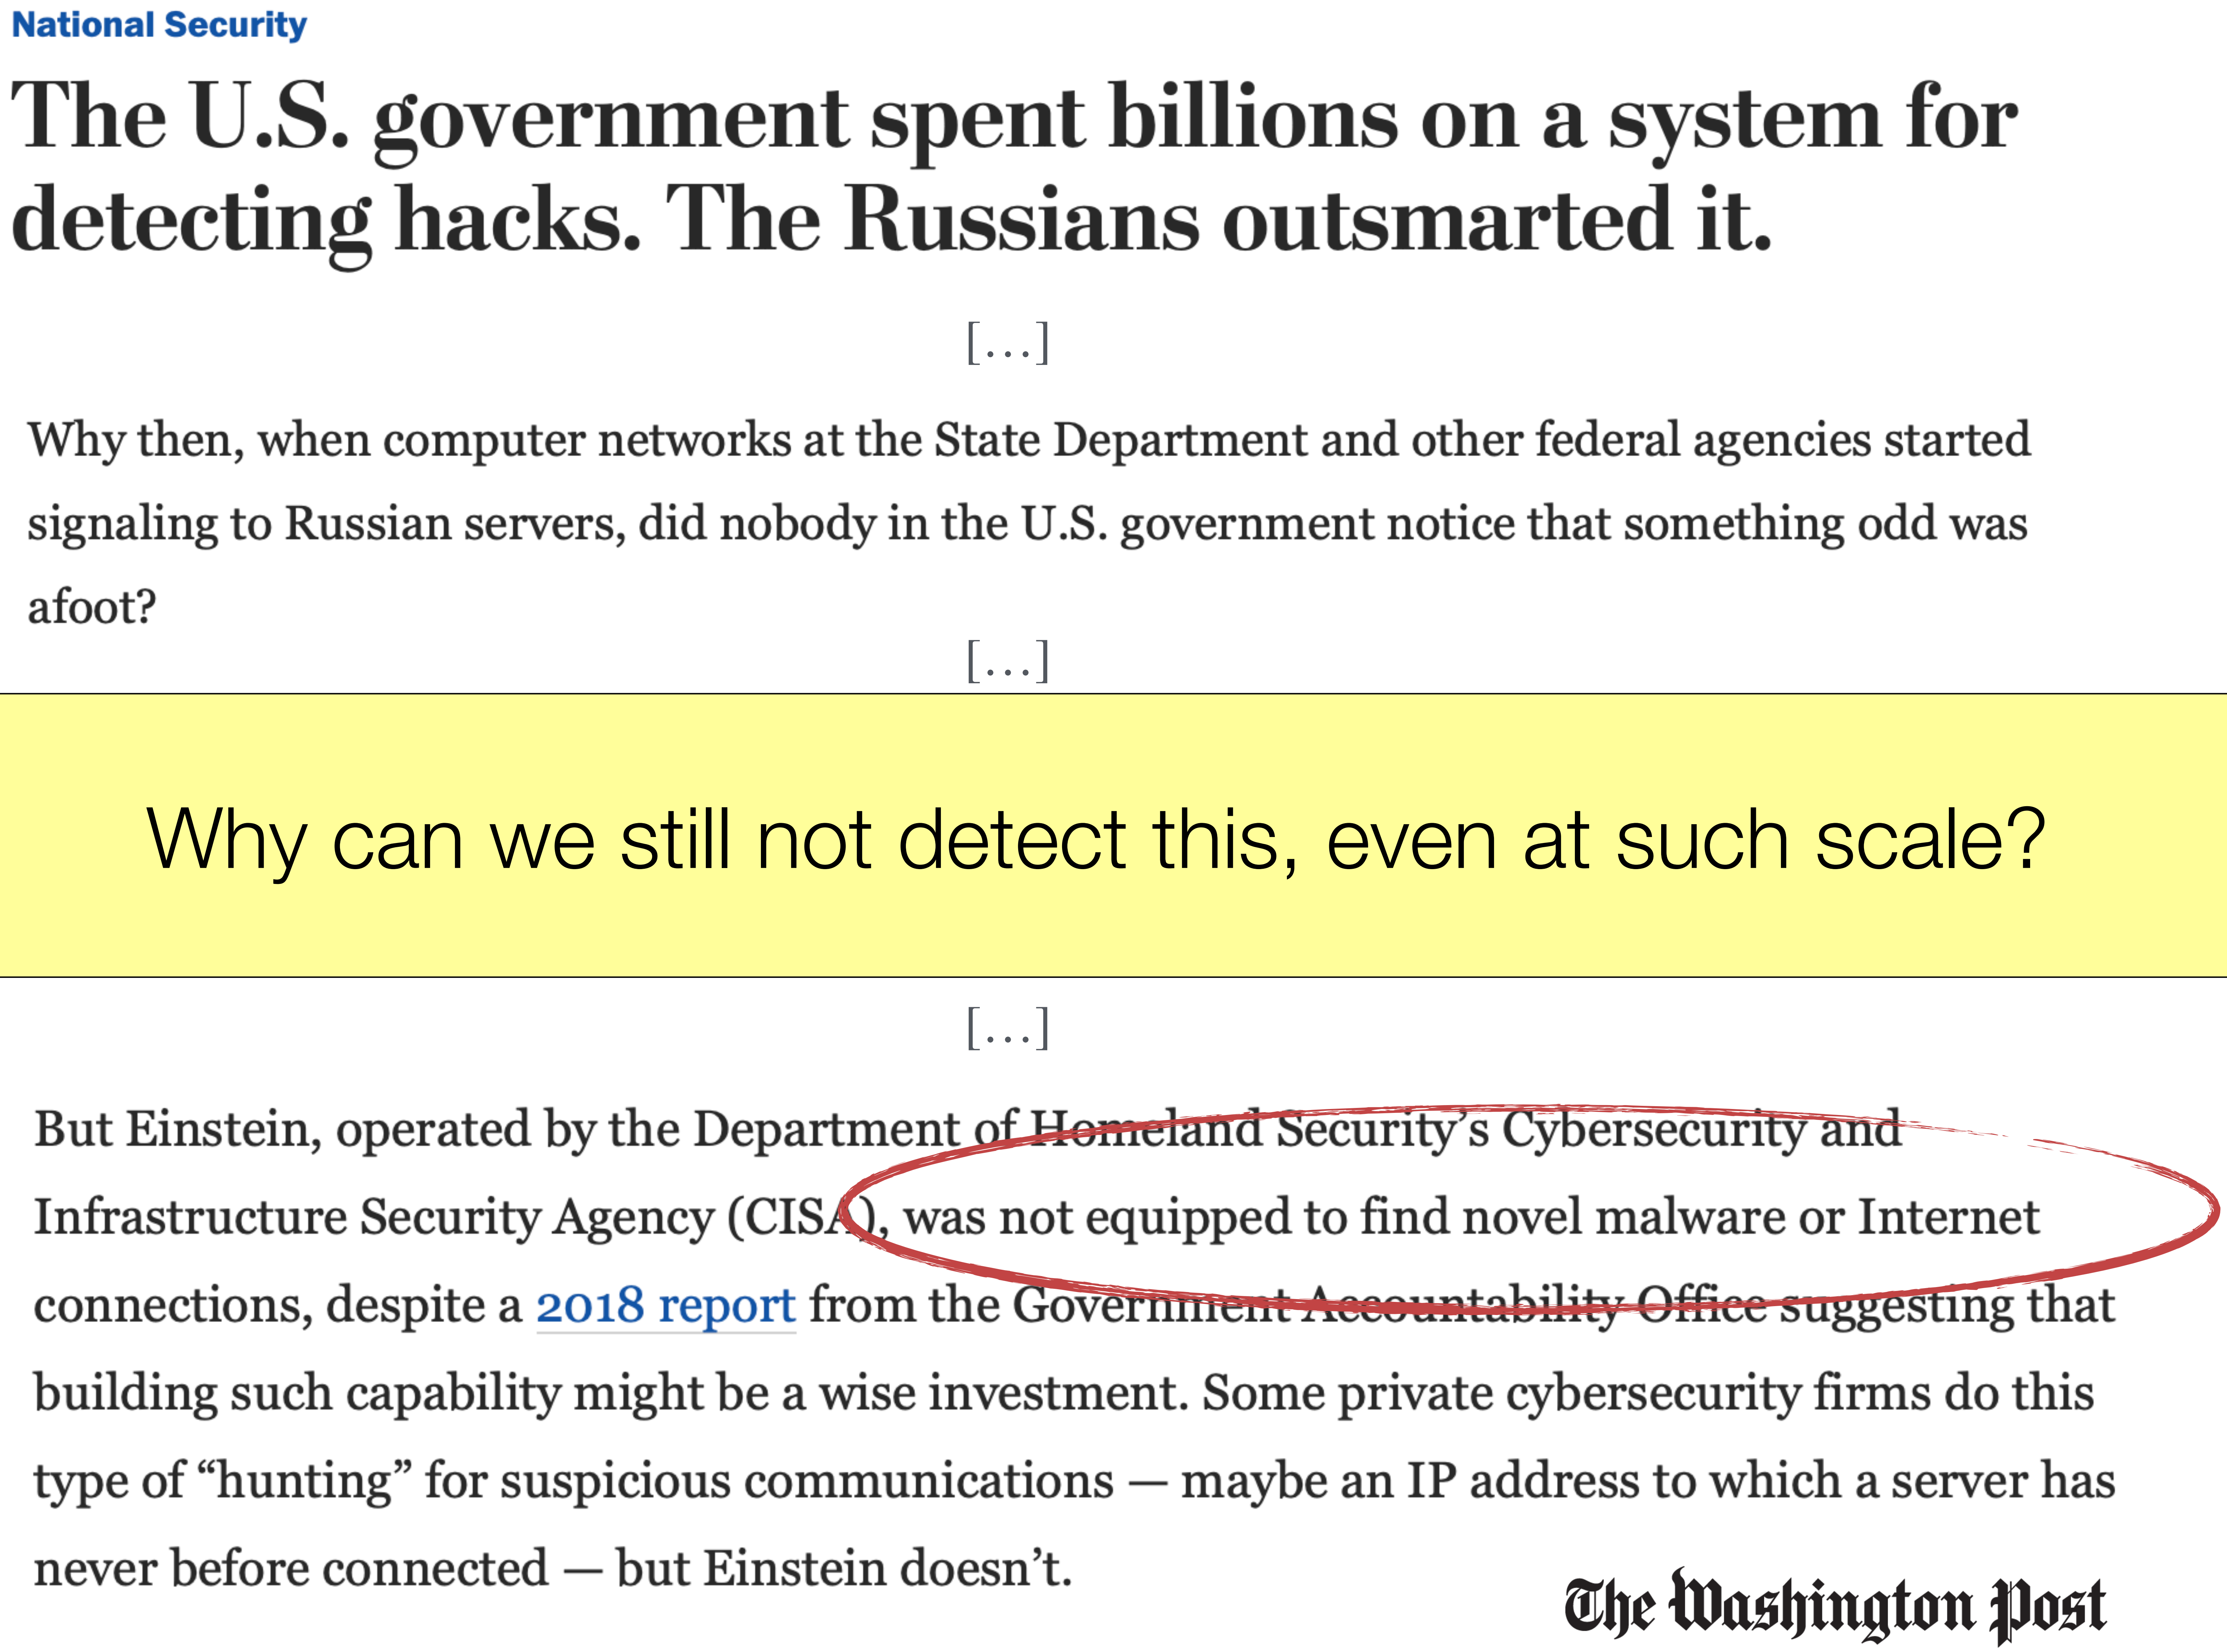

National Security

The U.S. government spent billions on a system for detecting hacks. The Russians outsmarted it.

[…]

Why then, when computer networks at the State Department and other federal agencies started signaling to Russian servers, did nobody in the U.S. government notice that something odd was afoot?

[…]

Why can we still not detect this, even at such scale?

[…]

But Einstein, operated by the Department of Homeland Security’s Cybersecurity and Infrastructure Security Agency (CISA), was not equipped to find novel malware or Internet connections, despite a 2018 report from the Government Accountability Office suggesting that building such capability might be a wise investment. Some private cybersecurity firms do this type of “hunting” for suspicious communications — maybe an IP address to which a server has never before connected — but Einstein doesn’t.

The Washington Post

Outline

1. Classic intrusion detection with machine learning
2. From intrusion detection to threat hunting
3. Beyond detection: A new role for AI

Classic Intrusion Detection

How can we detect (novel) attacks?

Analysis approaches

Misuse Detection

(using signatures)

Look for know attacks that we can describe

Analysis approaches

Misuse Detection
(using signatures)

Look for know attacks that we can describe

Anomaly Detection
(using machine learning)

Look for activity that’s “not normal”

Analysis approaches

Misuse Detection

(using signatures)

Look for know attacks that we can describe

Anomaly Detection

(using machine learning)

Look for activity that’s “not normal”

This is the Holy Grail of intrusion detection …

Early academic research

Early academic research

Two degrees of freedom

Input	Decide on features
ML	Select classifier

Early academic research

Two degrees of freedom

Input	Decide on features
ML	Select classifier

Network features used

- packet sizes
- IP addresses
- ports
- header fields
- timestamps
- inter-arrival times
- session size
- session duration
- session volume
- payload frequencies
- payload tokens
- payload pattern
- ...

Early academic research

Two degrees of freedom

Input	Decide on features
ML	Select classifier

Network features used

- packet sizes
- IP addresses
- ports
- header fields
- timestamps
- inter-arrival times
- session size
- session duration
- session volume
- payload frequencies
- payload tokens
- payload pattern
- ...

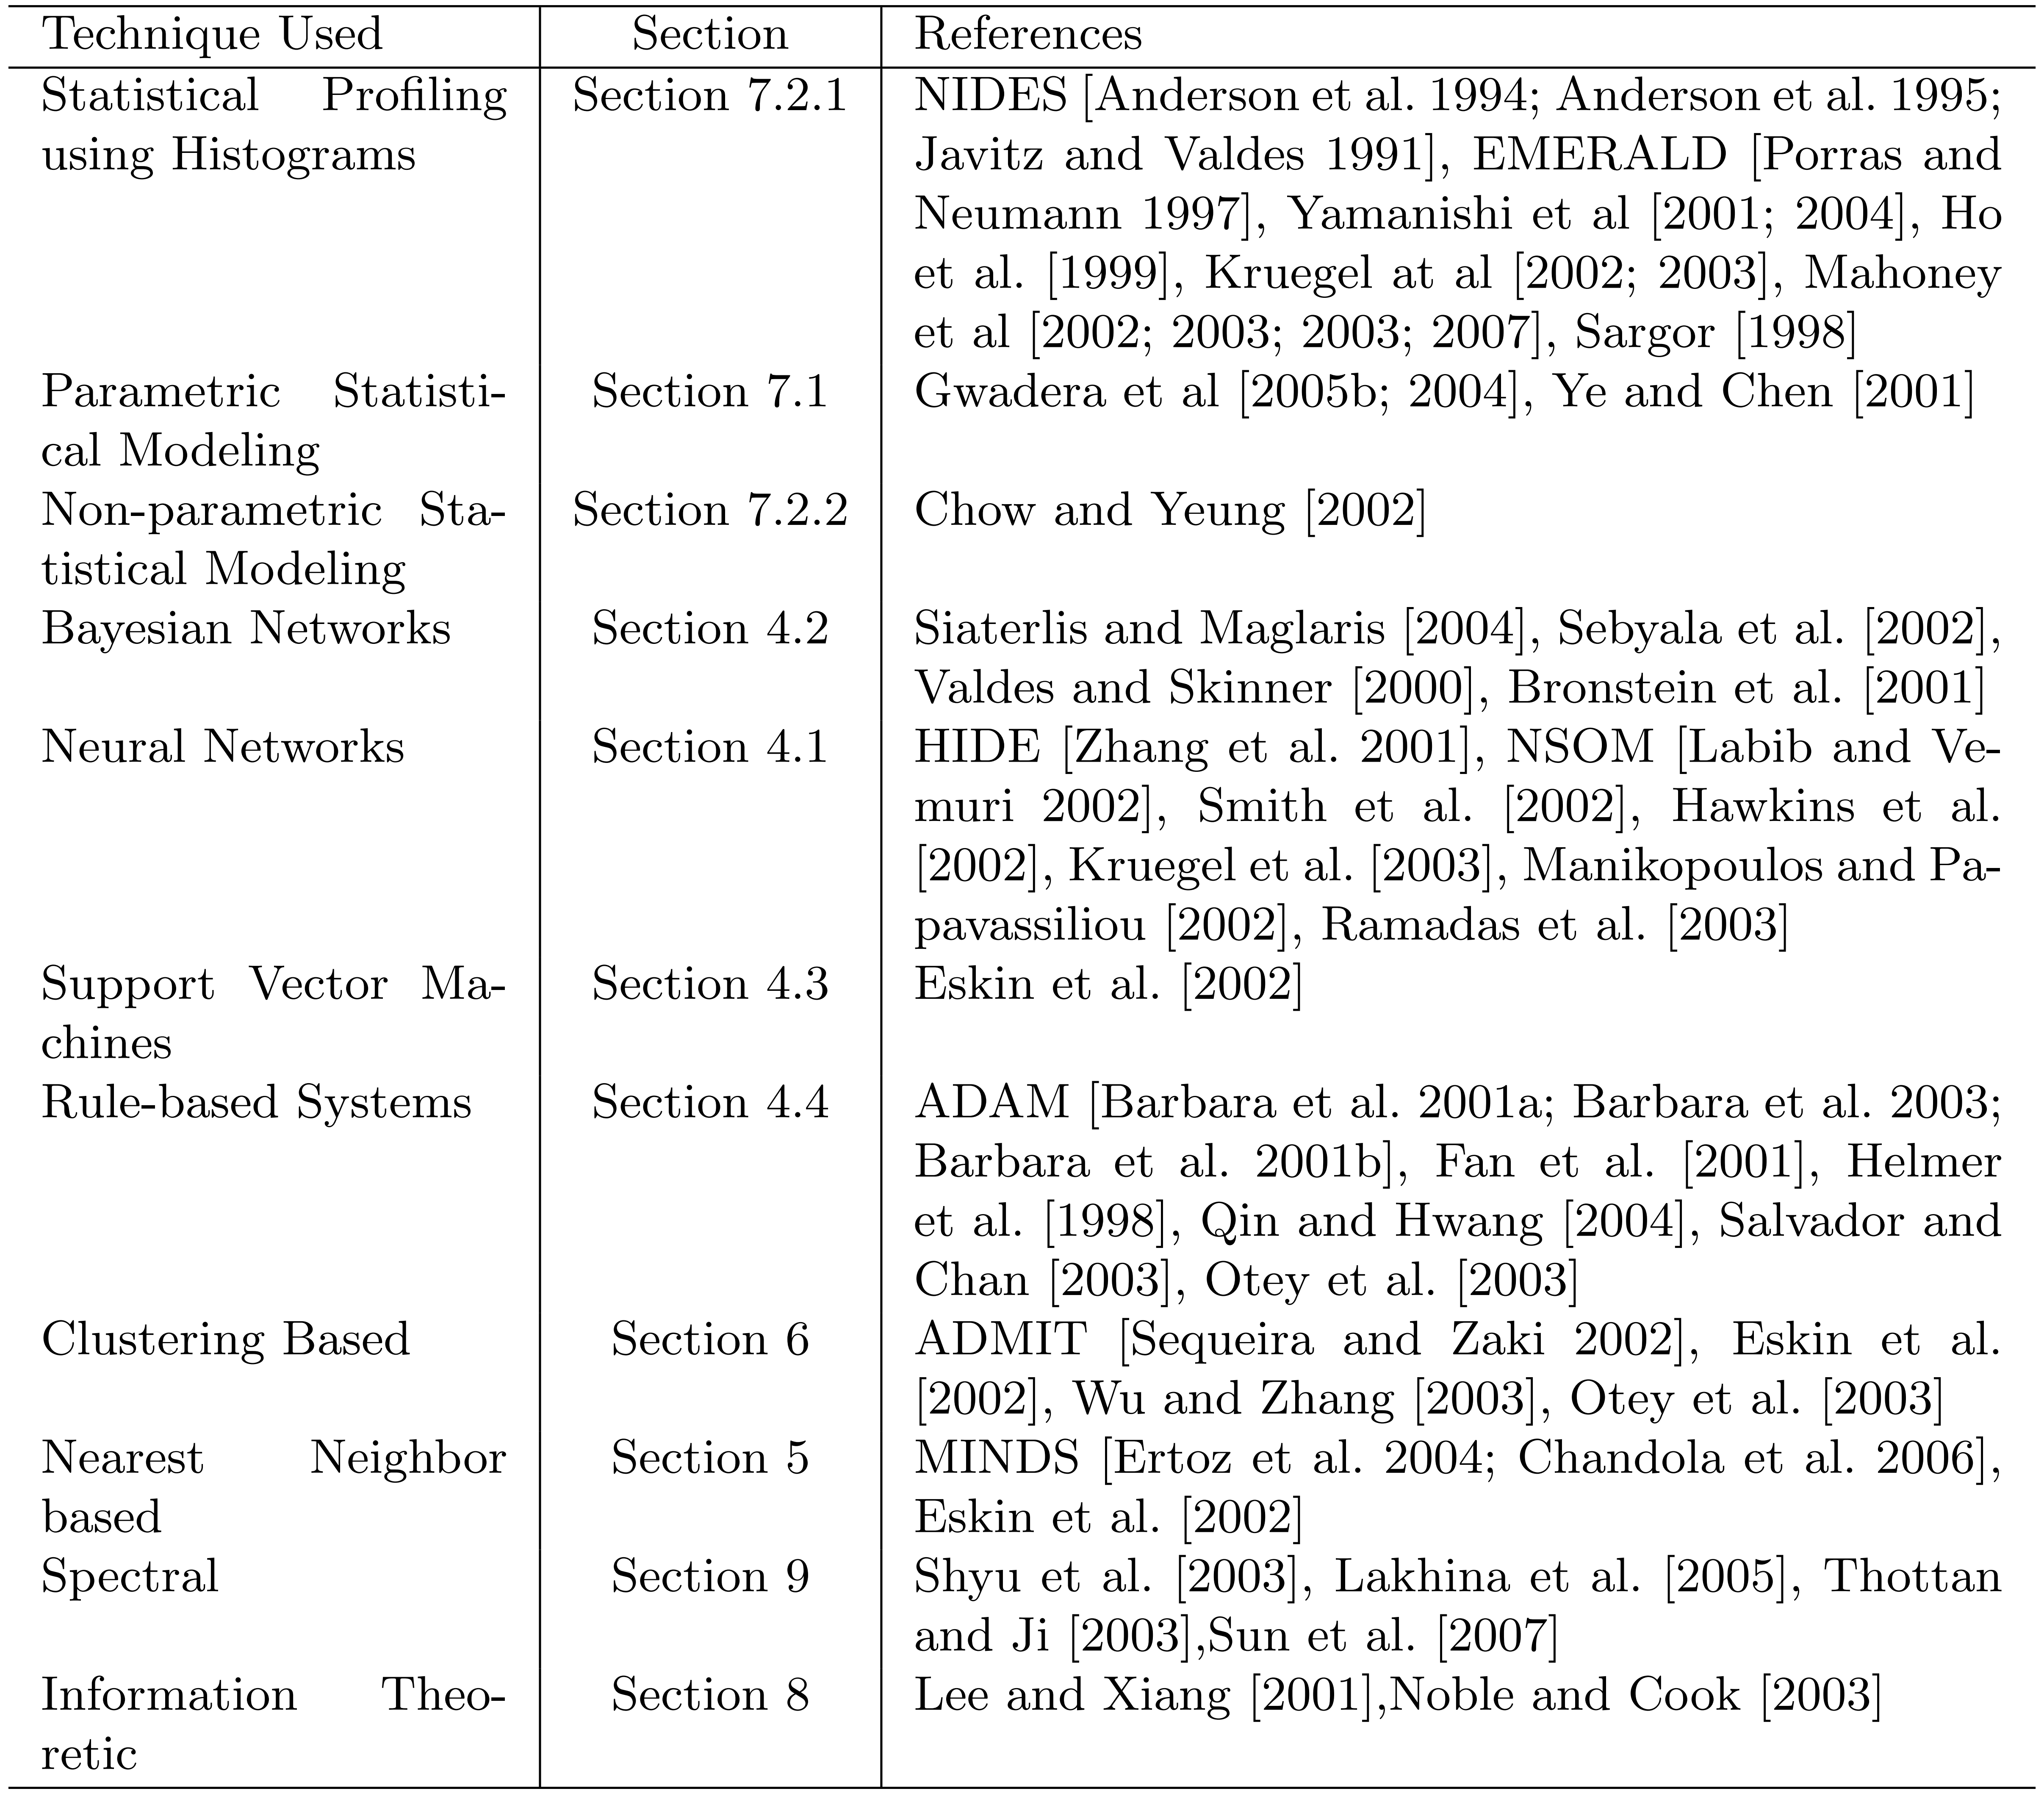

Technique Used	Section	References
Statistical Profiling using Histograms	Section 7.2.1	NIDES [Anderson et al. 1994; Anderson et al. 1995; Javitz and Valdes 1991], EMERALD [Porras and Neumann 1997], Yamanishi et al [2001; 2004], Ho et al. [1999], Kruegel at al [2002; 2003], Mahoney et al [2002; 2003; 2003; 2007], Sargor [1998]
Parametric Statistical Modeling	Section 7.1	Gwadera et al [2005b; 2004], Ye and Chen [2001]
Non-parametric Statistical Modeling	Section 7.2.2	Chow and Yeung [2002]
Bayesian Networks	Section 4.2	Siaterlis and Maglaris [2004], Sebyala et al. [2002], Valdes and Skinner [2000], Bronstein et al. [2001]
Neural Networks	Section 4.1	HIDE [Zhang et al. 2001], NSOM [Labib and Vemuri 2002], Smith et al. [2002], Hawkins et al. [2002], Kruegel et al. [2003], Manikopoulos and Papavassiliou [2002], Ramadas et al. [2003]
Support Vector Machines	Section 4.3	Eskin et al. [2002]
Rule-based Systems	Section 4.4	ADAM [Barbara et al. 2001a; Barbara et al. 2003; Barbara et al. 2001b], Fan et al. [2001], Helmer et al. [1998], Qin and Hwang [2004], Salvador and Chan [2003], Otey et al. [2003]
Clustering Based	Section 6	ADMIT [Sequeira and Zaki 2002], Eskin et al. [2002], Wu and Zhang [2003], Otey et al. [2003]
Nearest Neighbor based	Section 5	MINDS [Ertoz et al. 2004; Chandola et al. 2006], Eskin et al. [2002]
Spectral	Section 9	Shyu et al. [2003], Lakhina et al. [2005], Thottan and Ji [2003],Sun et al. [2007]
Information Theoretic	Section 8	Lee and Xiang [2001],Noble and Cook [2003]

Source: Chandola et al. 2009

Early academic research

Two degrees of freedom

Input	Decide on features
ML	Select classifier

Network features used

- packet sizes
- IP addresses
- ports
- header fields
- timestamps
- inter-arrival times
- session size
- session duration
- session volume
- payload frequencies
- payload tokens
- payload pattern
- ...

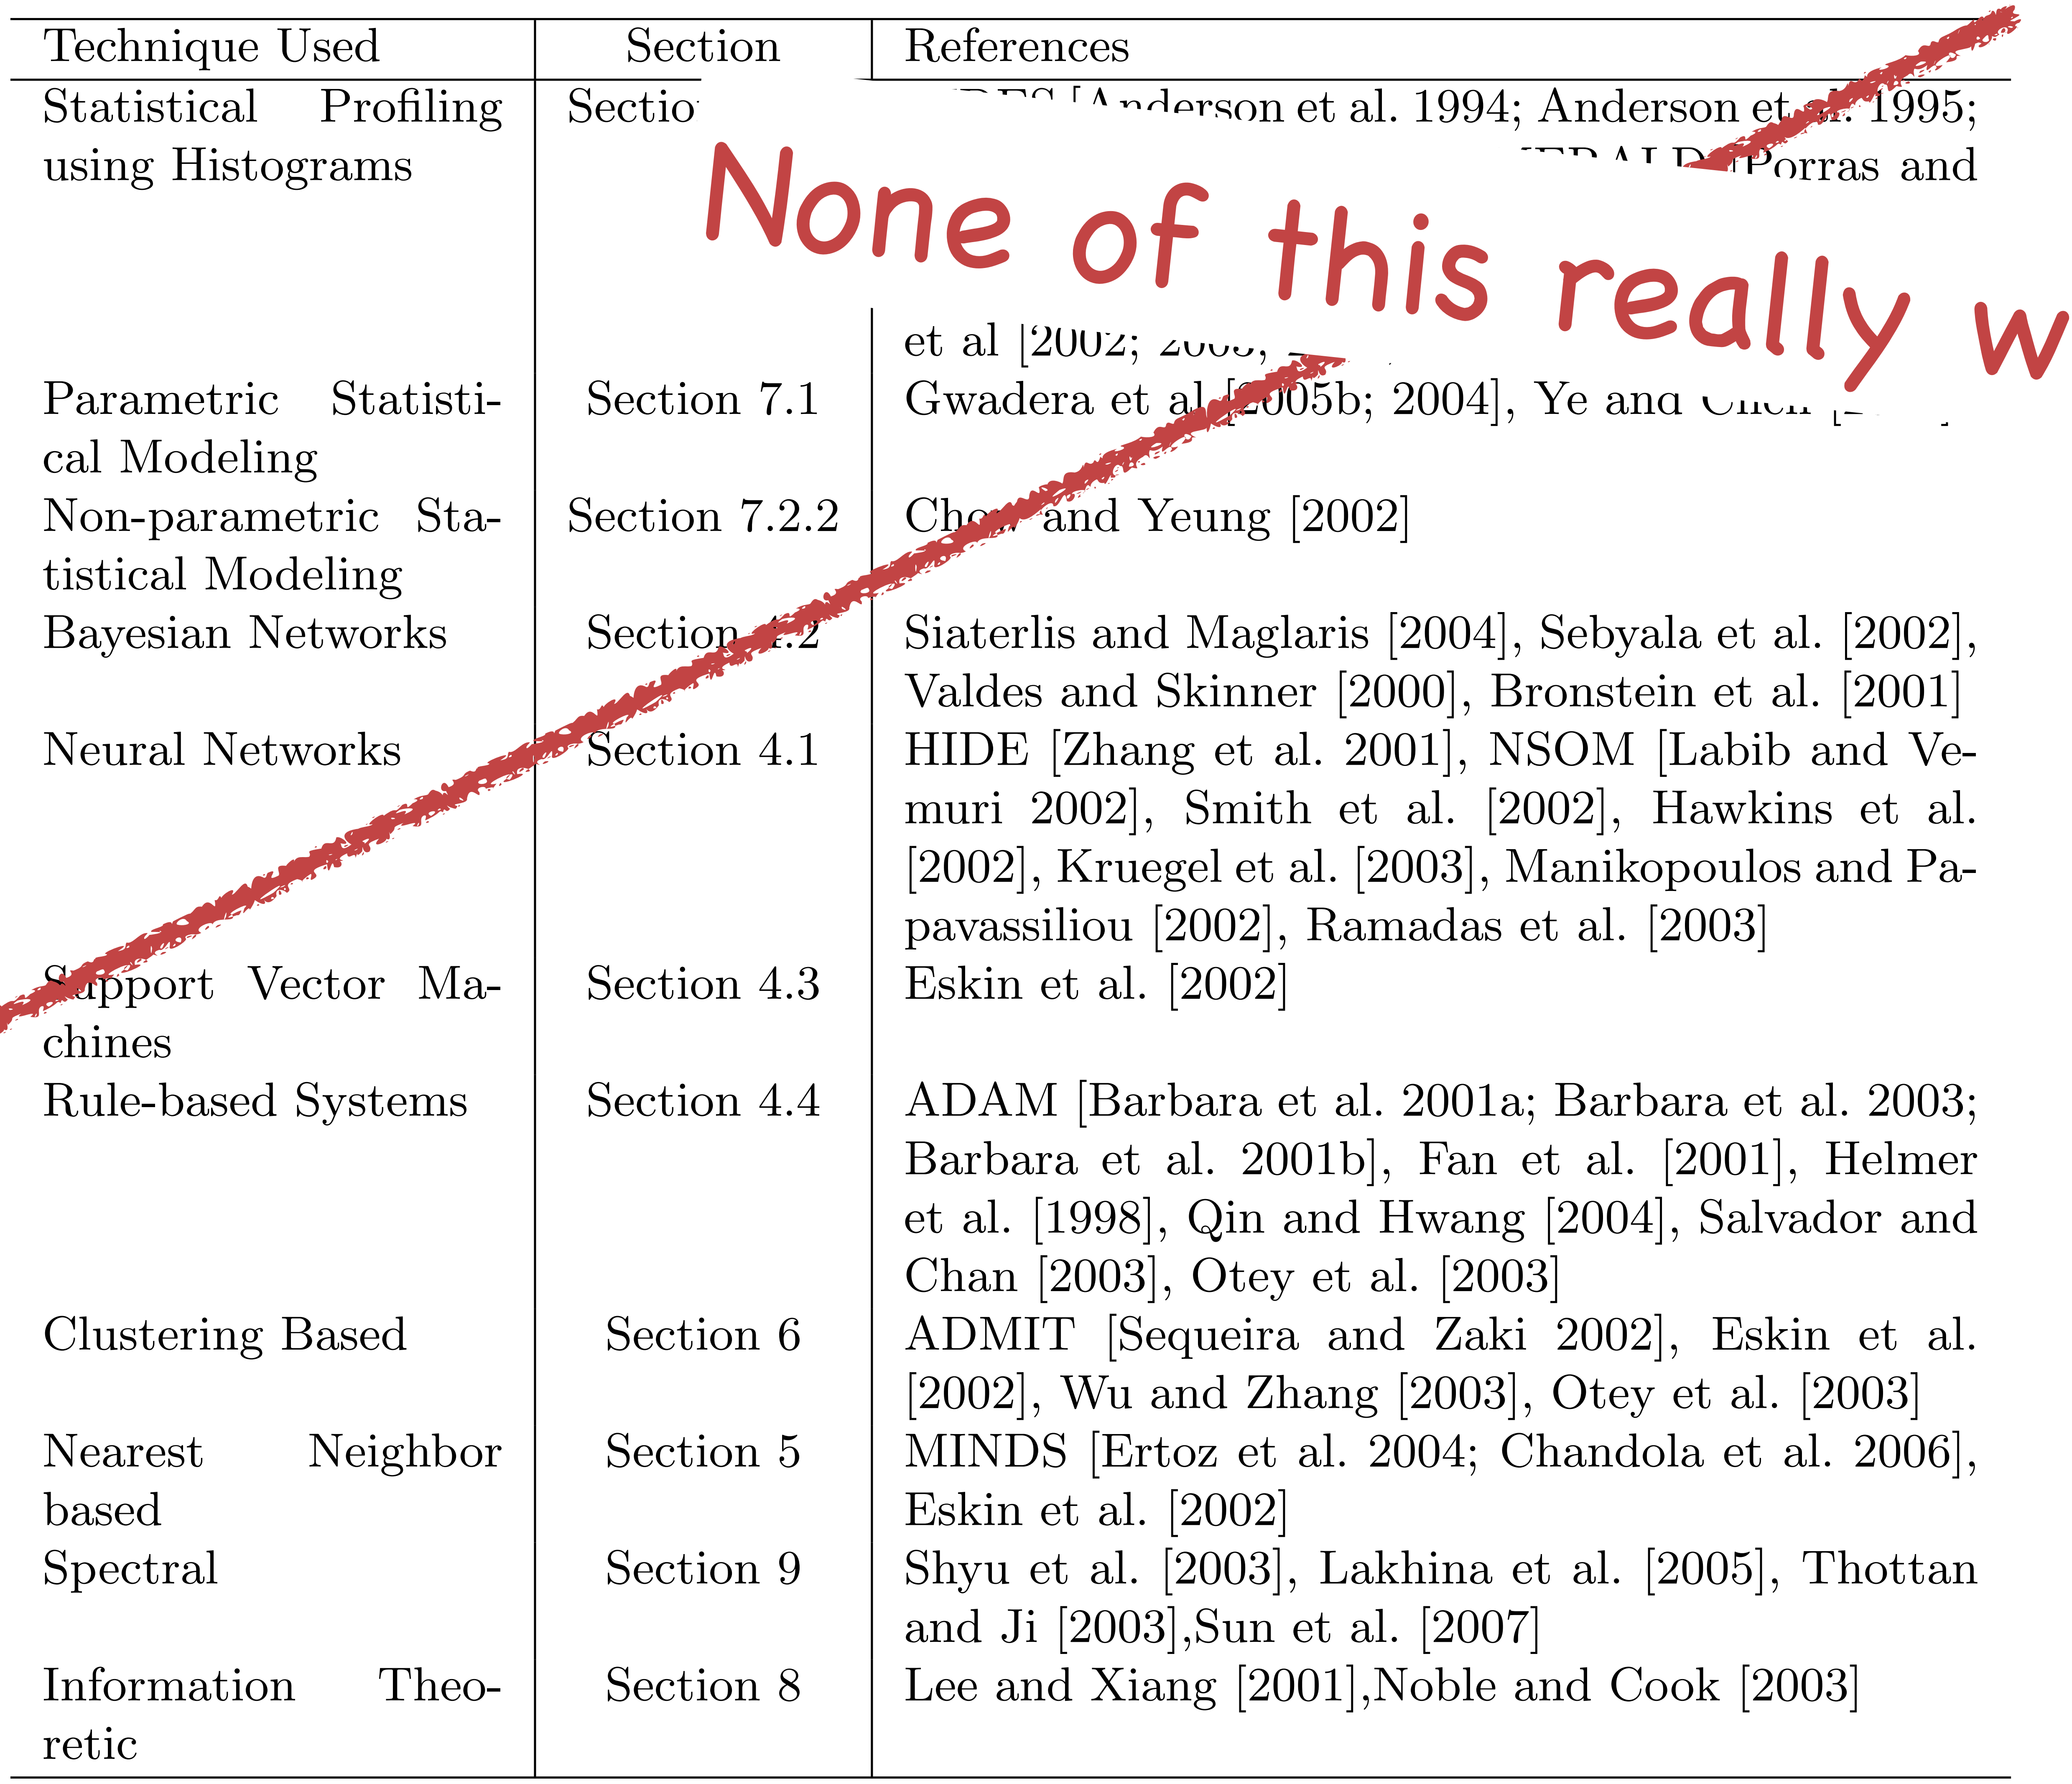

Technique Used	Section	References
Statistical Profiling using Histograms	Sectio[illegible]	[illegible] [Anderson et al. 1994; Anderson et al. 1995; [illegible] Porras and [illegible] et al [2002; [illegible]
Parametric Statistical Modeling	Section 7.1	Gwadera et al. [2005b; 2004], Ye and Chen [illegible]
Non-parametric Statistical Modeling	Section 7.2.2	Chow and Yeung [2002]
Bayesian Networks	Section 4.2	Siaterlis and Maglaris [2004], Sebyala et al. [2002], Valdes and Skinner [2000], Bronstein et al. [2001]
Neural Networks	Section 4.1	HIDE [Zhang et al. 2001], NSOM [Labib and Vemuri 2002], Smith et al. [2002], Hawkins et al. [2002], Kruegel et al. [2003], Manikopoulos and Papavassiliou [2002], Ramadas et al. [2003]
Support Vector Machines	Section 4.3	Eskin et al. [2002]
Rule-based Systems	Section 4.4	ADAM [Barbara et al. 2001a; Barbara et al. 2003; Barbara et al. 2001b], Fan et al. [2001], Helmer et al. [1998], Qin and Hwang [2004], Salvador and Chan [2003], Otey et al. [2003]
Clustering Based	Section 6	ADMIT [Sequeira and Zaki 2002], Eskin et al. [2002], Wu and Zhang [2003], Otey et al. [2003]
Nearest Neighbor based	Section 5	MINDS [Ertoz et al. 2004; Chandola et al. 2006], Eskin et al. [2002]
Spectral	Section 9	Shyu et al. [2003], Lakhina et al. [2005], Thottan and Ji [2003],Sun et al. [2007]
Information Theoretic	Section 8	Lee and Xiang [2001],Noble and Cook [2003]

None of this really works

Source: Chandola et al. 2009

Early academic research

Two degrees of freedom

Input	Decide on features
ML	Select classifier

Network features used

- packet sizes
- IP addresses
- ports
- header fields
- timestamps
- inter-arrival times
- session size
- session duration
- session volume
- payload frequencies
- payload tokens
- payload pattern
- ...

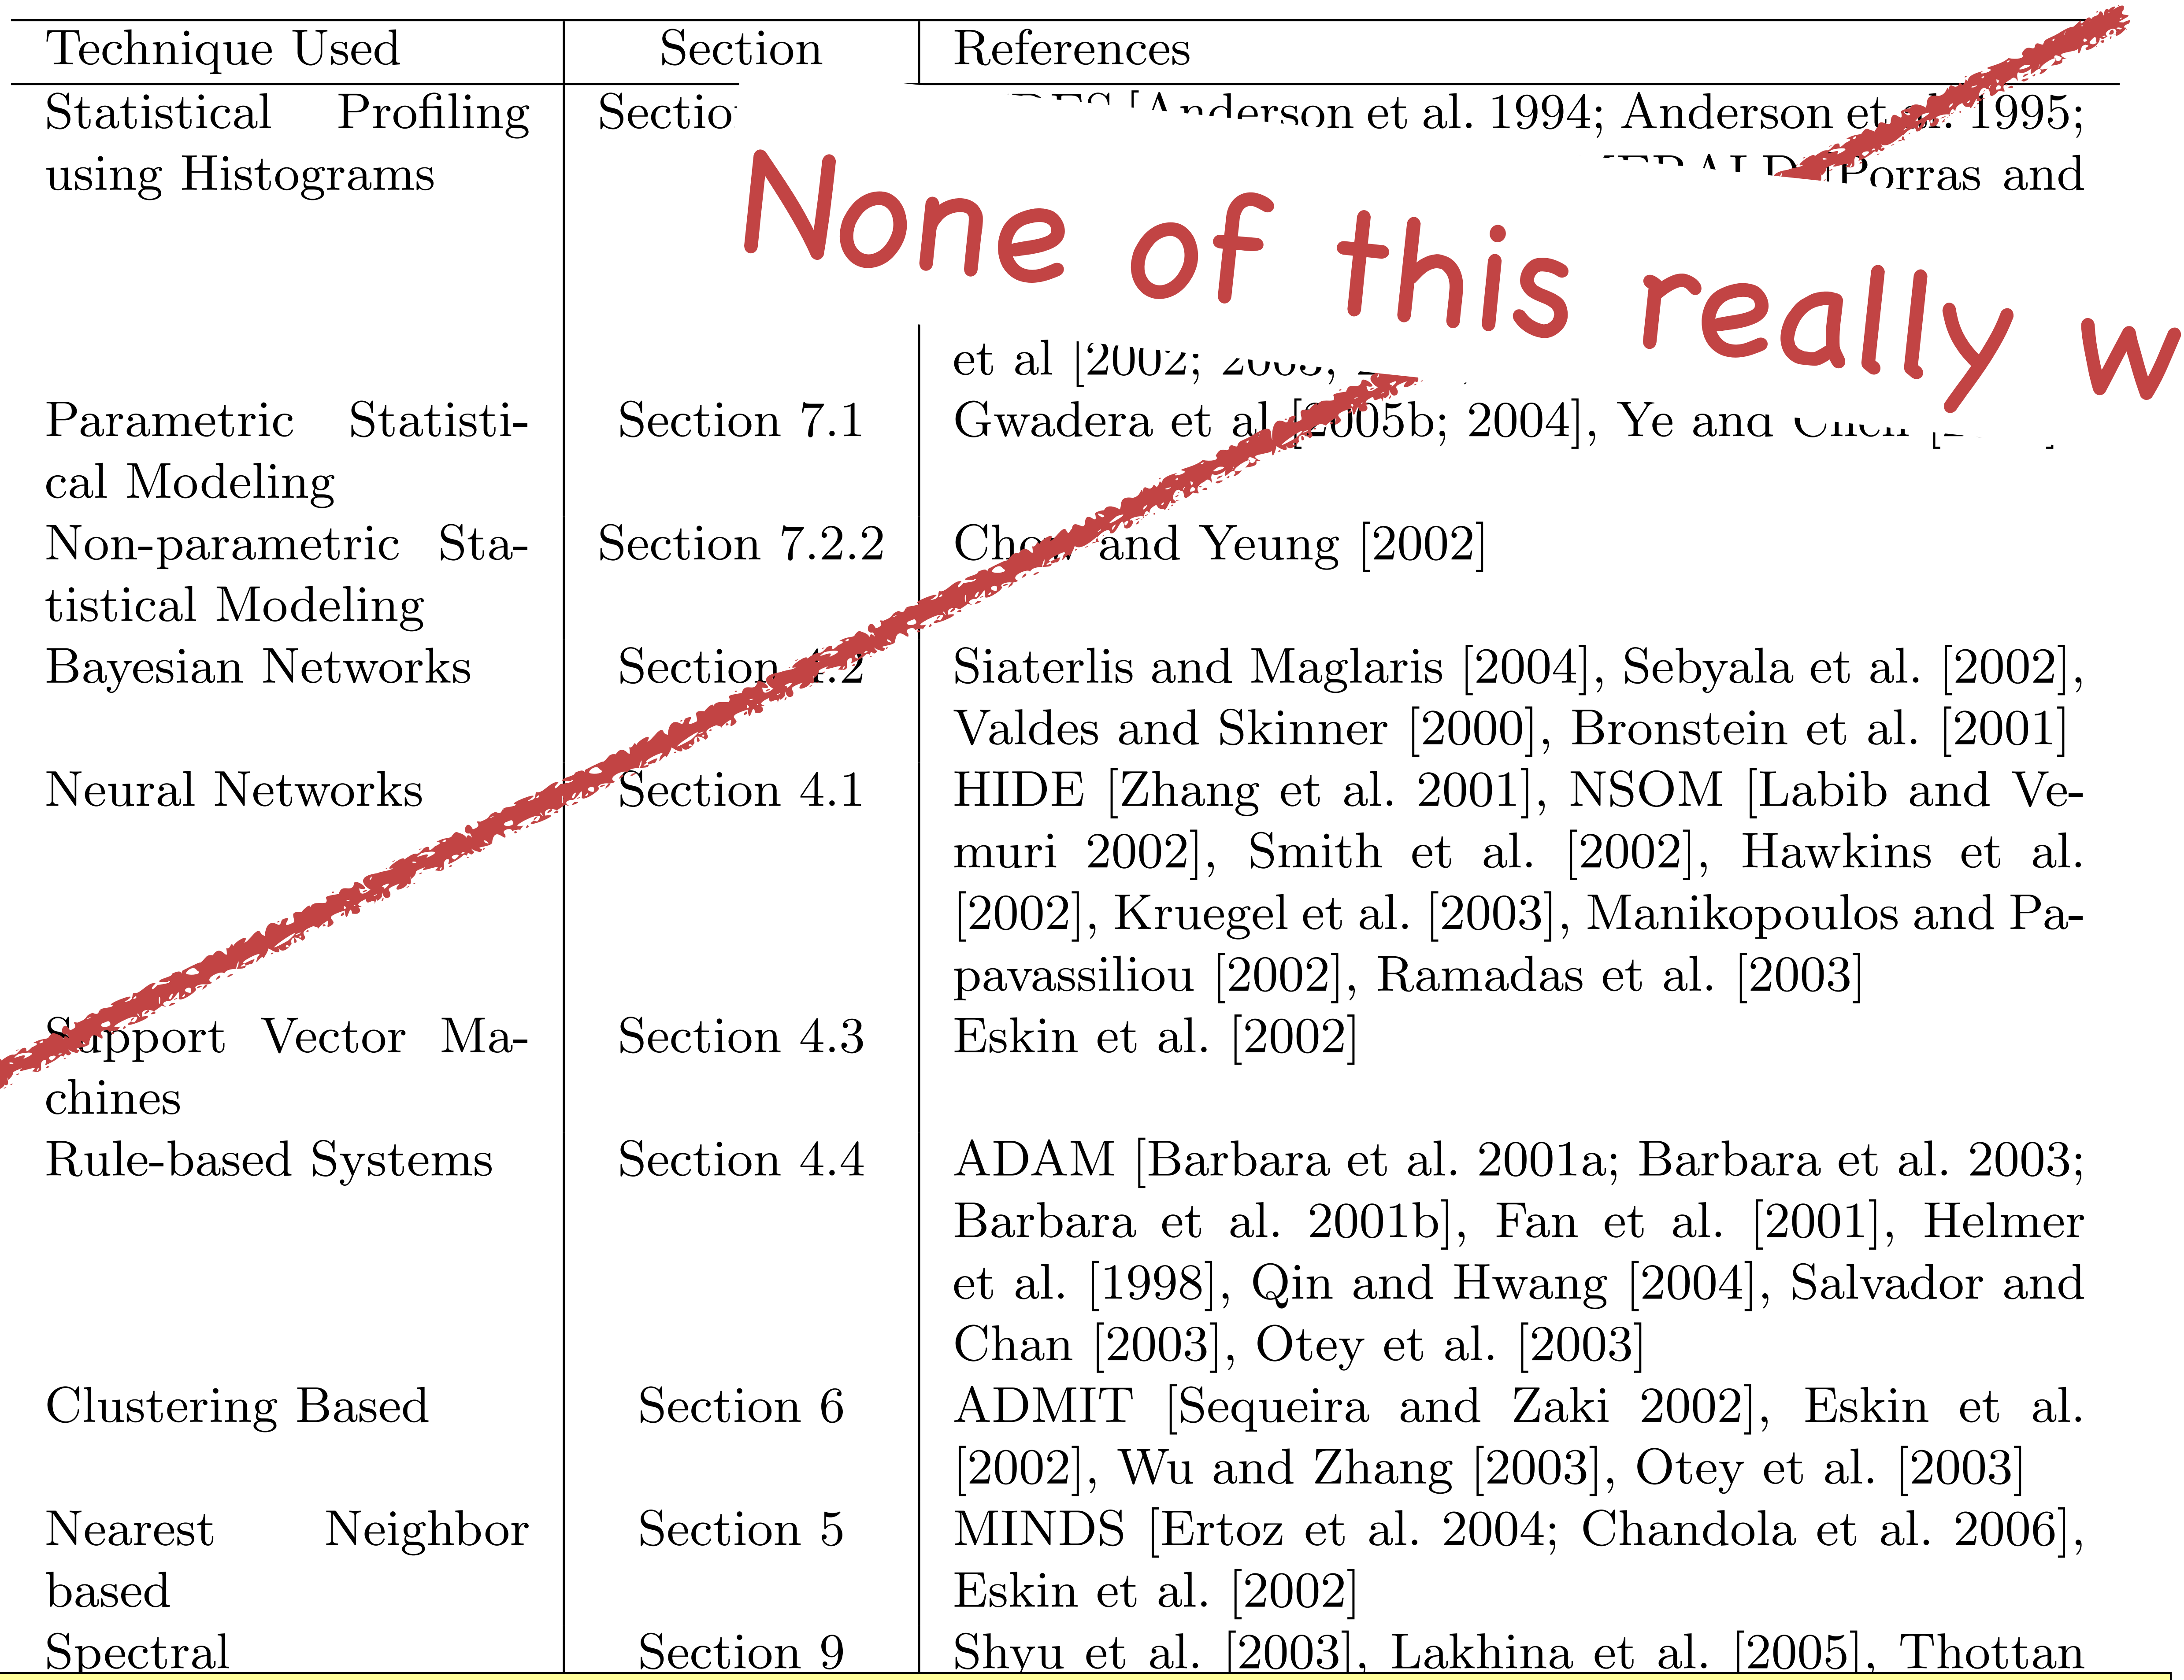

Technique Used	Section	References
Statistical Profiling using Histograms	Sectio	[illegible] [Anderson et al. 1994; Anderson et al. 1995; [illegible] [illegible] [Porras and [illegible] et al [2002; [illegible]
Parametric Statistical Modeling	Section 7.1	Gwadera et al [2005b; 2004], Ye and [illegible]
Non-parametric Statistical Modeling	Section 7.2.2	Chow and Yeung [2002]
Bayesian Networks	Section 4.2	Siaterlis and Maglaris [2004], Sebyala et al. [2002], Valdes and Skinner [2000], Bronstein et al. [2001]
Neural Networks	Section 4.1	HIDE [Zhang et al. 2001], NSOM [Labib and Vemuri 2002], Smith et al. [2002], Hawkins et al. [2002], Kruegel et al. [2003], Manikopoulos and Papavassiliou [2002], Ramadas et al. [2003]
Support Vector Machines	Section 4.3	Eskin et al. [2002]
Rule-based Systems	Section 4.4	ADAM [Barbara et al. 2001a; Barbara et al. 2003; Barbara et al. 2001b], Fan et al. [2001], Helmer et al. [1998], Qin and Hwang [2004], Salvador and Chan [2003], Otey et al. [2003]
Clustering Based	Section 6	ADMIT [Sequeira and Zaki 2002], Eskin et al. [2002], Wu and Zhang [2003], Otey et al. [2003]
Nearest Neighbor based	Section 5	MINDS [Ertoz et al. 2004; Chandola et al. 2006], Eskin et al. [2002]
Spectral	Section 9	Shyu et al. [2003], Lakhina et al. [2005], Thottan

None of this really works

Why is machine learning so ineffective in this domain?

Machine learning in other domains

Machine learning in other domains

Machine Translation

Optical Character Recognition

Product Recommendations

Spam Detection

Classification Problems

Machine learning in other domains

Machine Translation

Optical Character Recognition

Product Recommendations

Spam Detection

Classification Problems

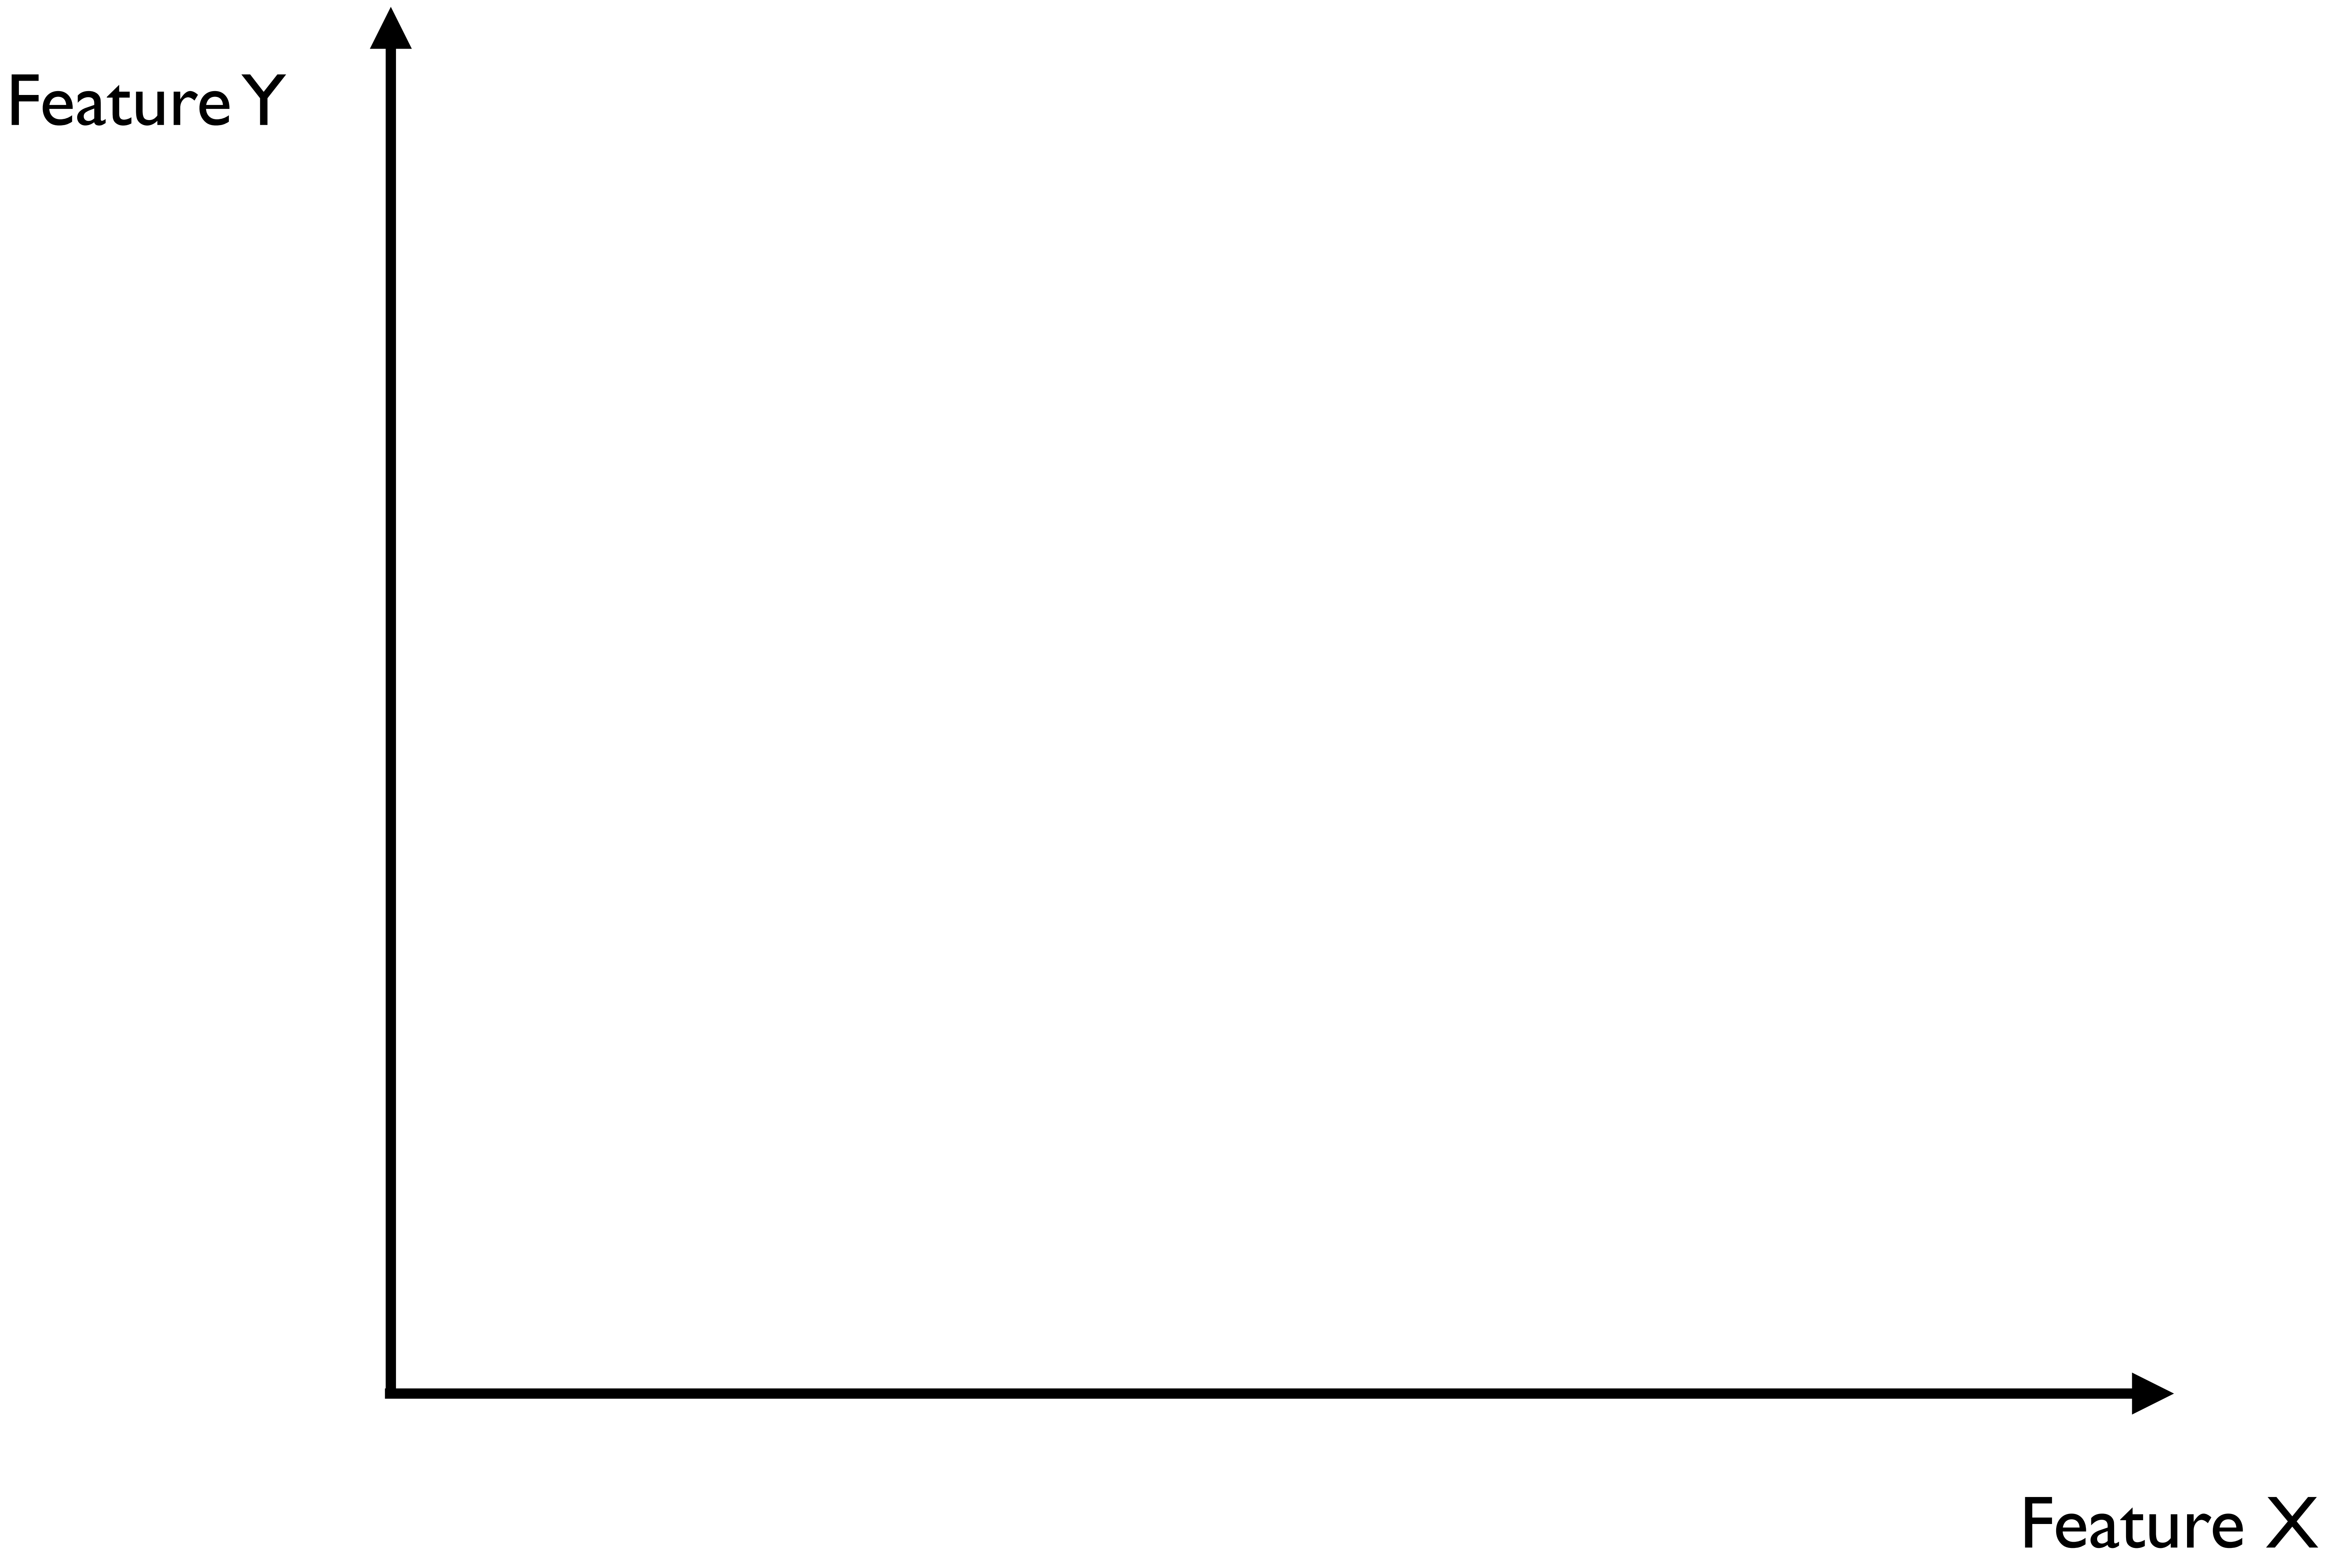

Machine learning in other domains

Machine Translation

Optical Character Recognition

Product Recommendations

Spam Detection

Classification Problems

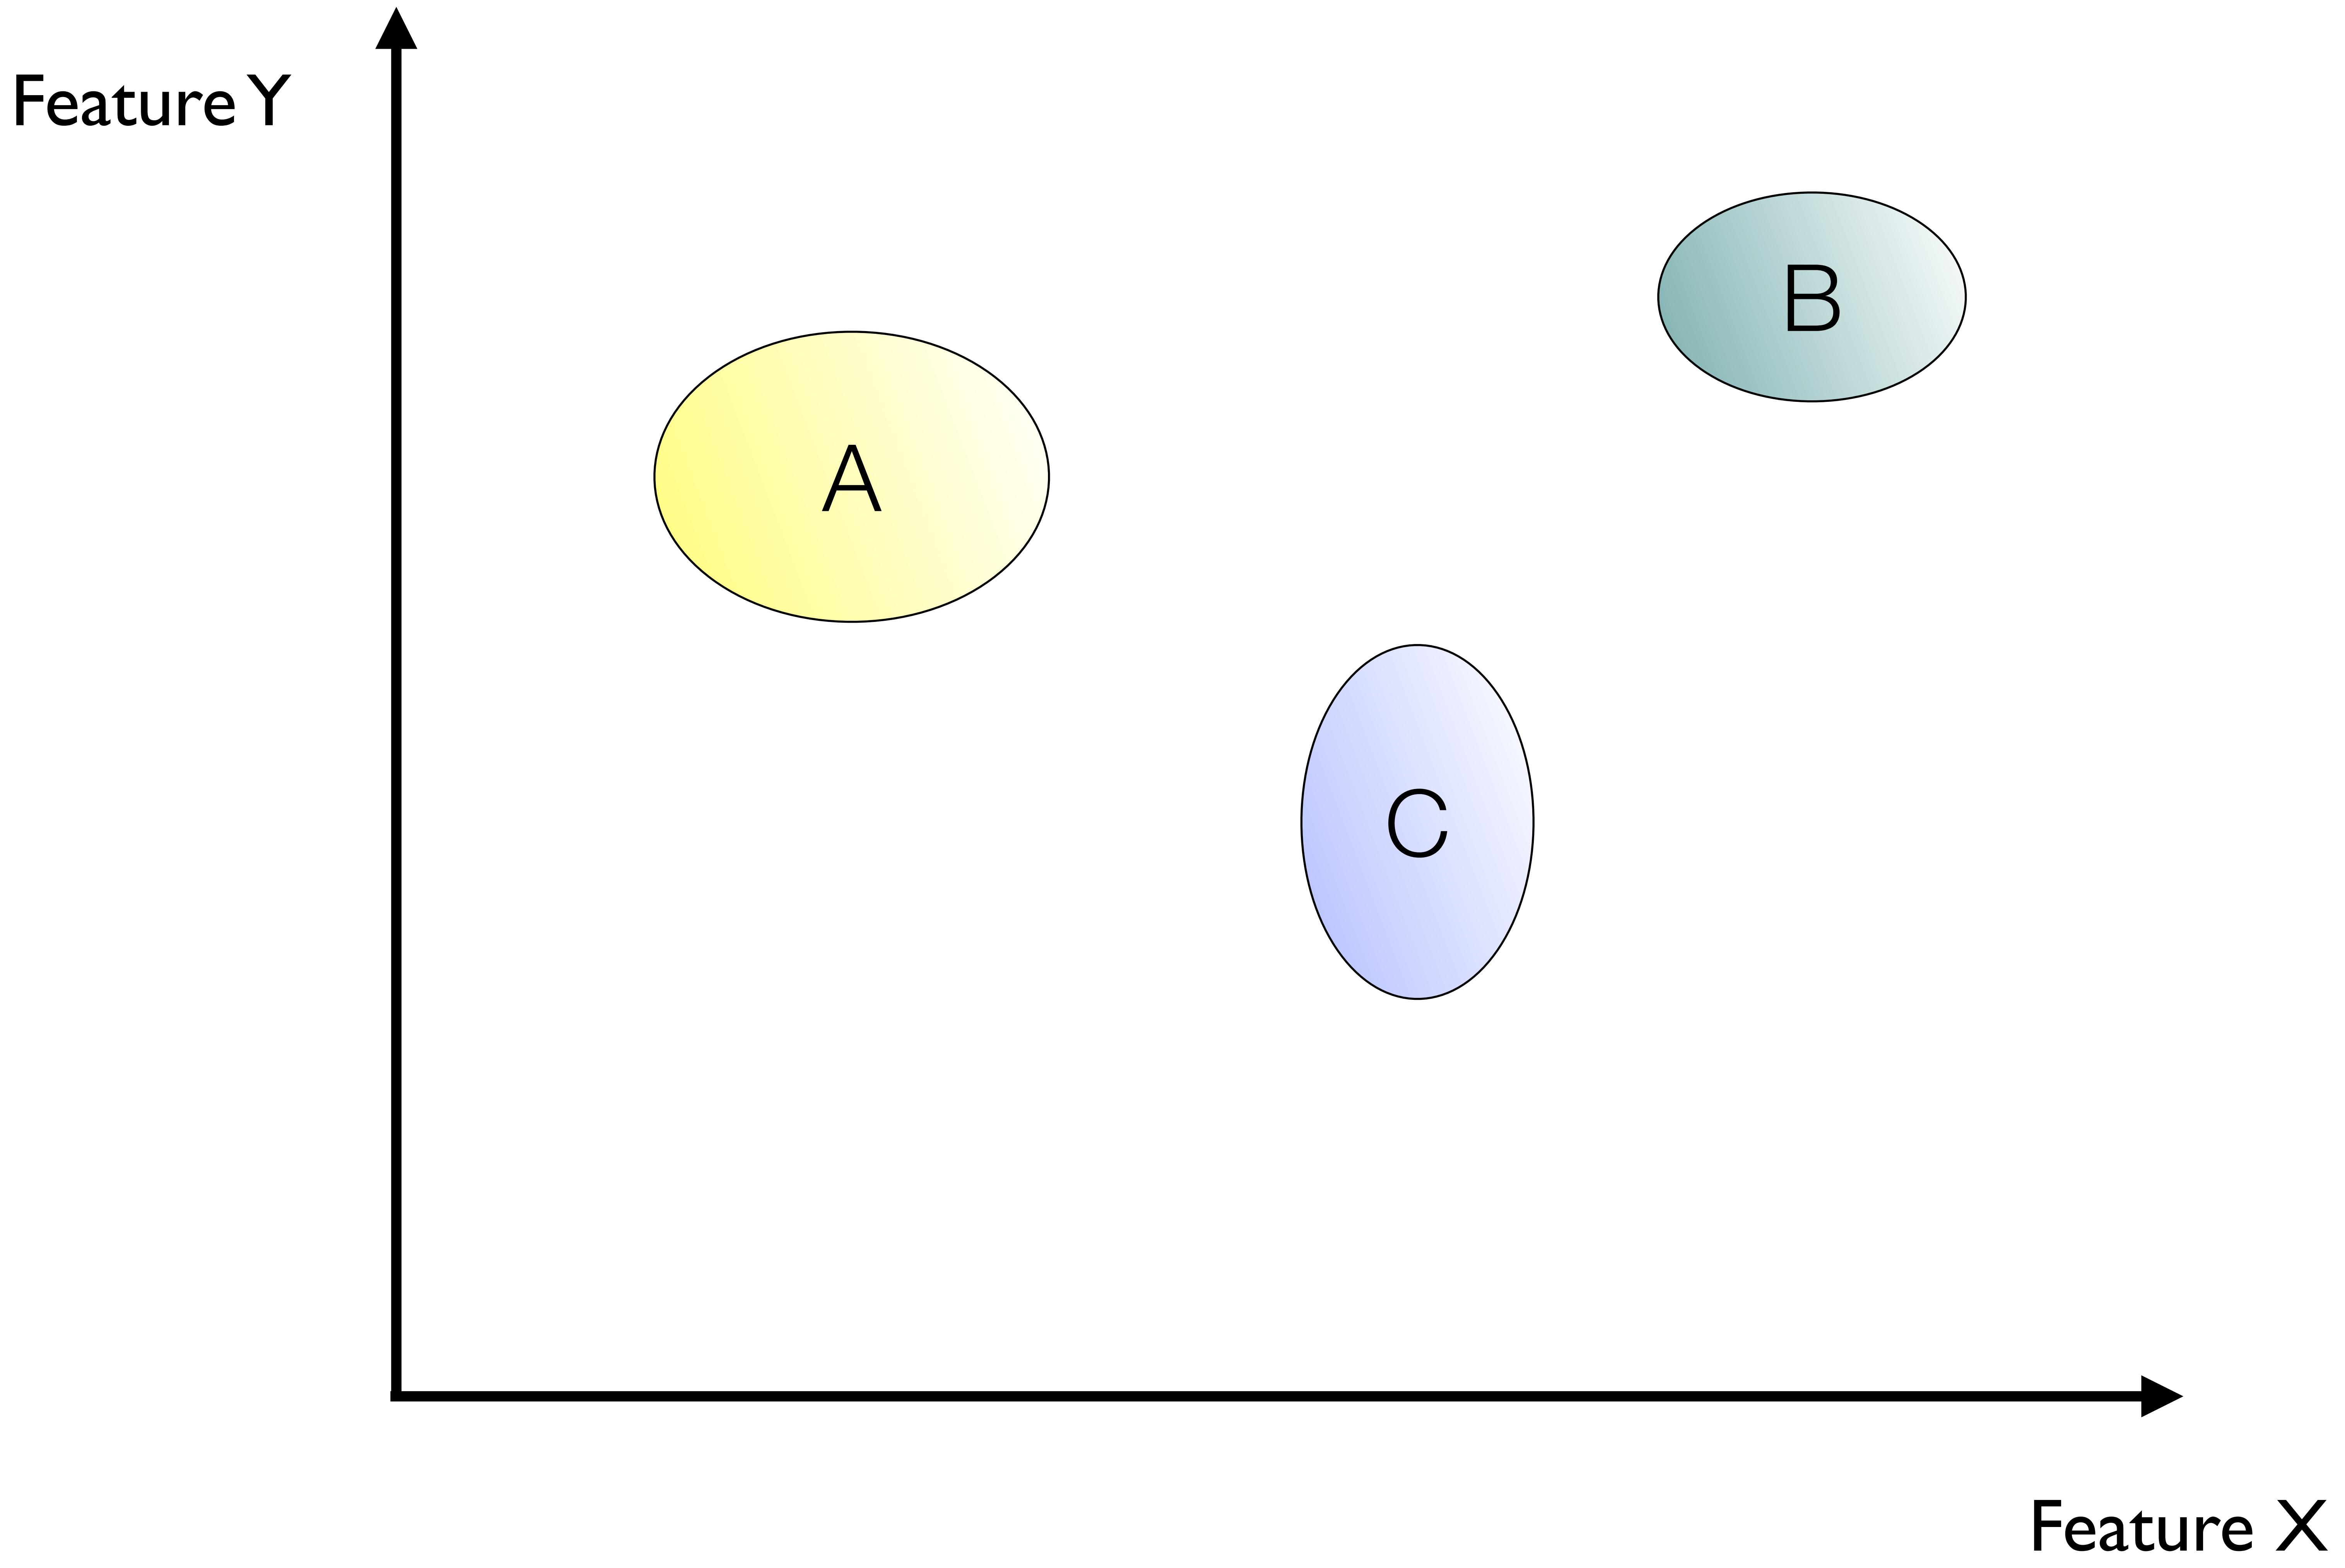

Machine learning in other domains

Machine Translation

Optical Character Recognition

Product Recommendations

Spam Detection

Classification Problems

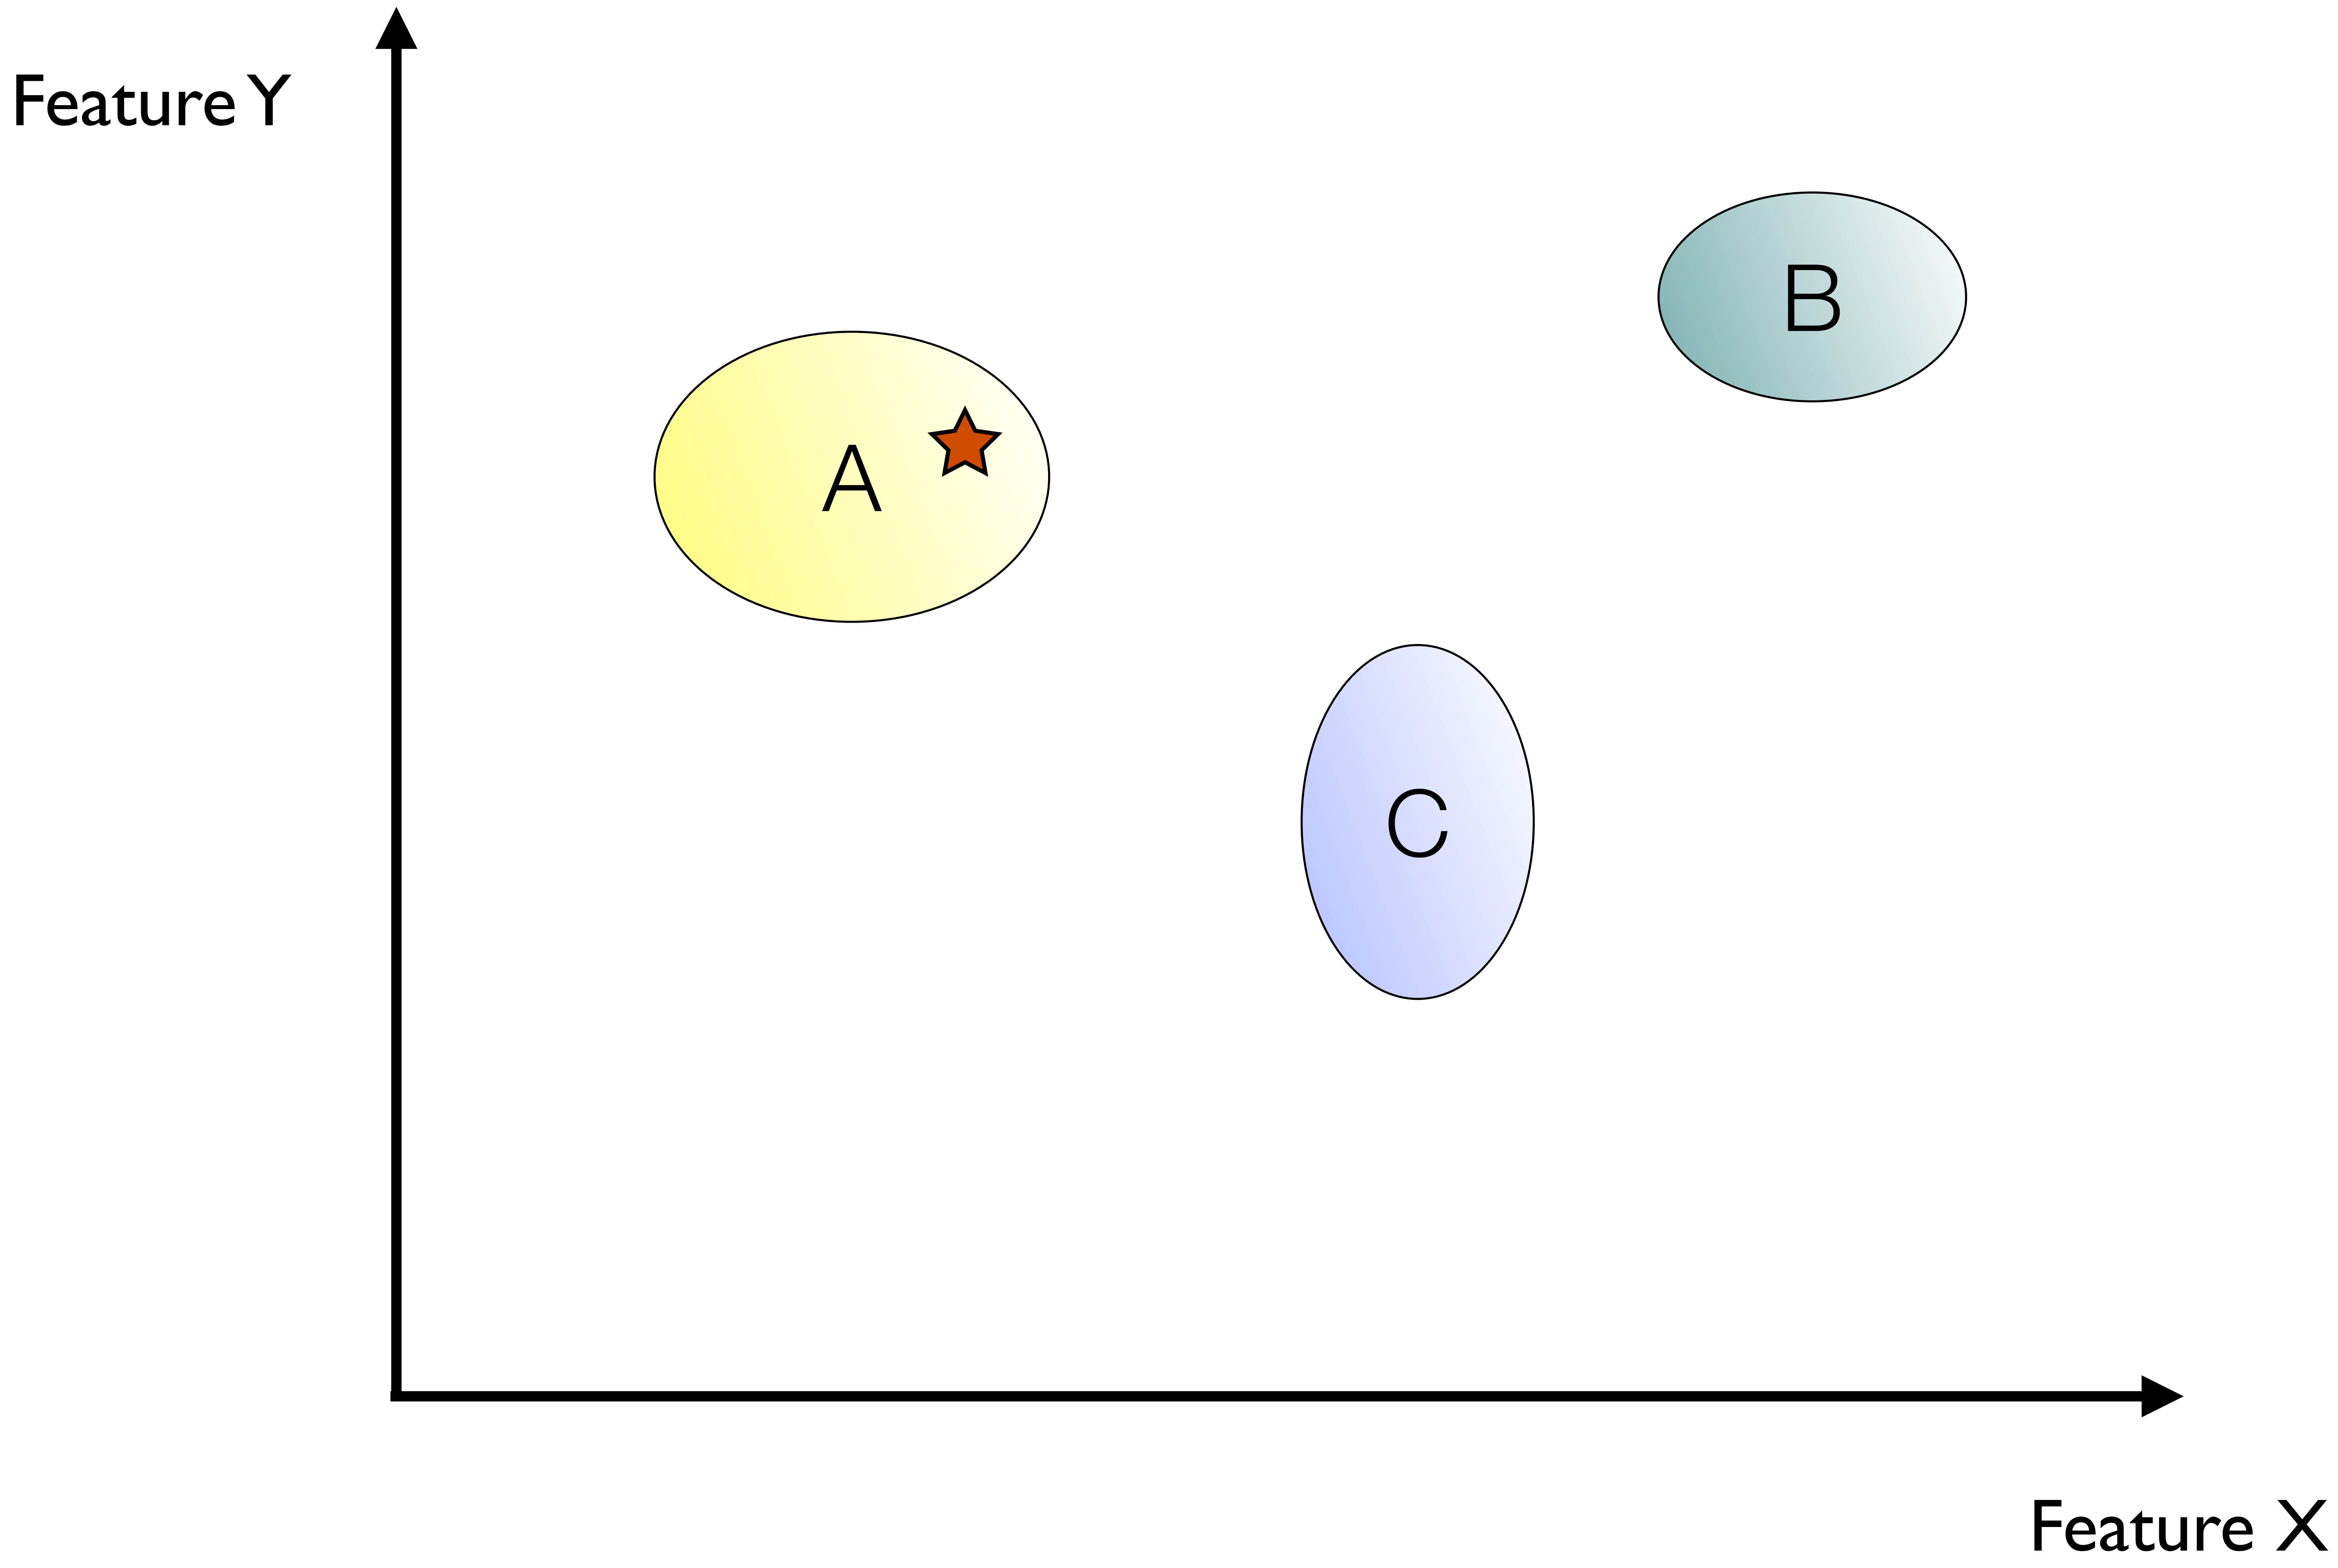

Machine learning in other domains

Machine Translation

Optical Character Recognition

Product Recommendations

Spam Detection

Classification Problems

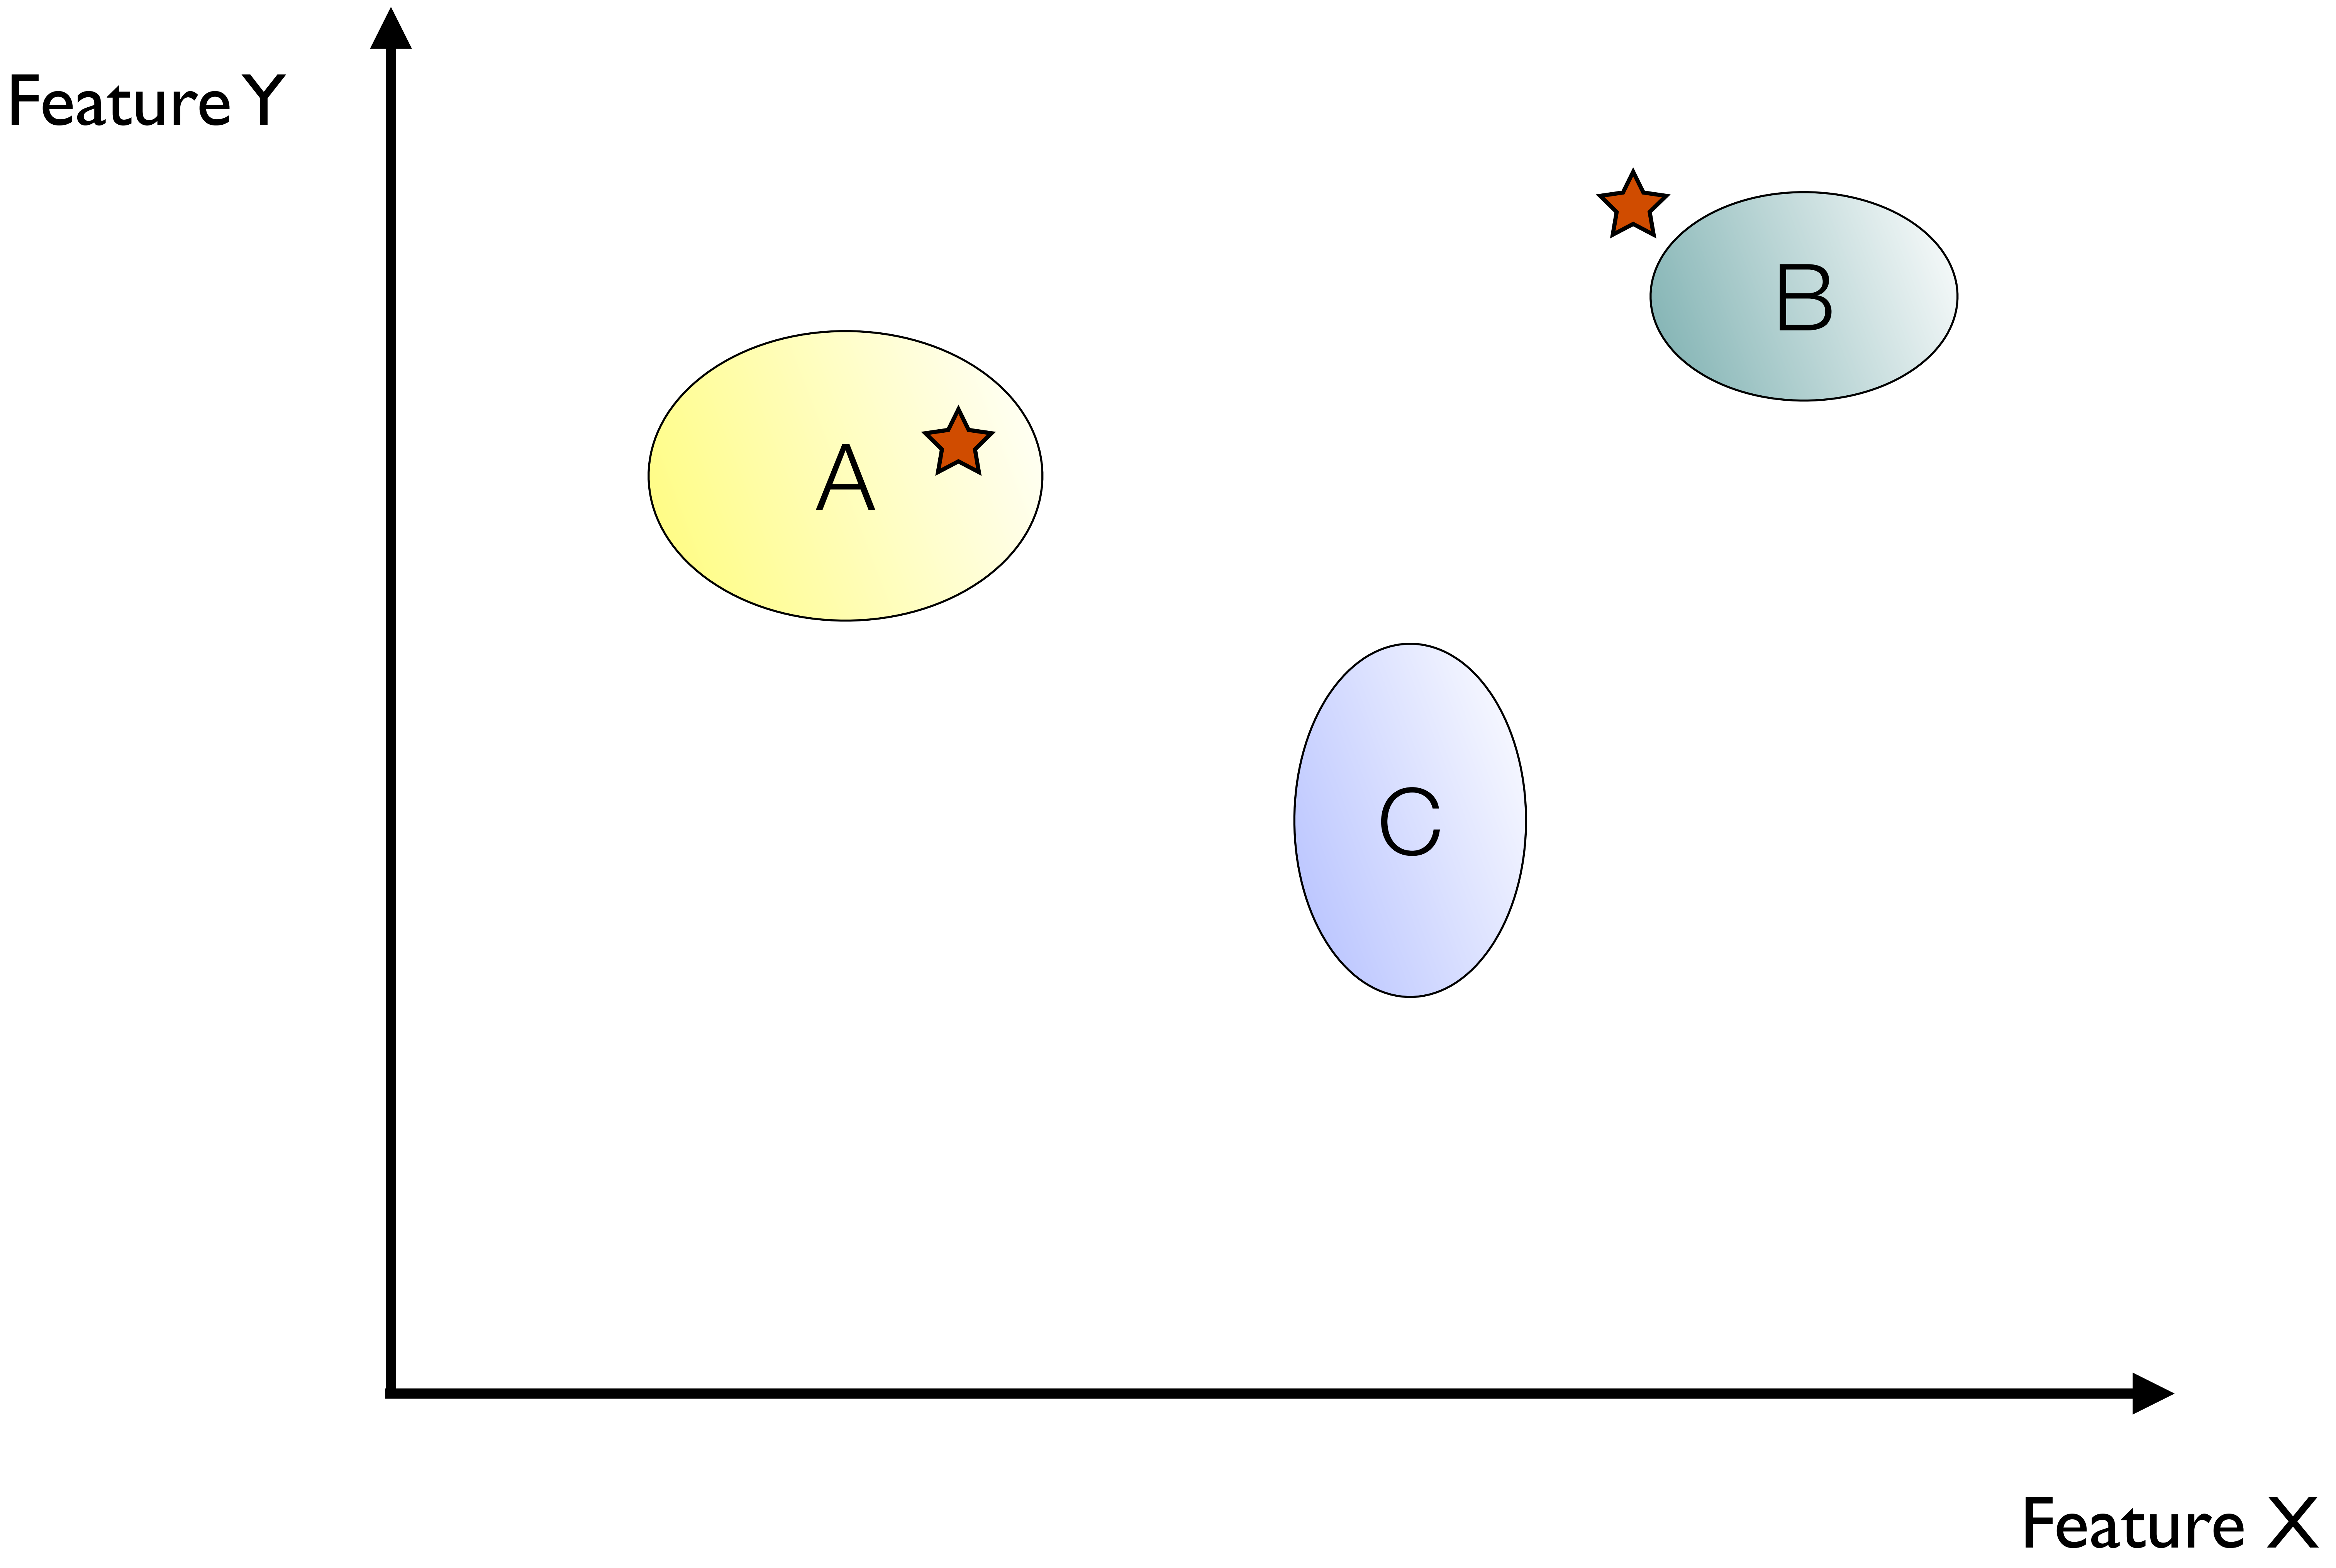

Machine learning in other domains

Machine Translation

Optical Character Recognition

Product Recommendations

Spam Detection

Classification Problems

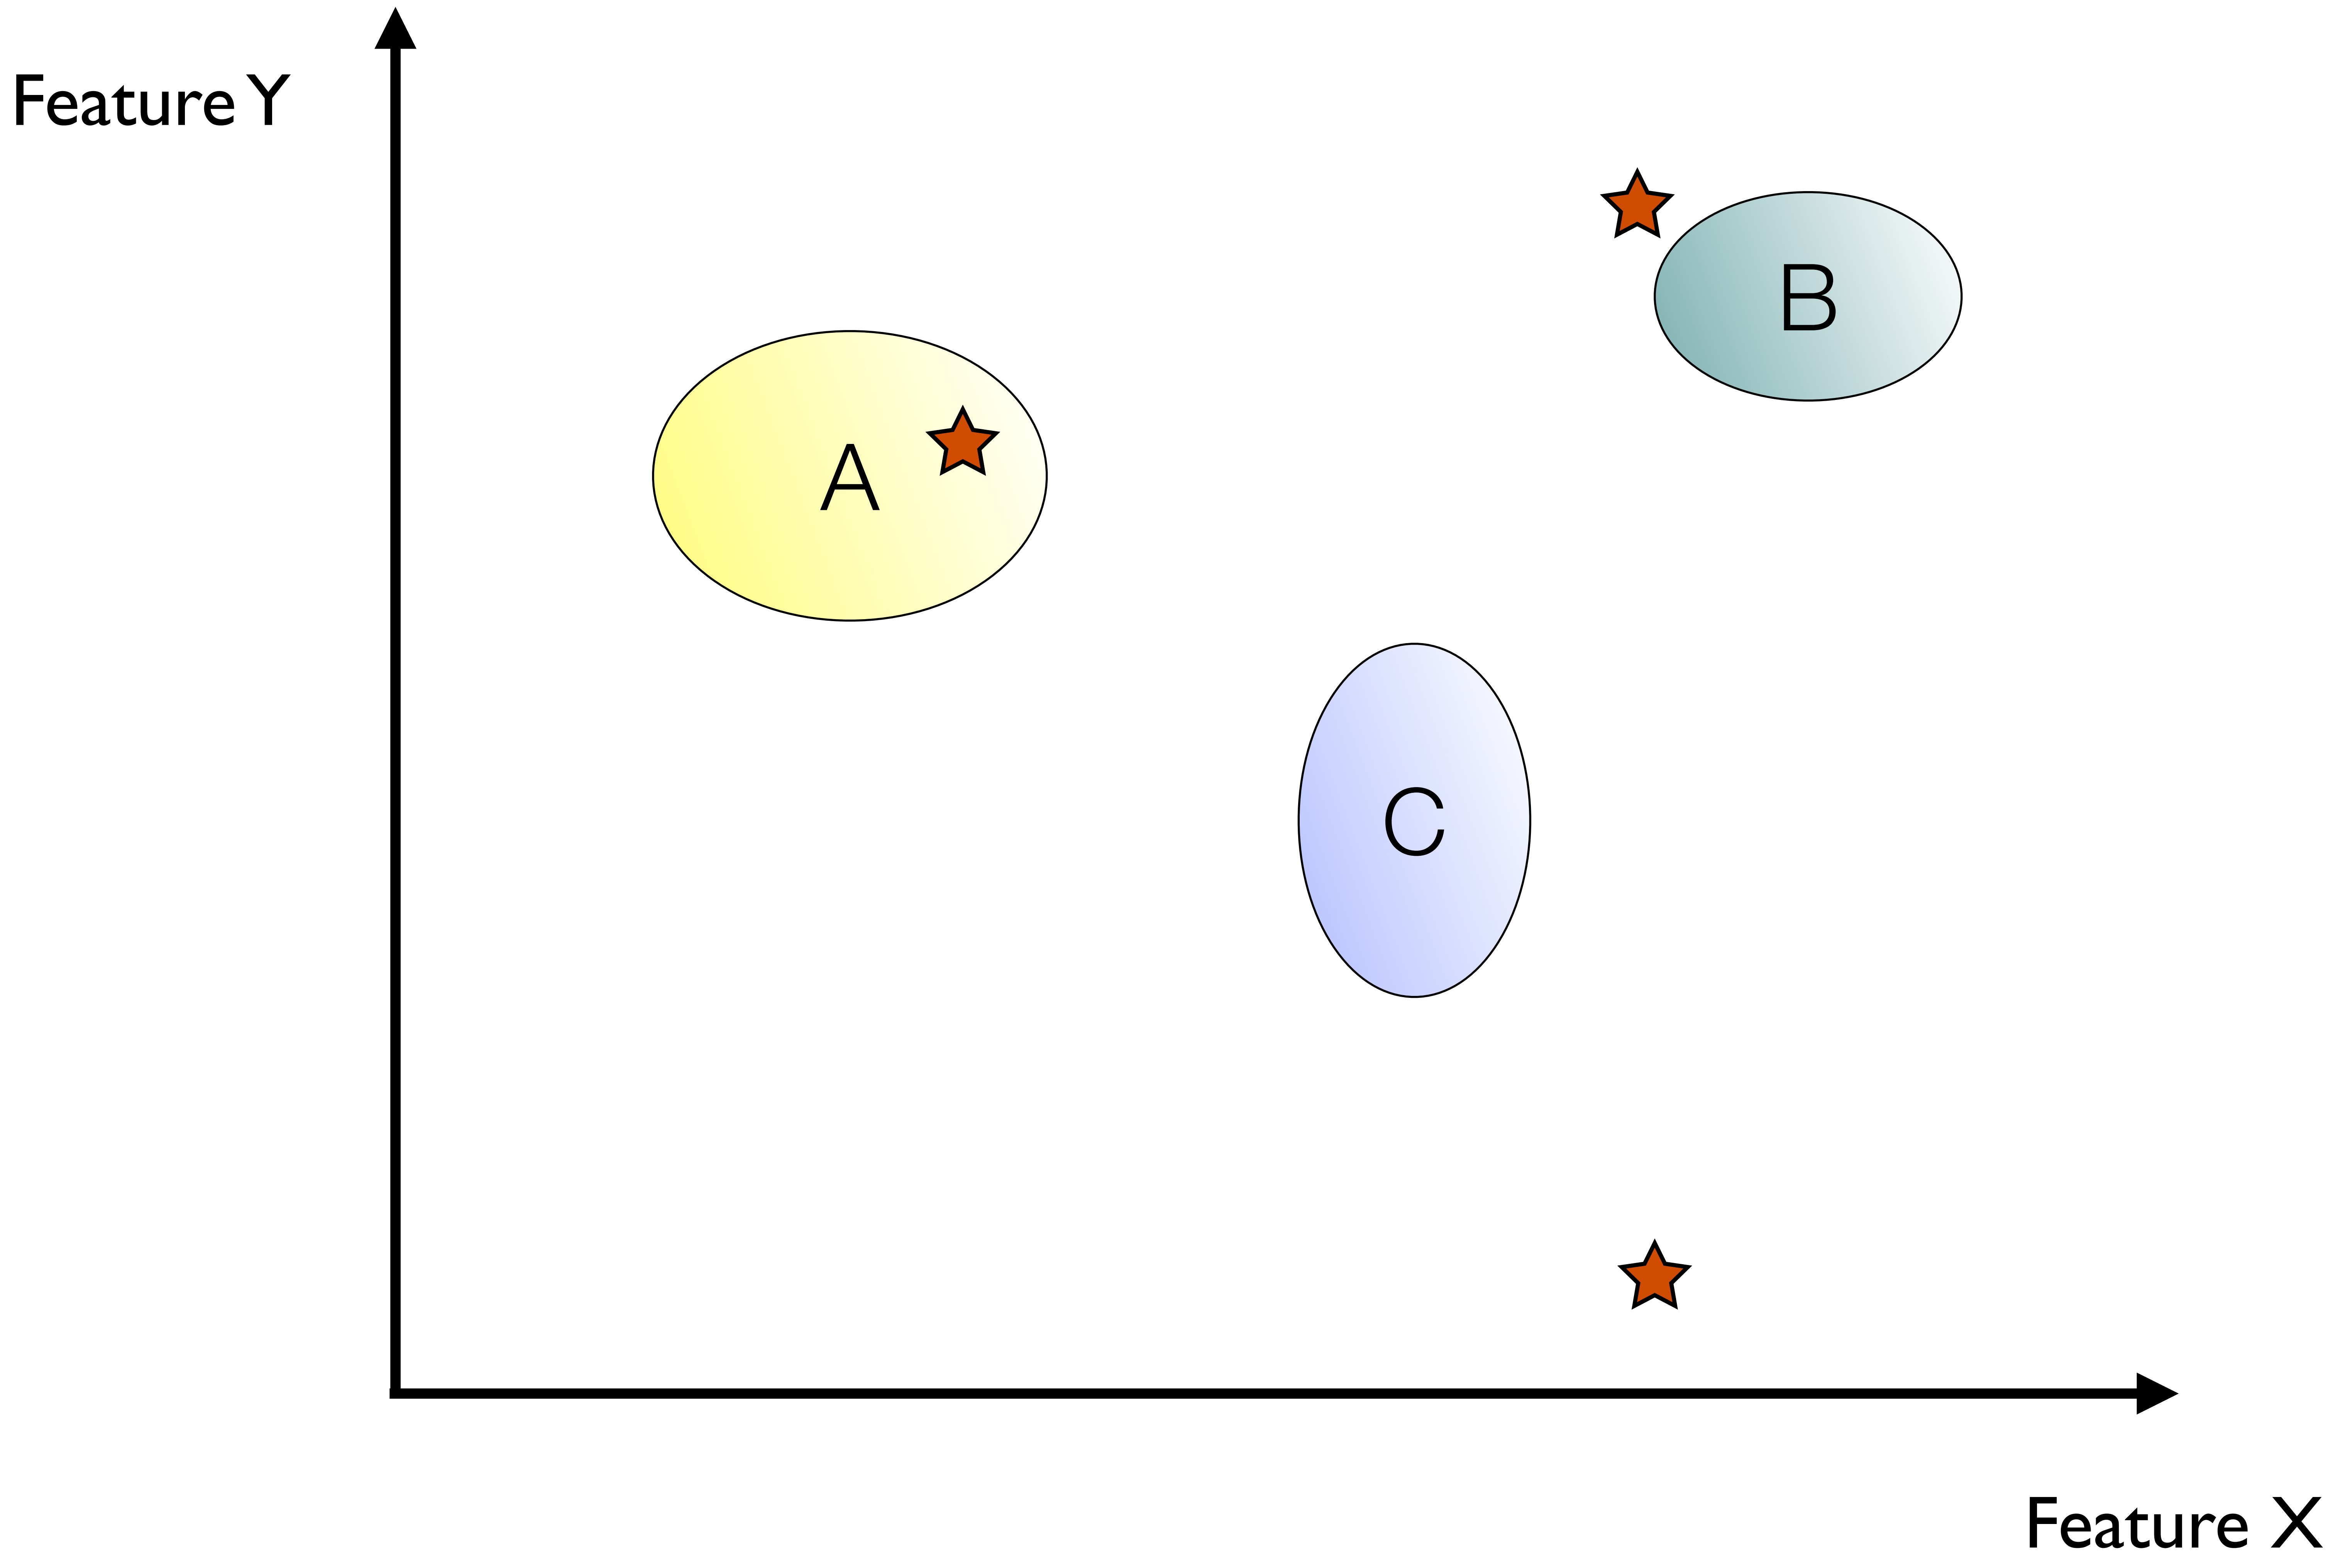

Machine learning in other domains

Machine Translation

Optical Character Recognition

Product Recommendations

Spam Detection

Classification Problems

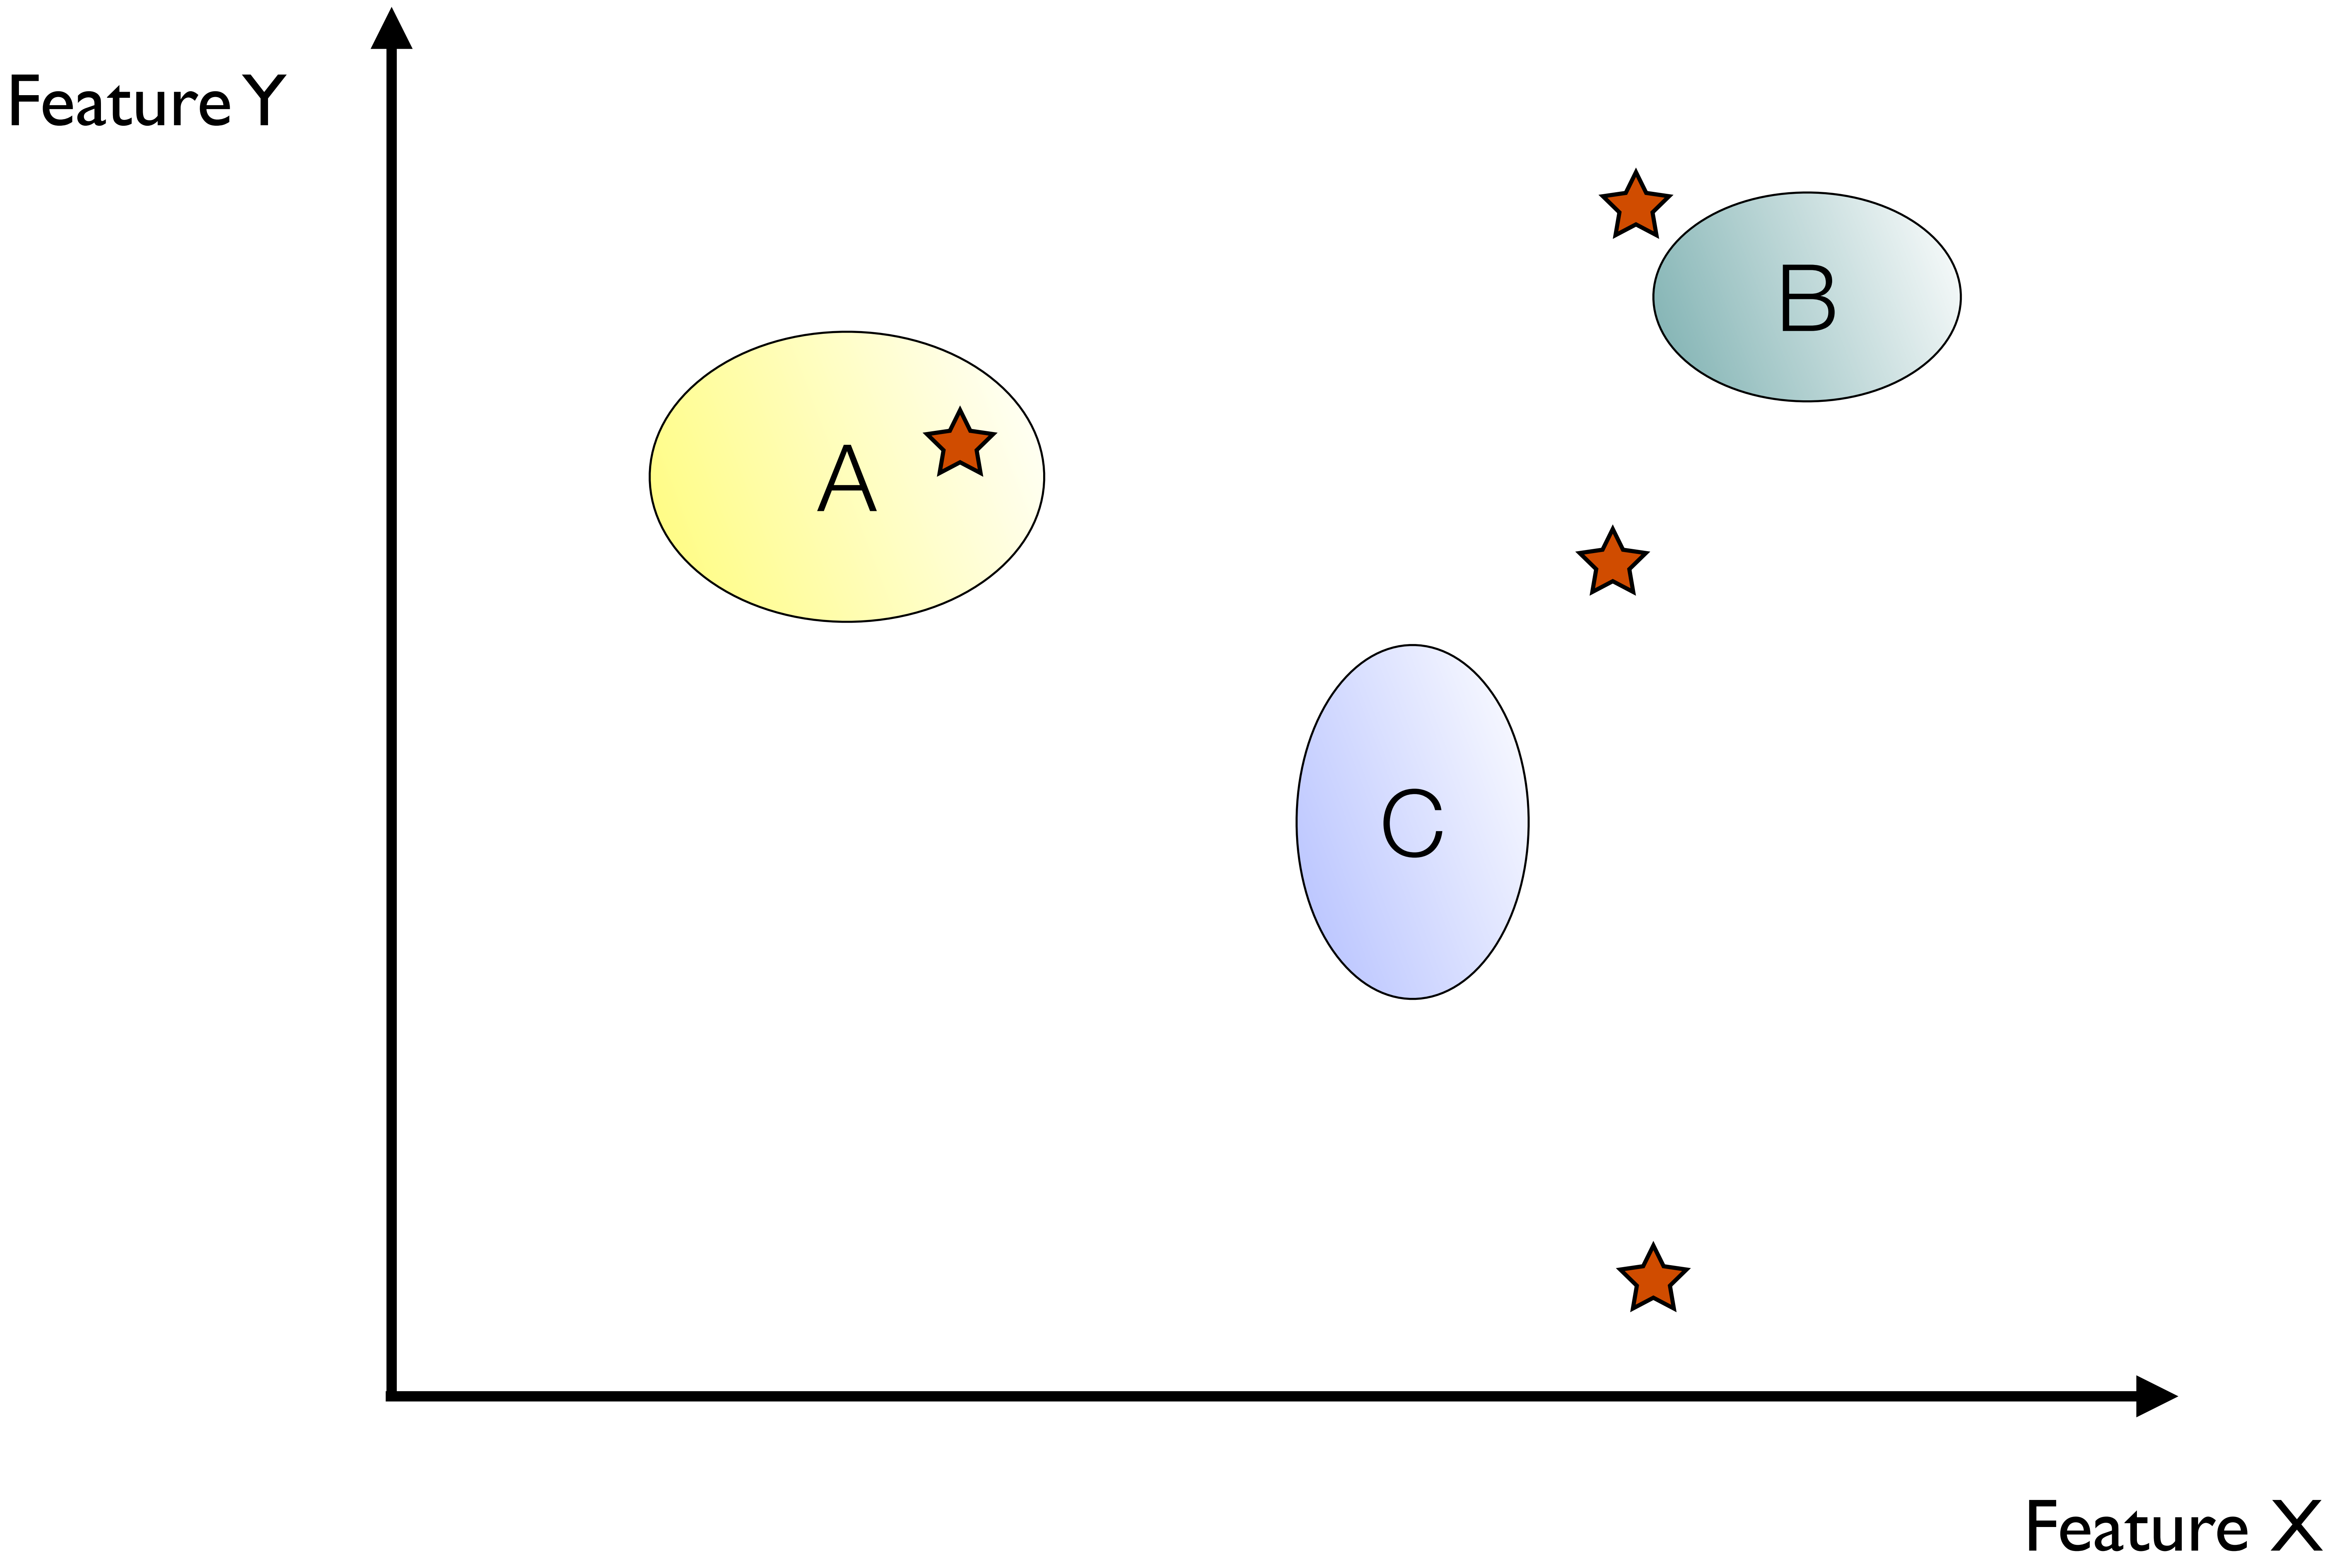

Machine learning in other domains

Machine Translation

Optical Character Recognition

Product Recommendations

Spam Detection

Classification Problems

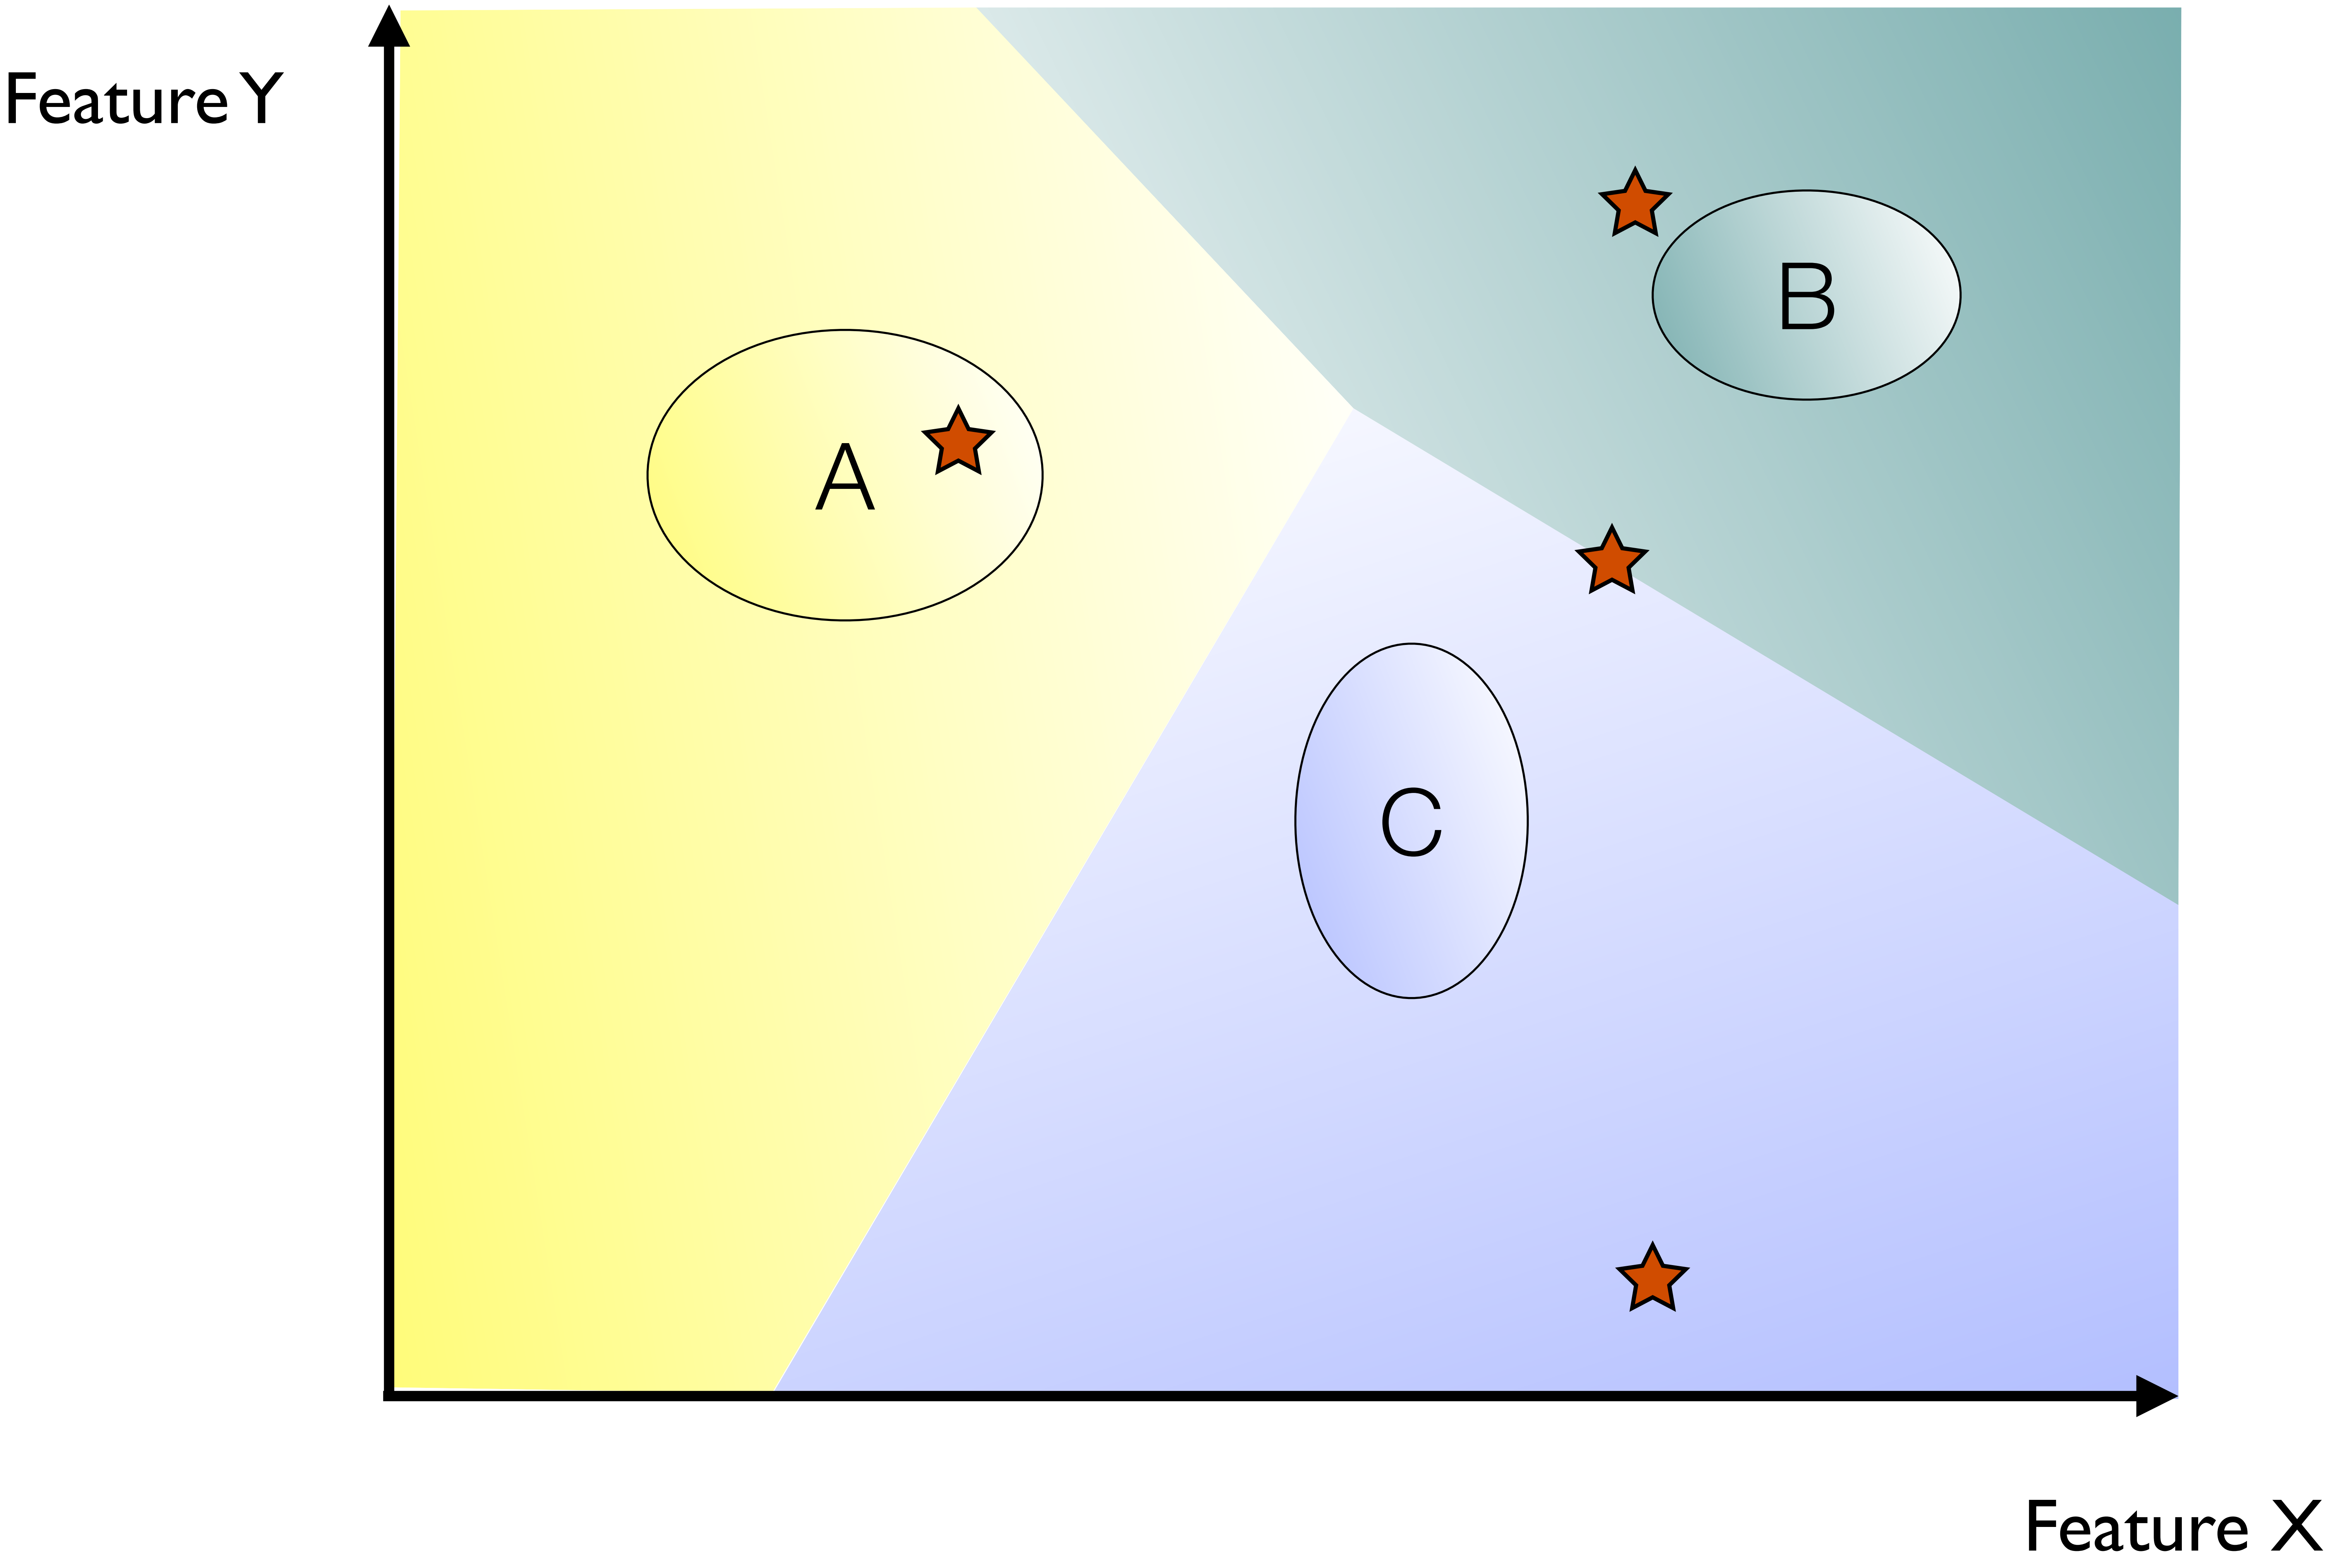

Machine learning in other domains

Machine Translation

Optical Character Recognition

Product Recommendations

Spam Detection

Classification Problems

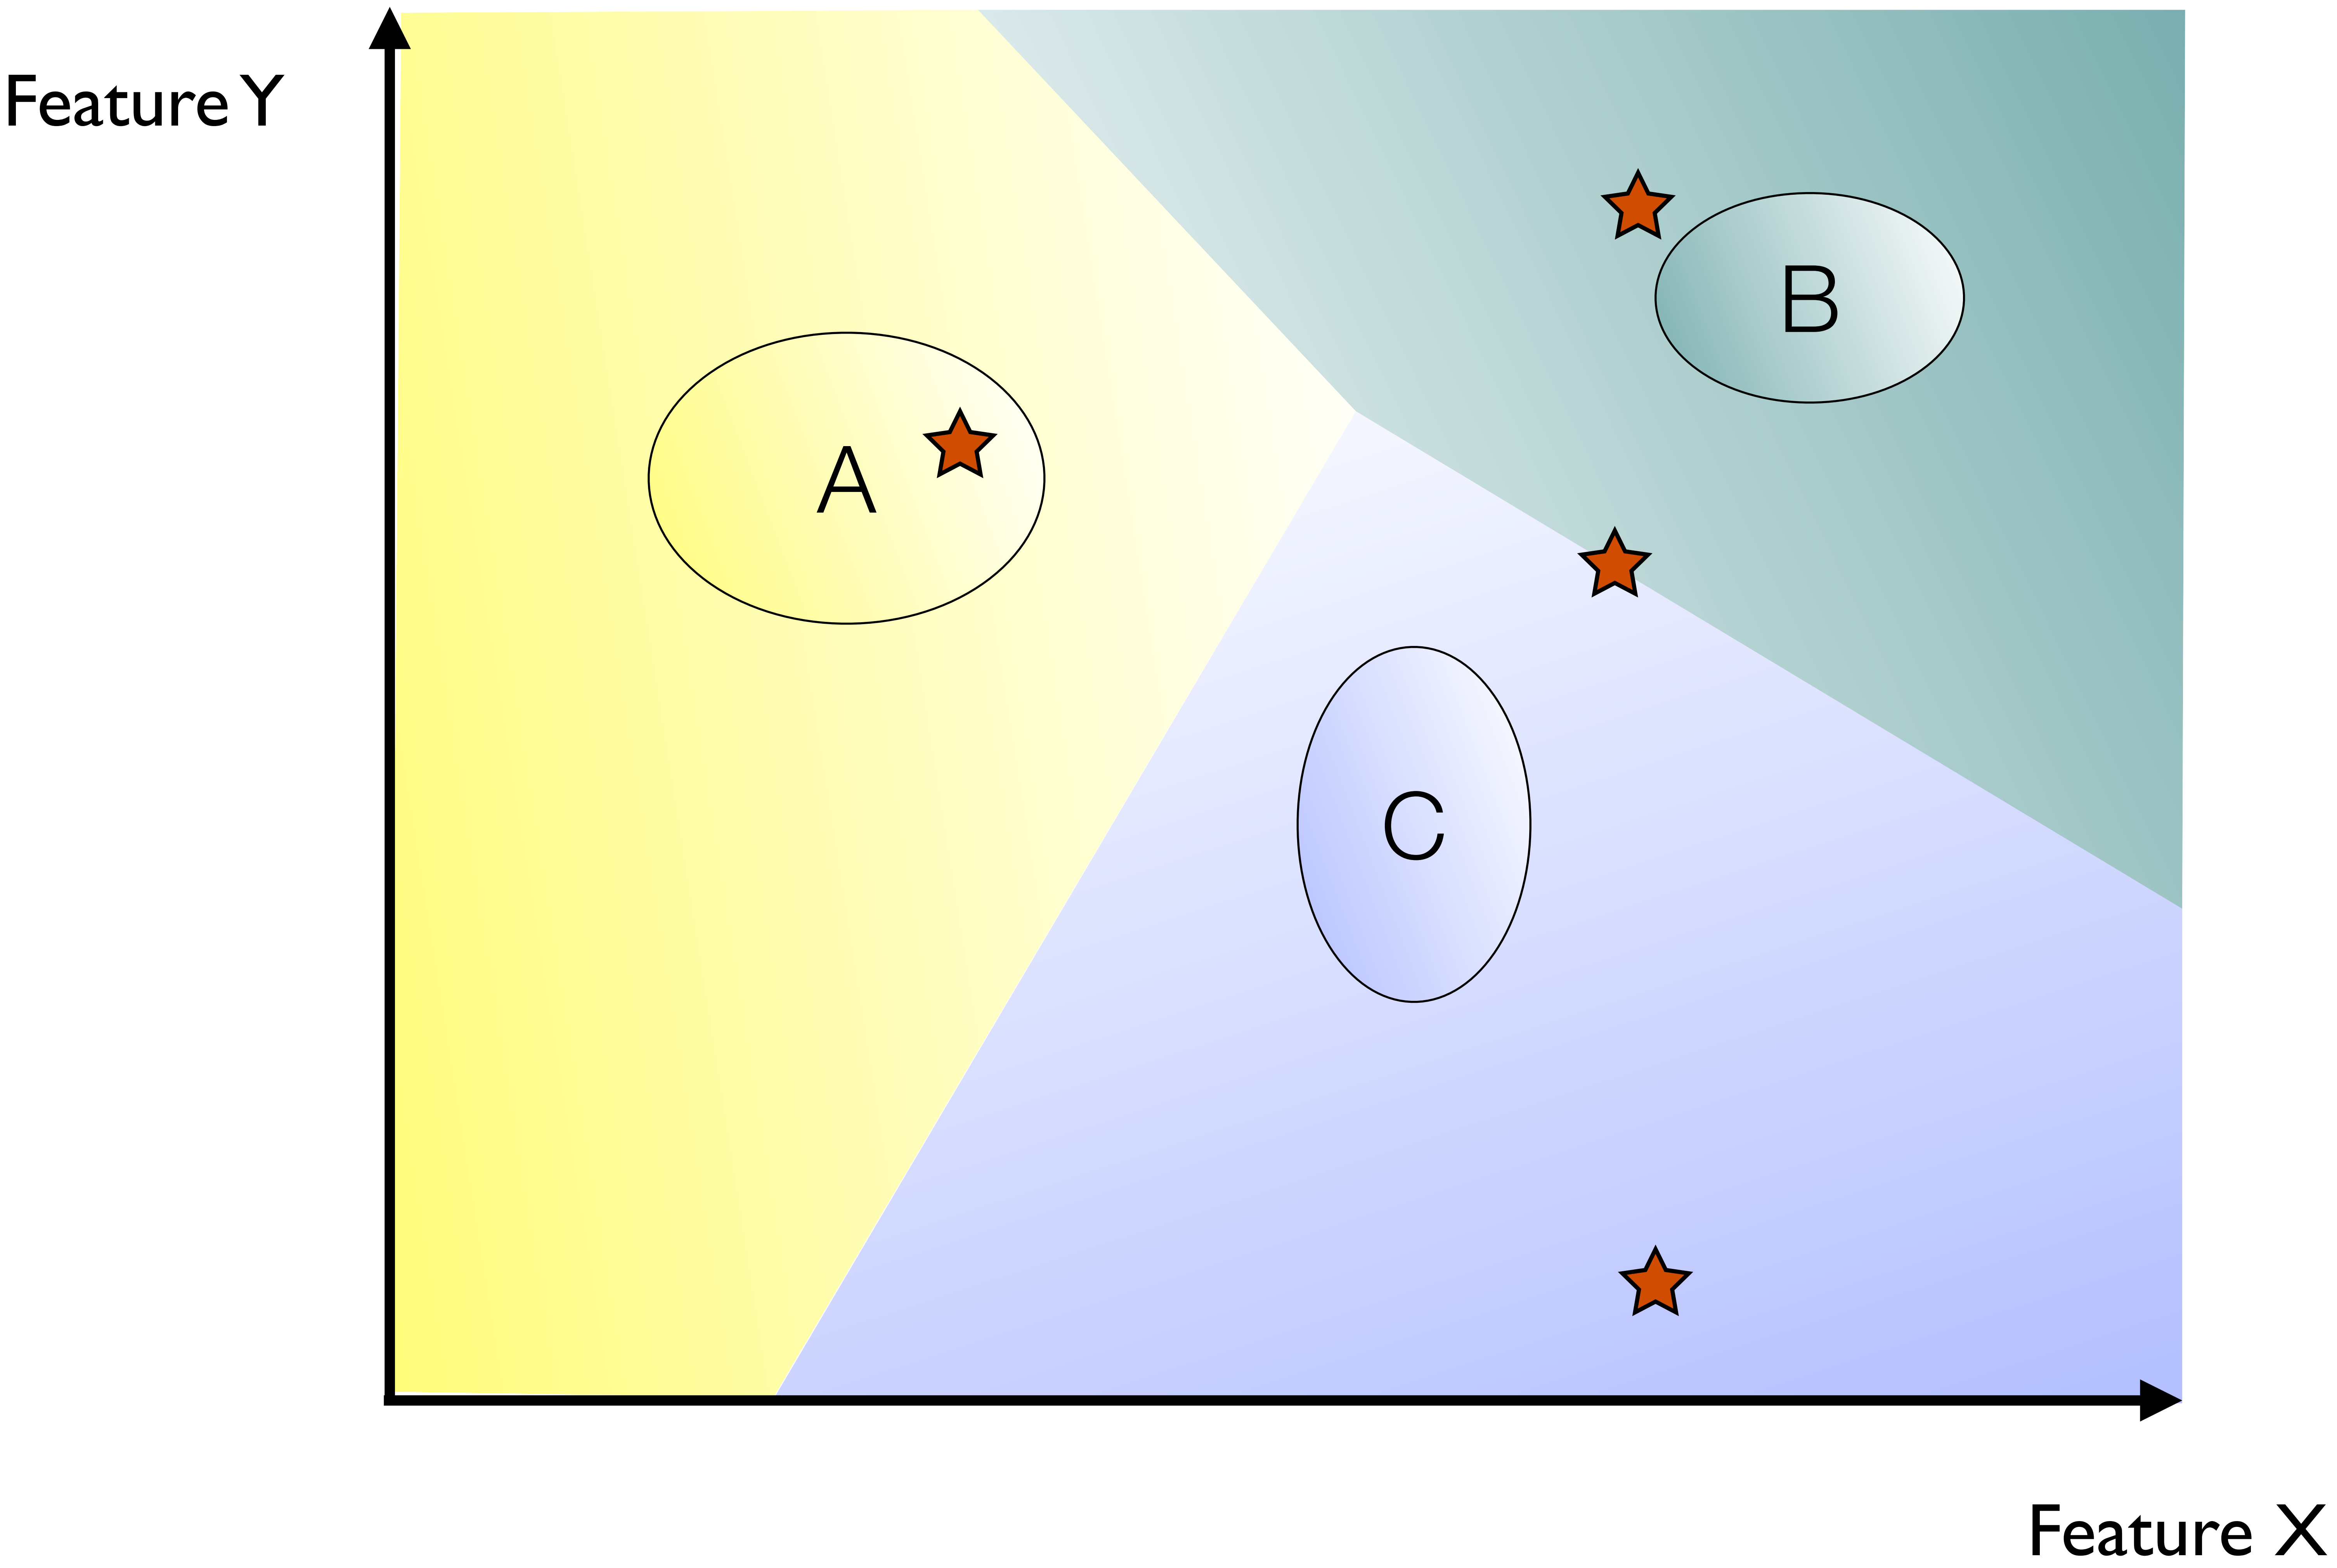

Trained with specimen of all categories -> Very robust even at scale

Outlier detection

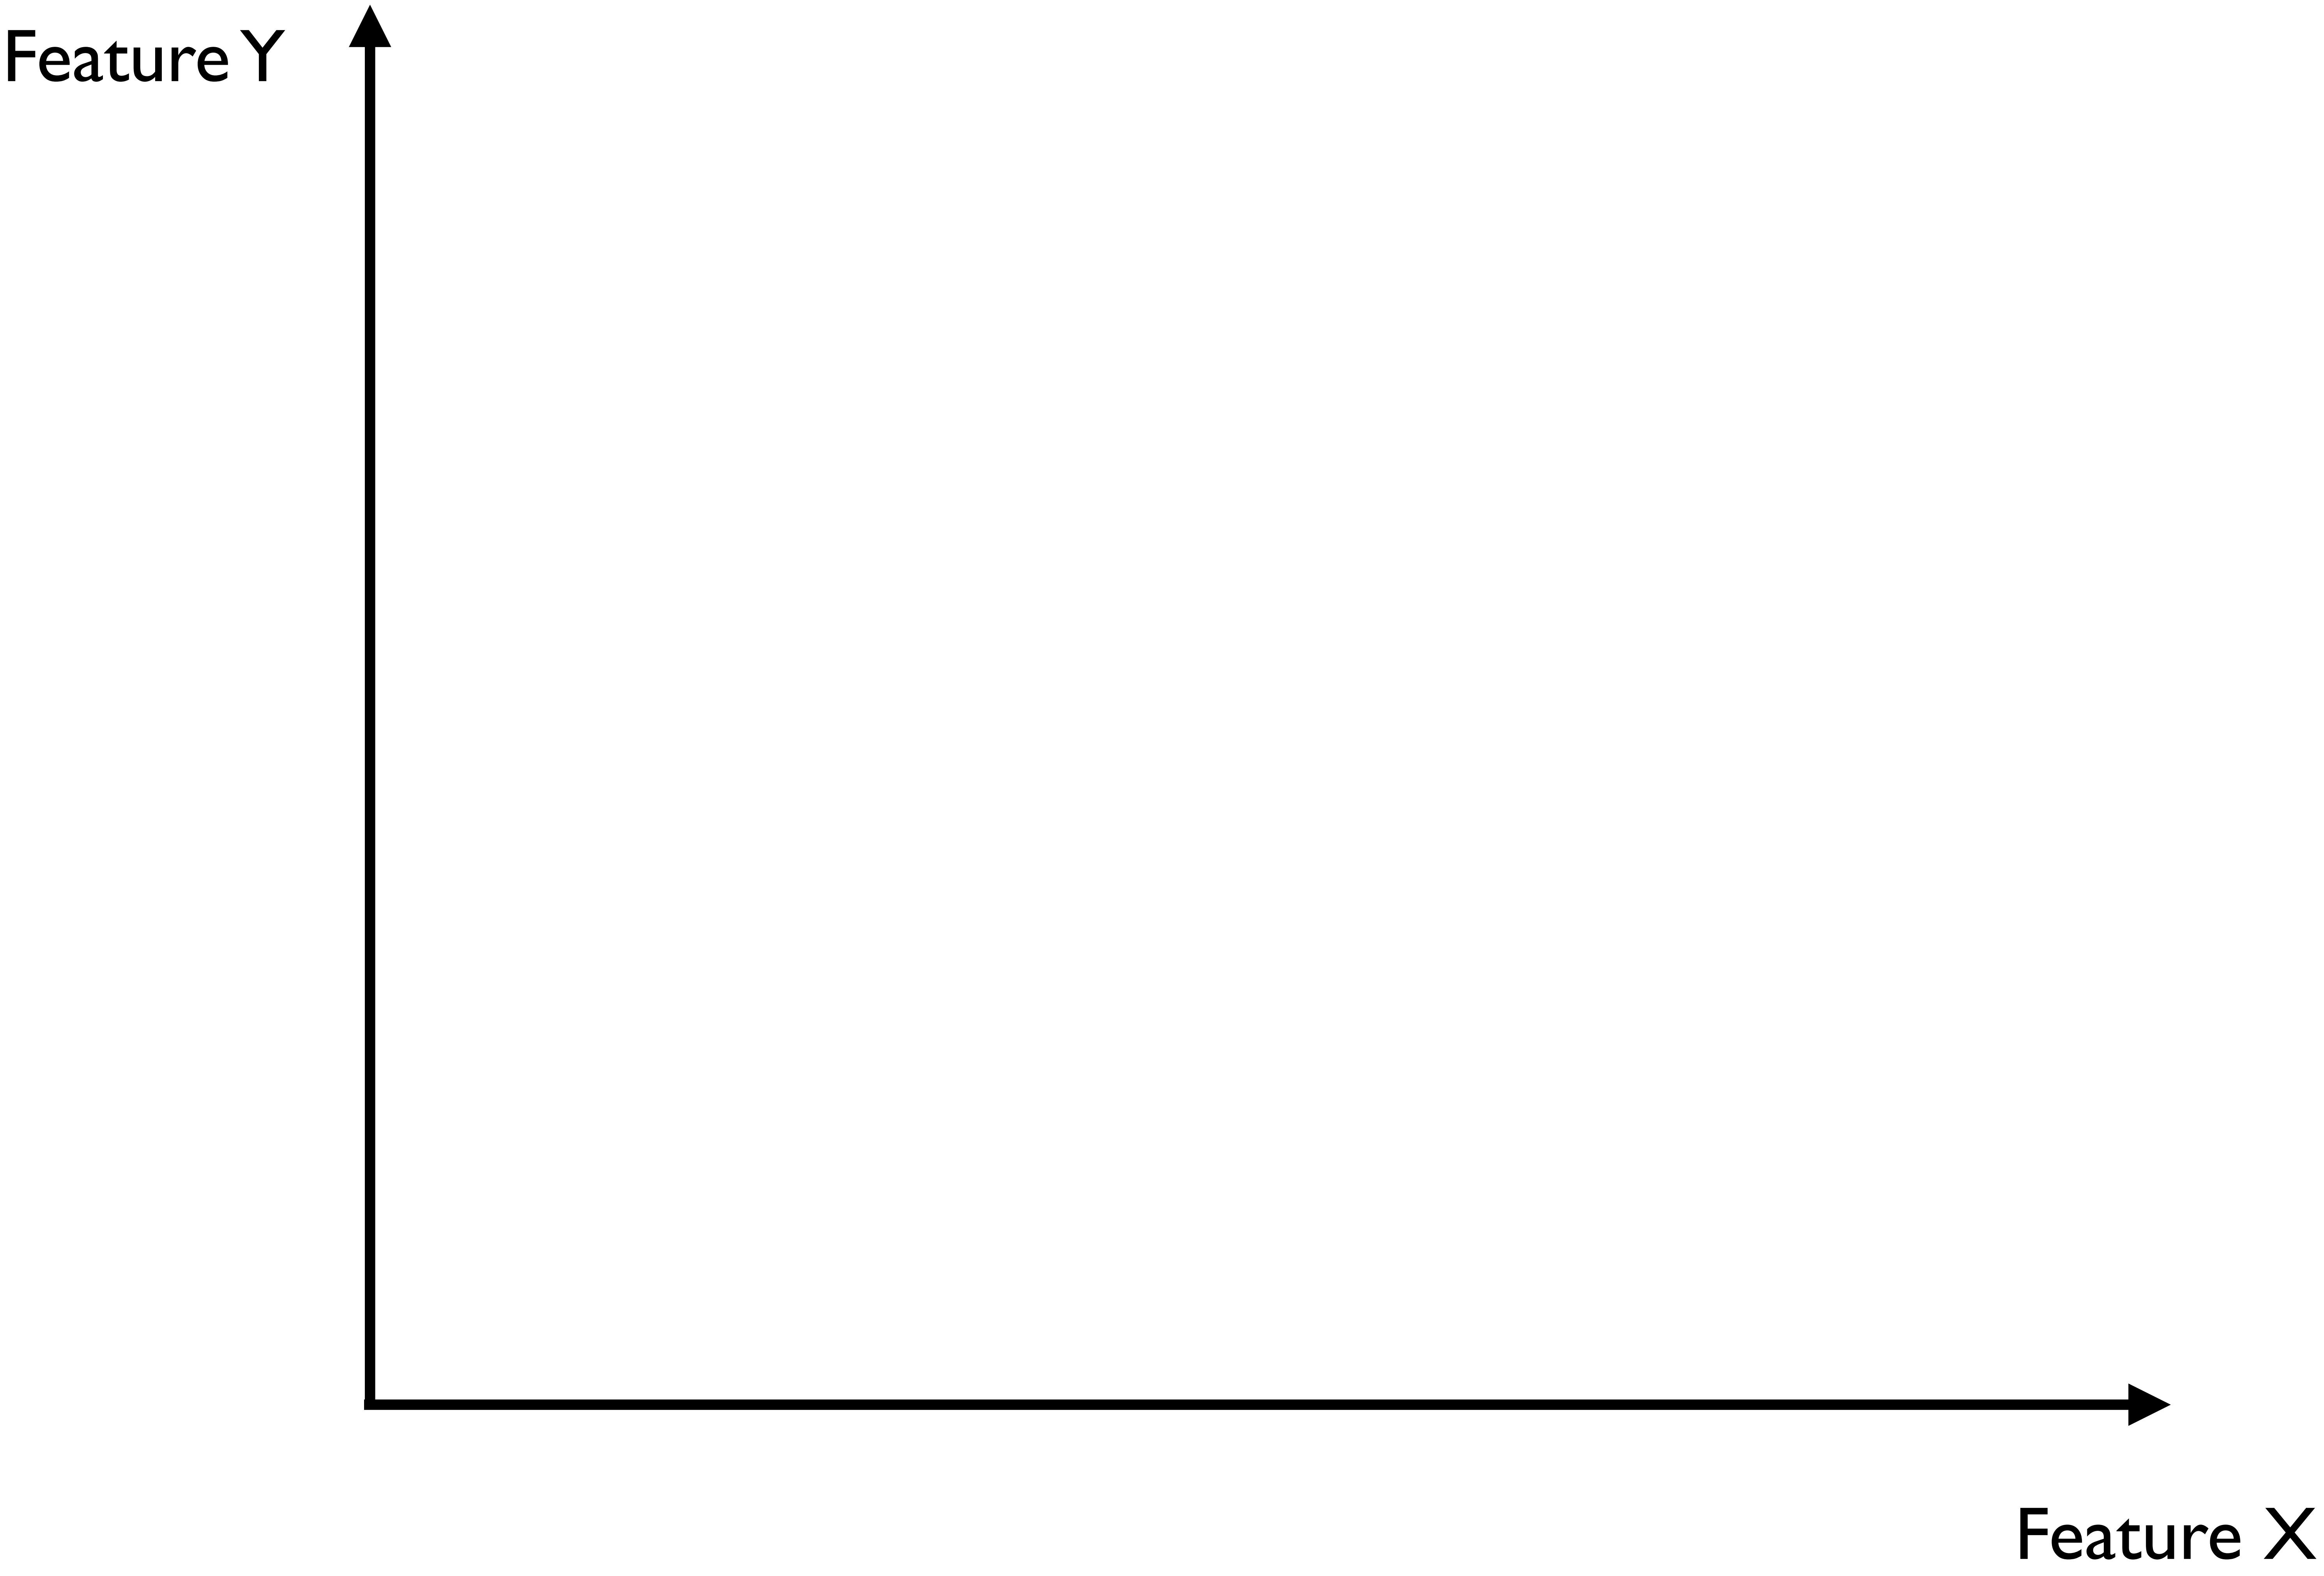

Outlier detection

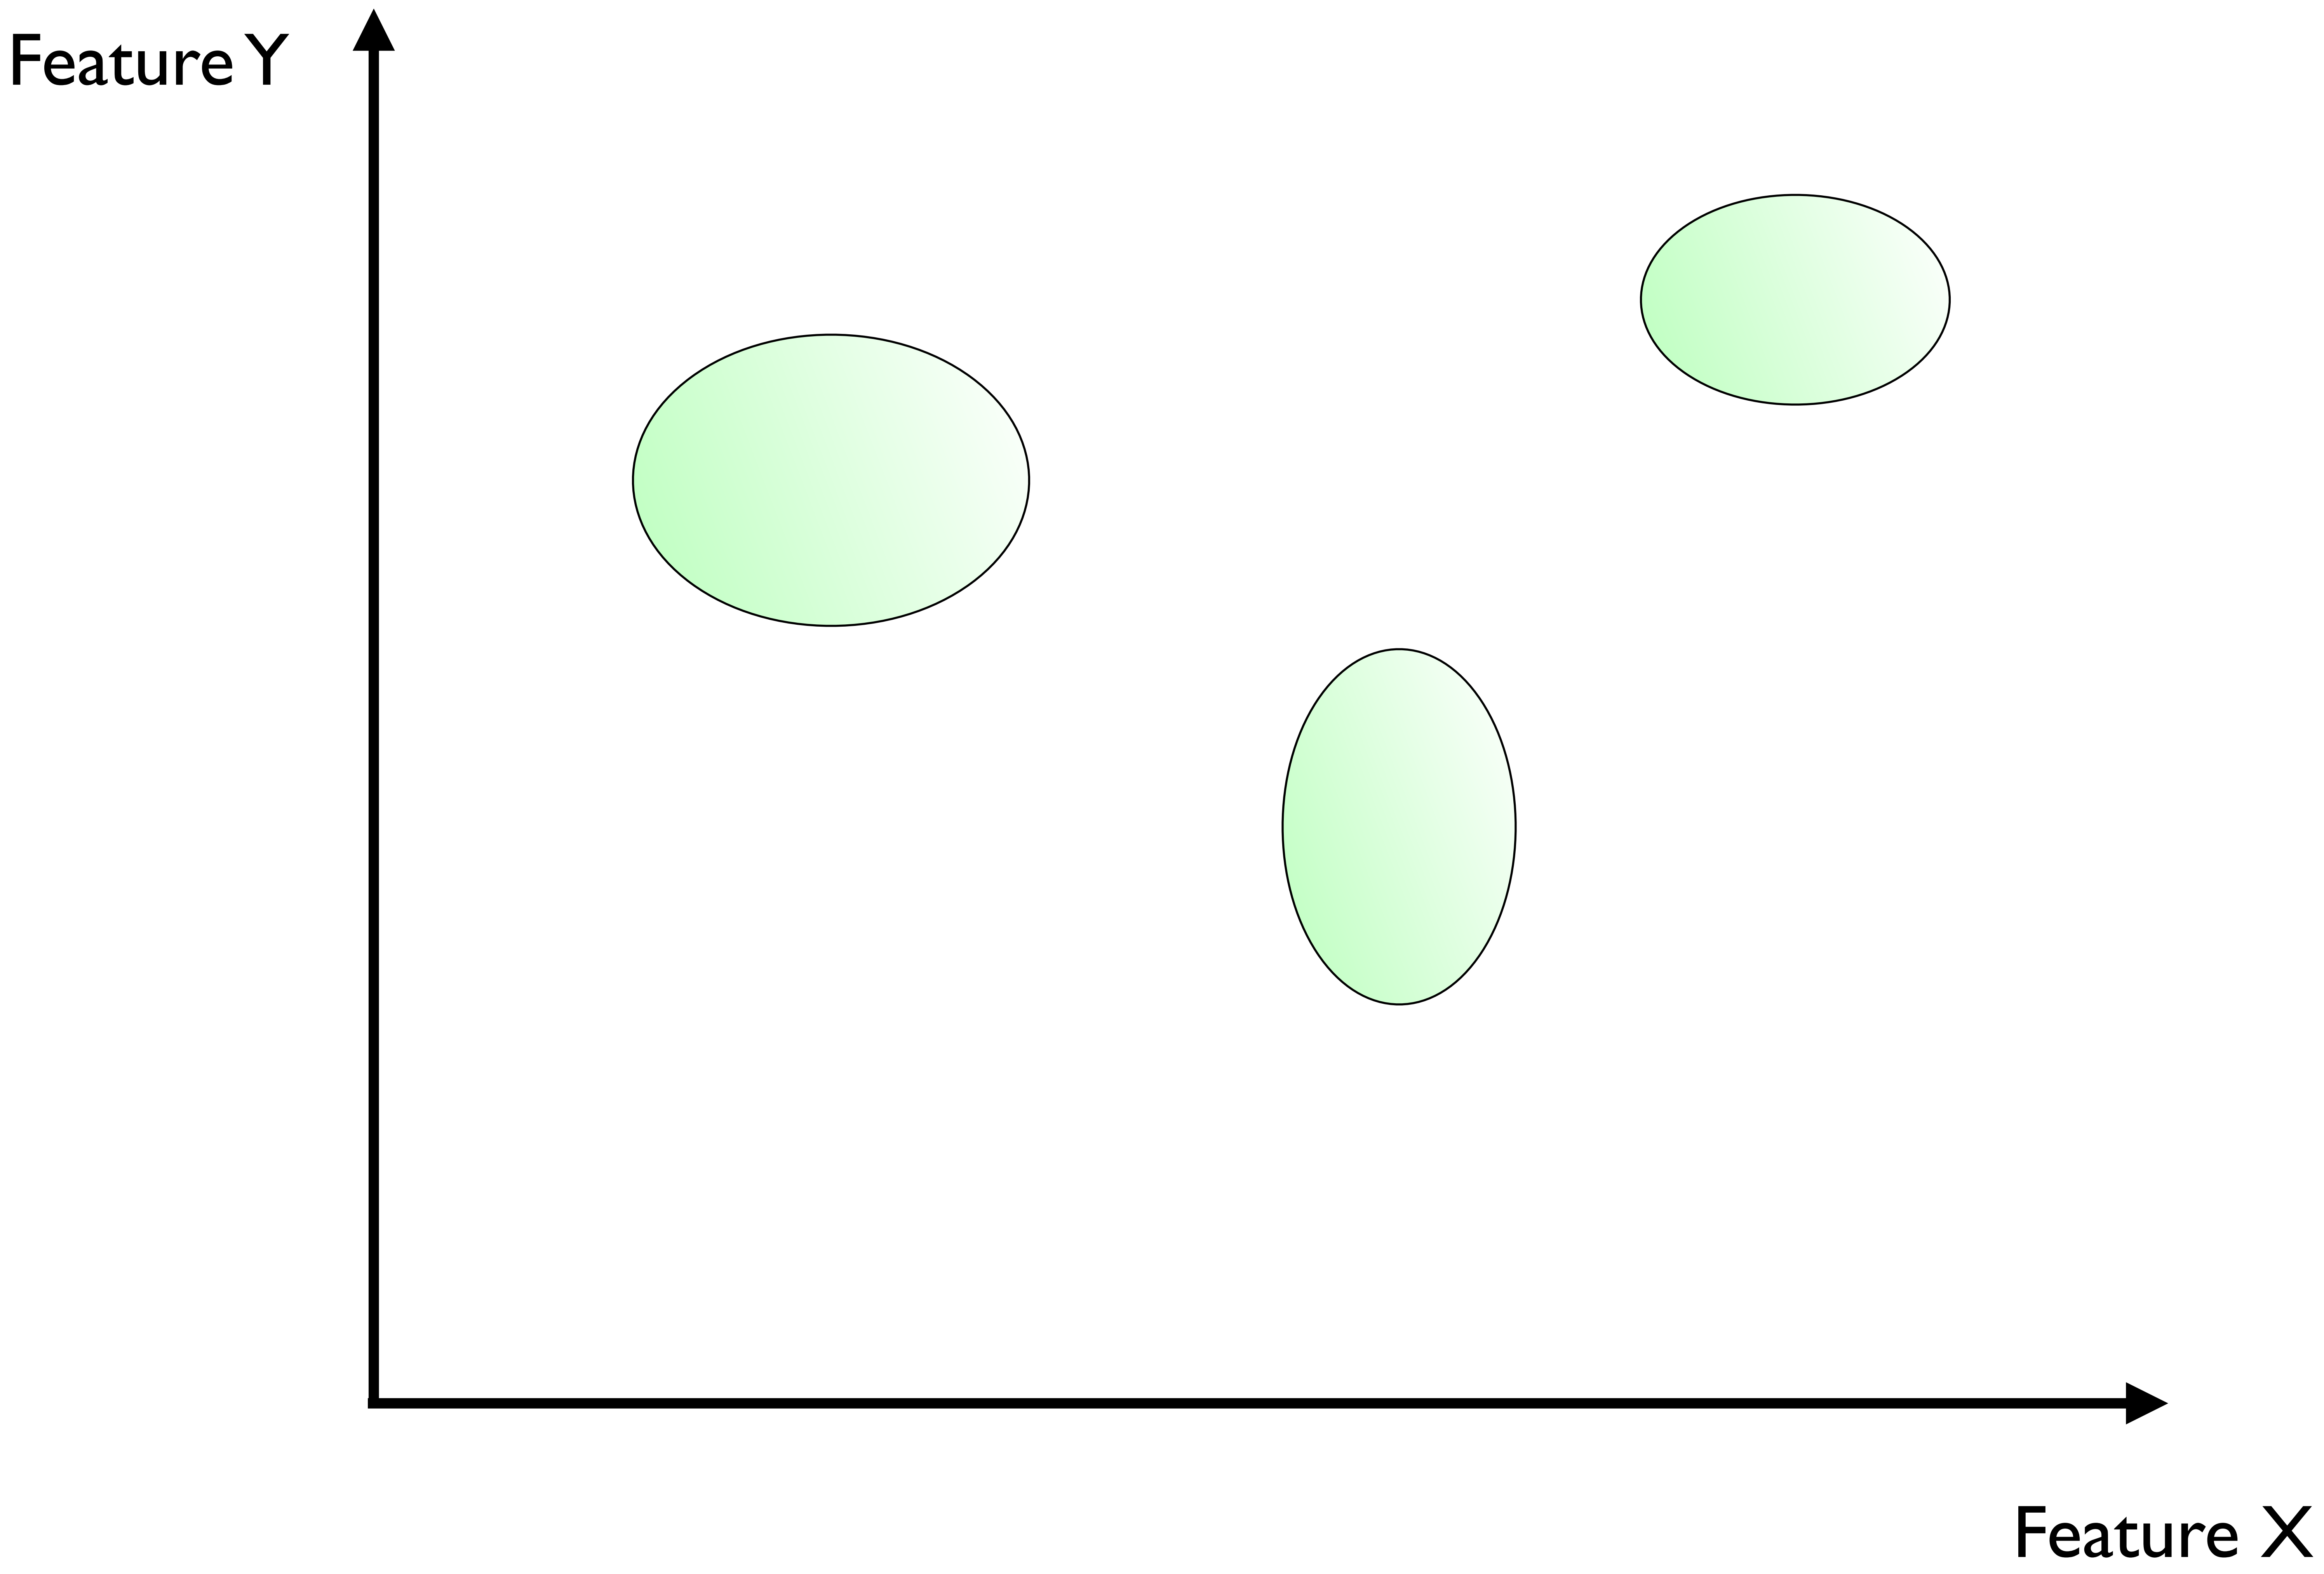

Outlier detection

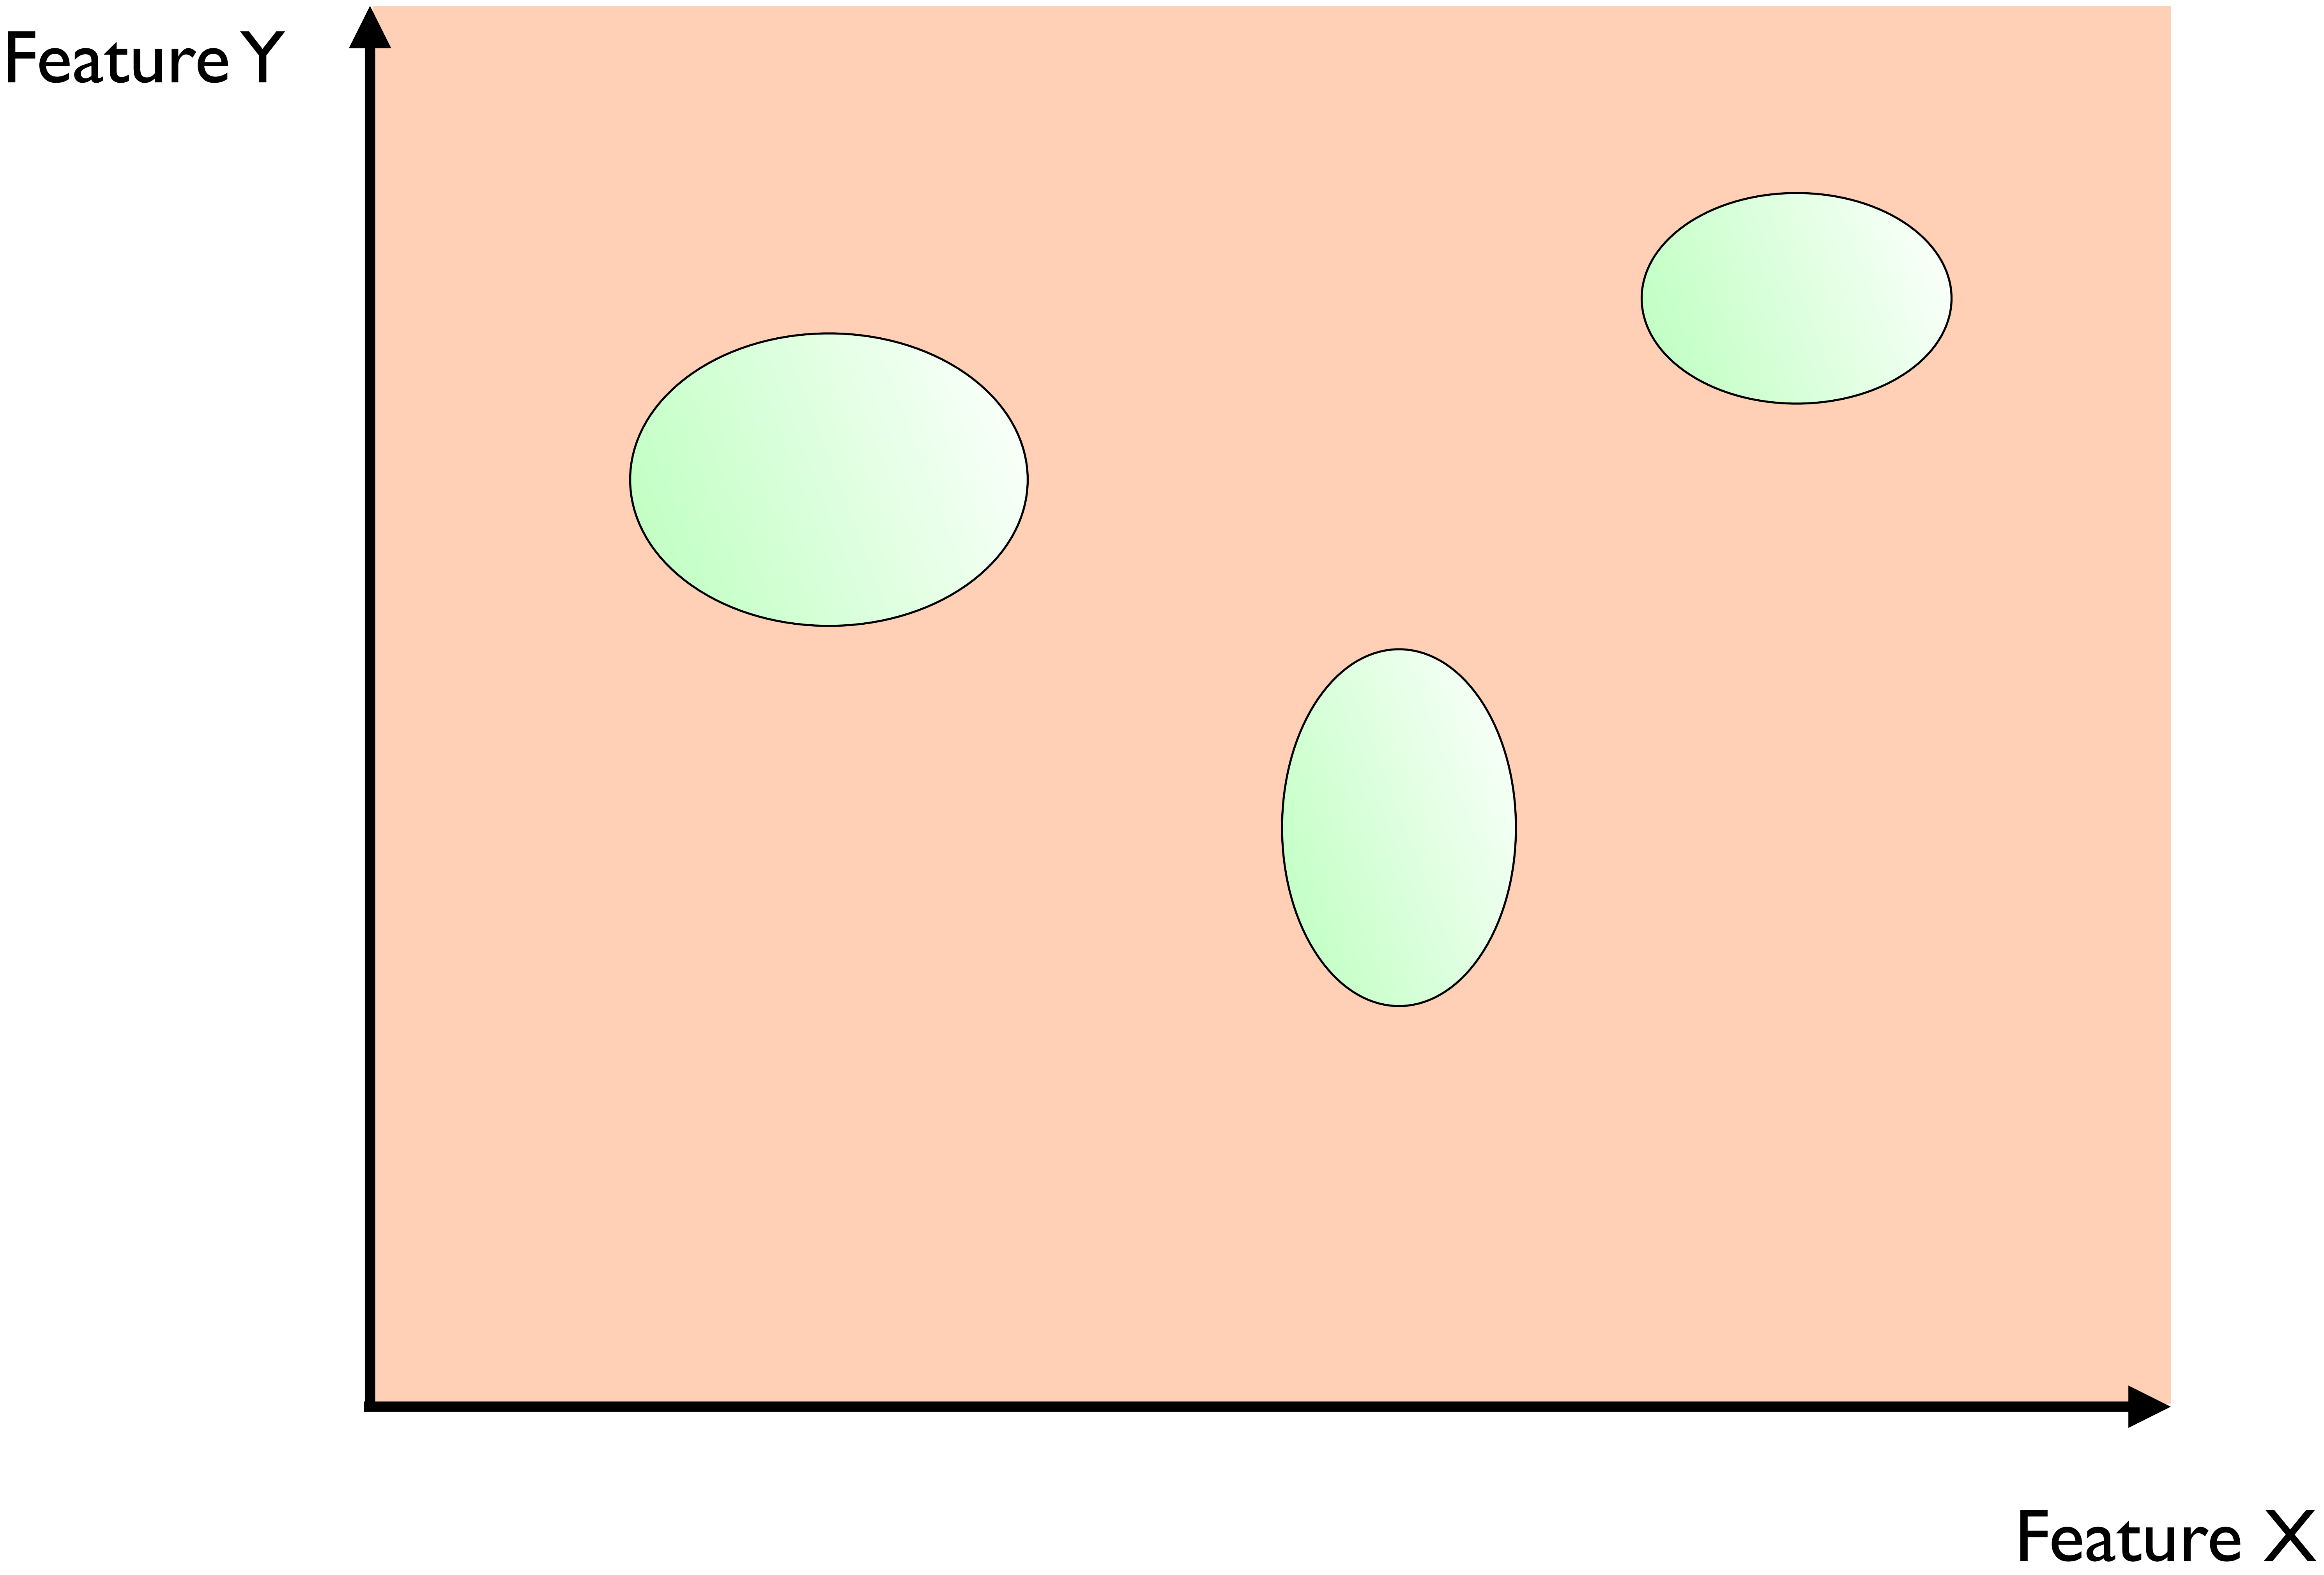

Outlier detection

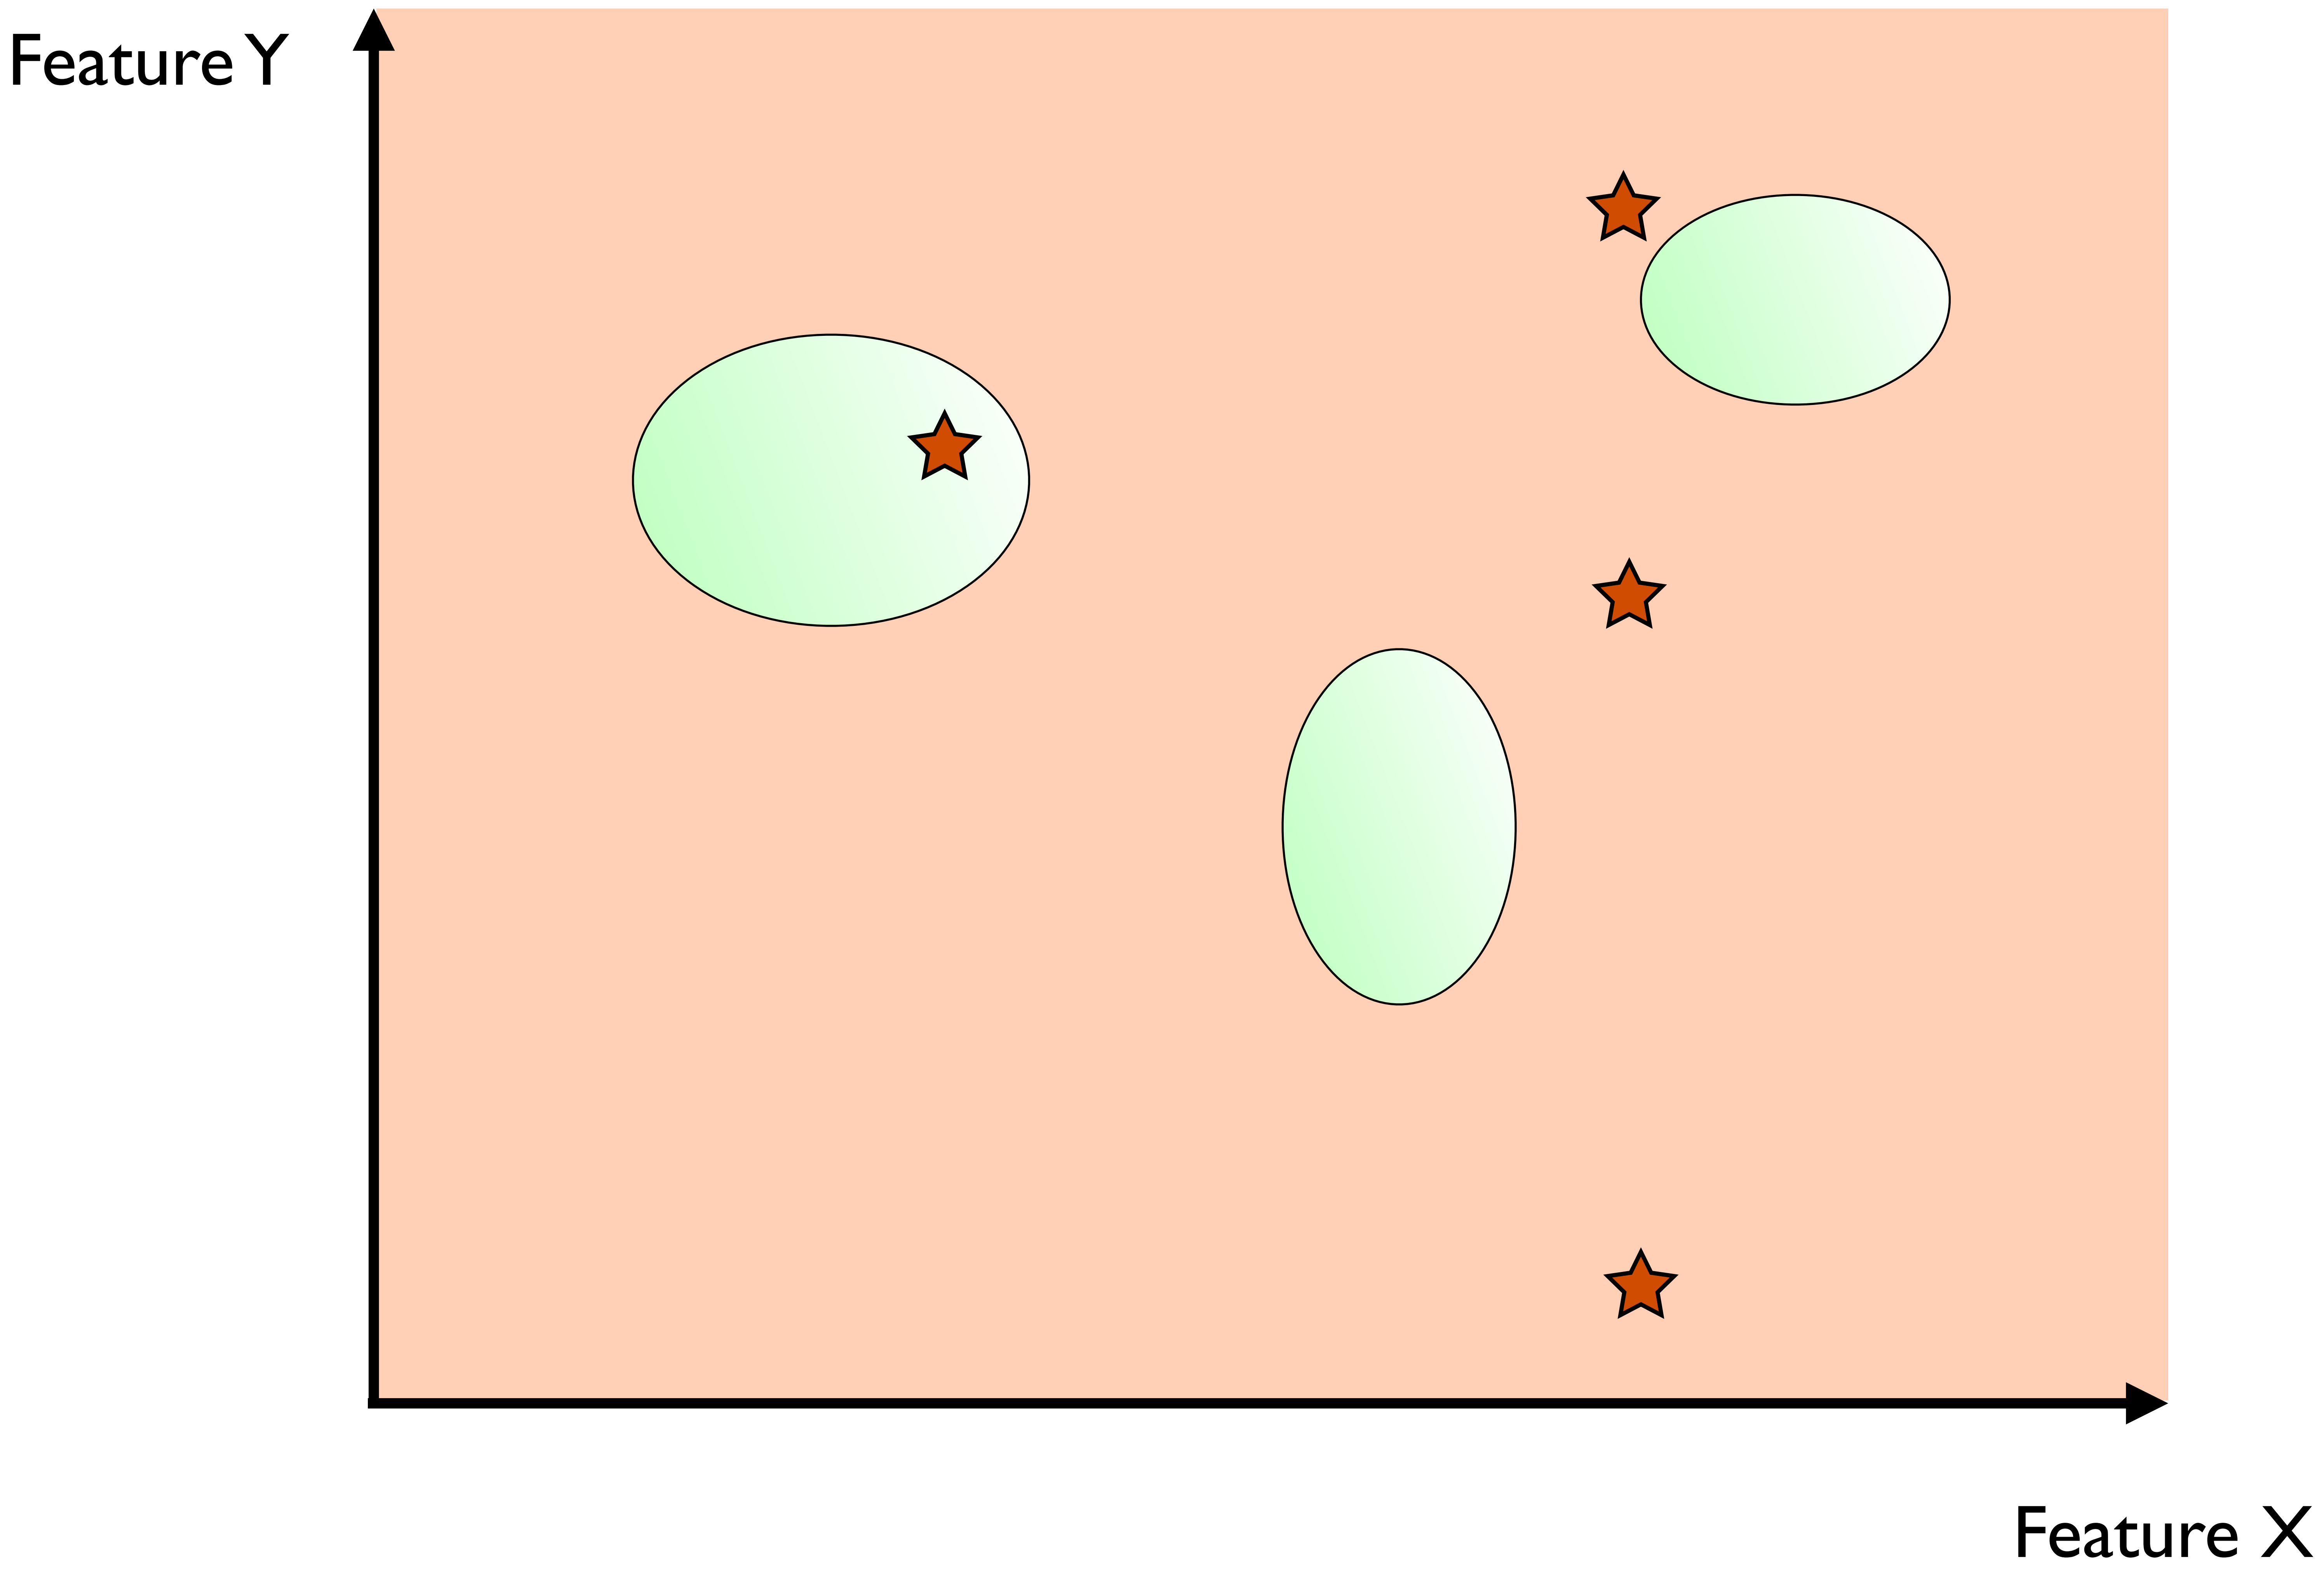

Outlier detection

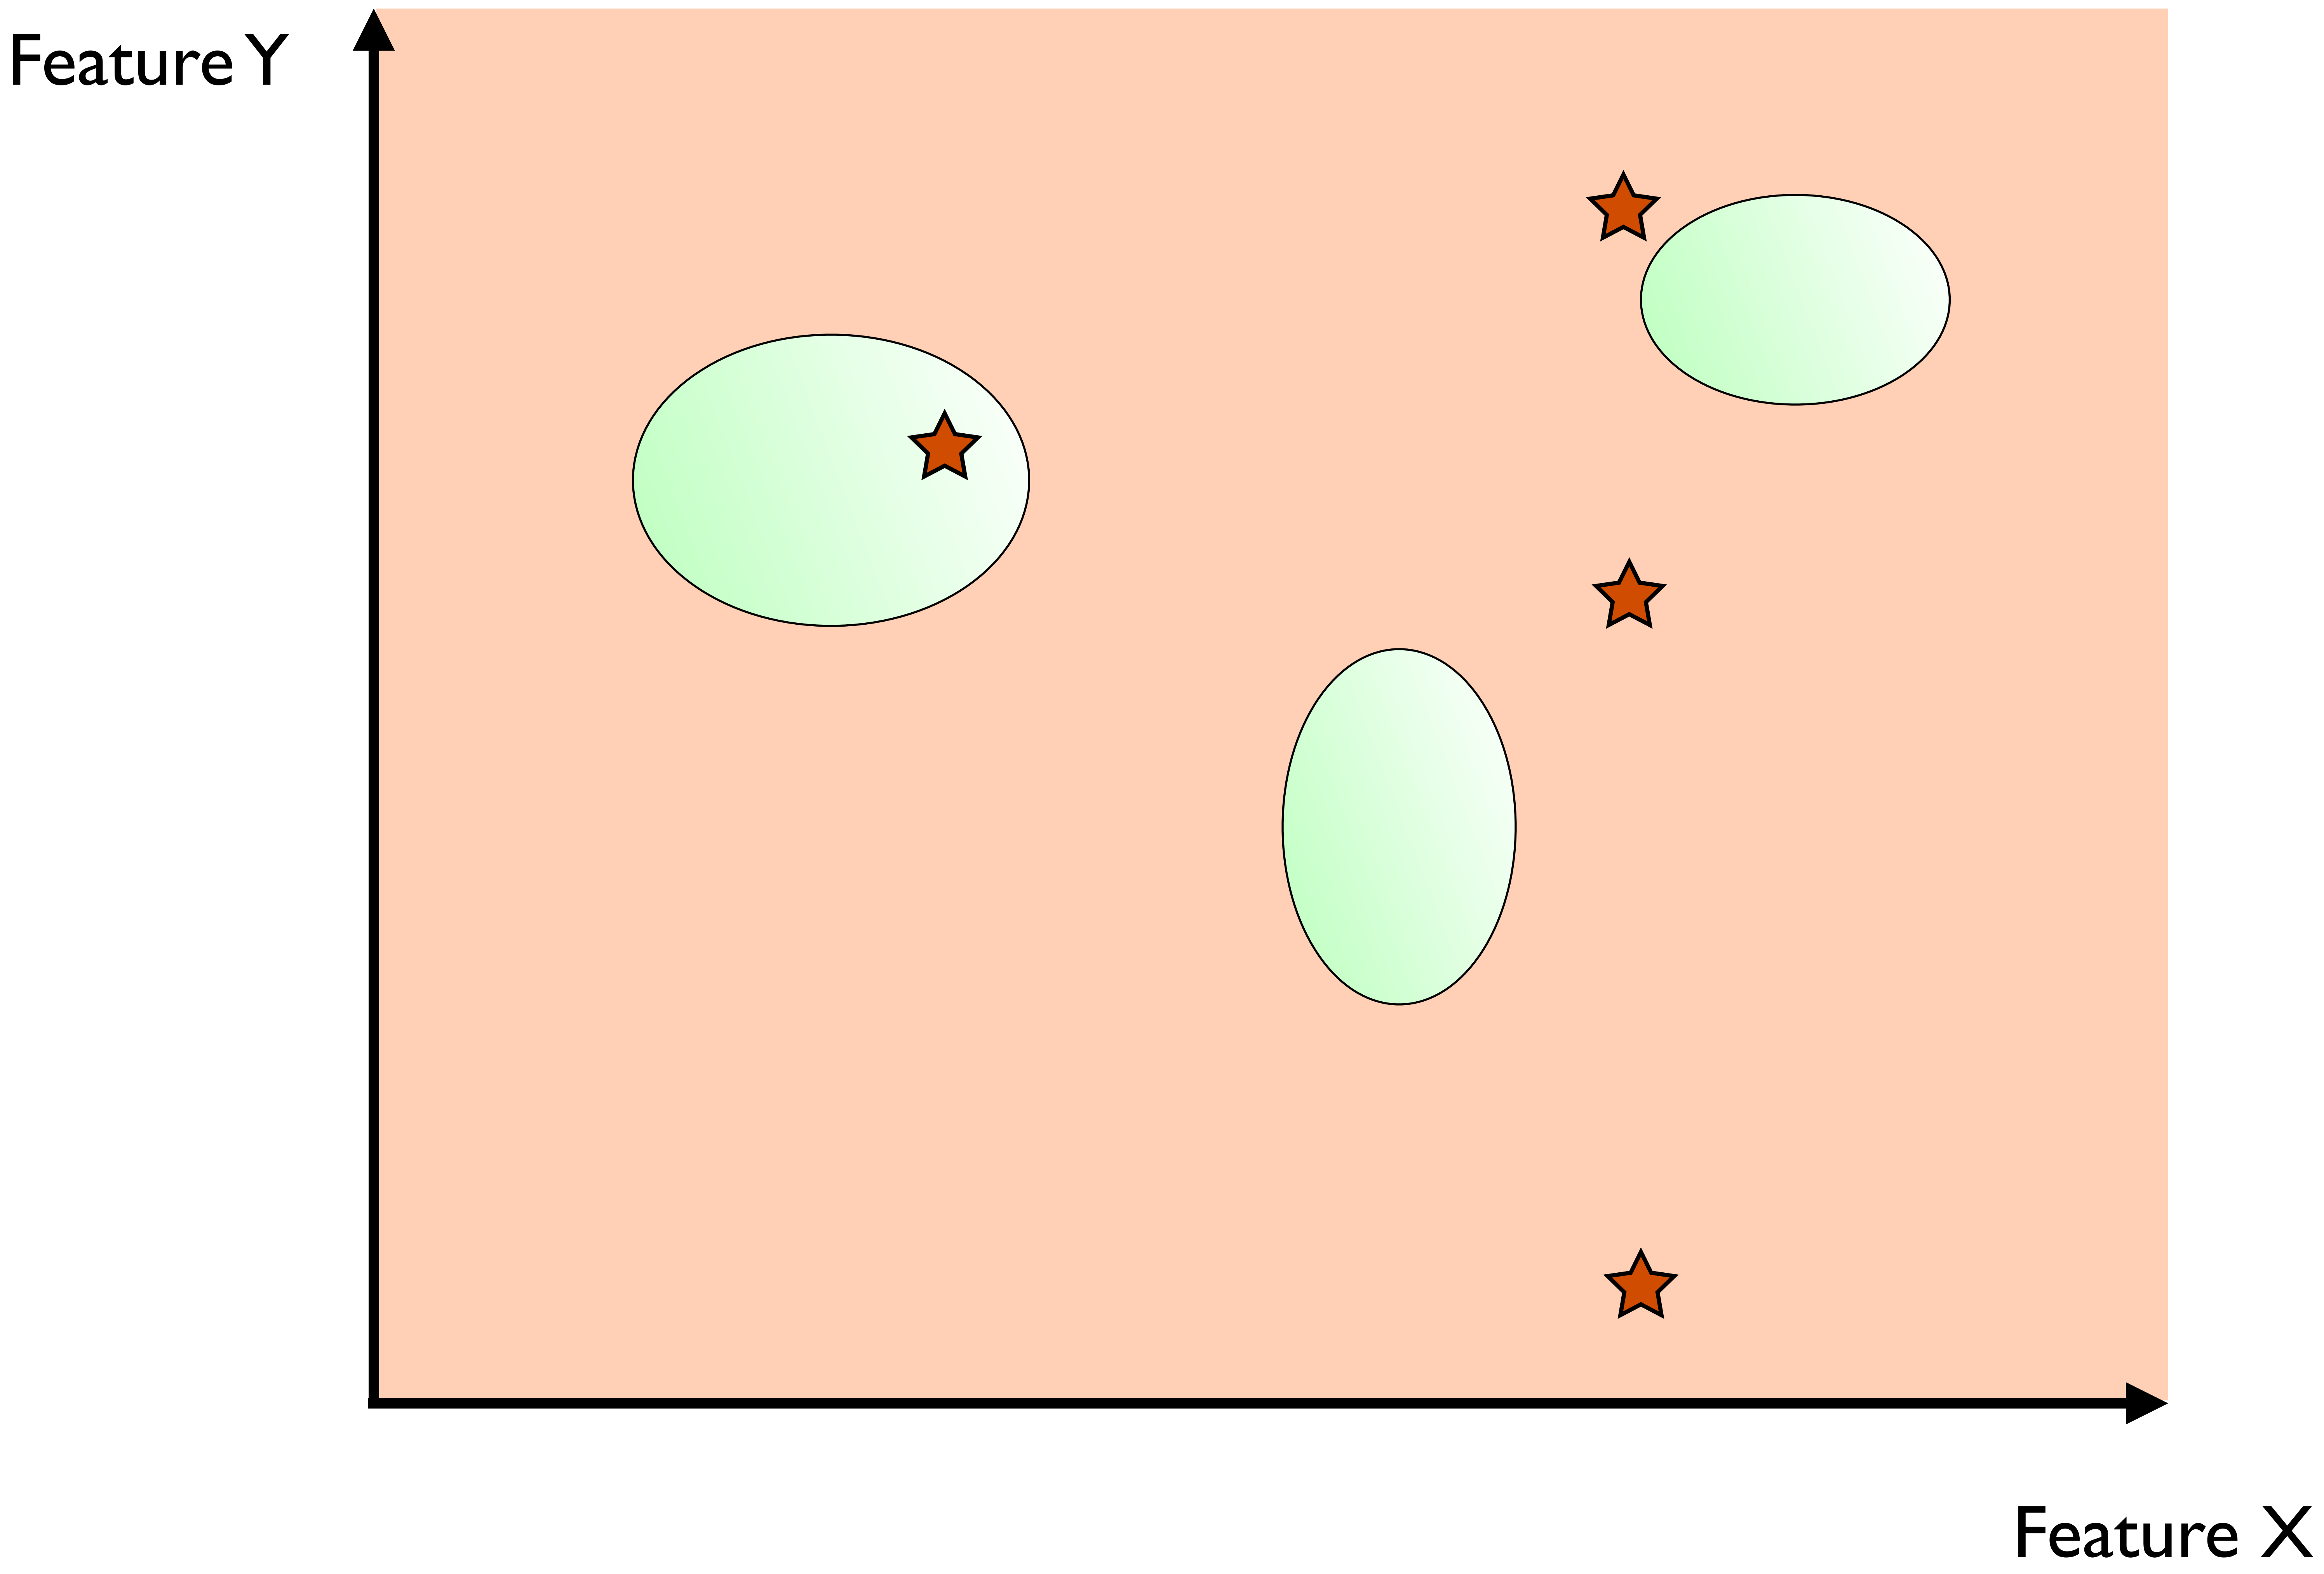

Training with the opposite we're looking for -> No margin for errors

Why is machine learning so ineffective in this domain?

- Machine learning isn’t good at finding outliers
    - In other domains, one looks for activity that’s similar to what’s been trained with

Why is machine learning so ineffective in this domain?

- Machine learning isn’t good at finding outliers
  - In other domains, one looks for activity that’s similar to what’s been trained with
- No stable notion of normality
  - Network environments exhibit enormous variability & noise; “not yet seen” is normal

Why is machine learning so ineffective in this domain?

Machine learning isn’t good at finding outliers

In other domains, one looks for activity that’s similar to what’s been trained with

No stable notion of normality

Network environments exhibit enormous variability & noise; “not yet seen” is normal

Semantic gap

Features do not tie back to operational semantics

Why is machine learning so ineffective in this domain?

Machine learning isn’t good at finding outliers

In other domains, one looks for activity that’s similar to what’s been trained with

No stable notion of normality

Network environments exhibit enormous variability & noise; “not yet seen” is normal

Semantic gap

Features do not tie back to operational semantics

High cost of errors

There are “too few attacks” → base rate fallacy

Increasing precision: narrow classifiers

Increasing precision: narrow classifiers

Activity known to remain quite stable

- Service availability
- SSL certificates
- Executables on a server

Increasing precision: narrow classifiers

Activity known to remain quite stable

- Service availability
- SSL certificates
- Executables on a server

Individual features with characteristic distributions

- URL parameters
- DNS lookups
- Communication timing (e.g., interactive logins)

Increasing precision: narrow classifiers

- Activity known to remain quite stable
  - Service availability
  - SSL certificates
  - Executables on a server
- Individual features with characteristic distributions
  - URL parameters
  - DNS lookups
  - Communication timing (e.g., interactive logins)
- Variations of known attacks
  - Pre-canned attack tools
  - Phishing emails

So, why is detecting novel attacks so difficult?

We’re limited to finding what we can describe, one way or the other.

So, why is detecting novel attacks so difficult?

We're limited to finding what we can describe, one way or the other.

Misuse Detection	By definition: We need a library of attacks
Anomaly Detection	Need to target something we understand

So, why is detecting novel attacks so difficult?

We're limited to finding what we can describe, one way or the other.

Misuse Detection	By definition: We need a library of attacks
Anomaly Detection	Need to target something we understand

Corollary: The more sophisticated the attacker, the less likely we'll be detecting what they are doing.

From Intrusion Detection to Threat Hunting

When are attackers found

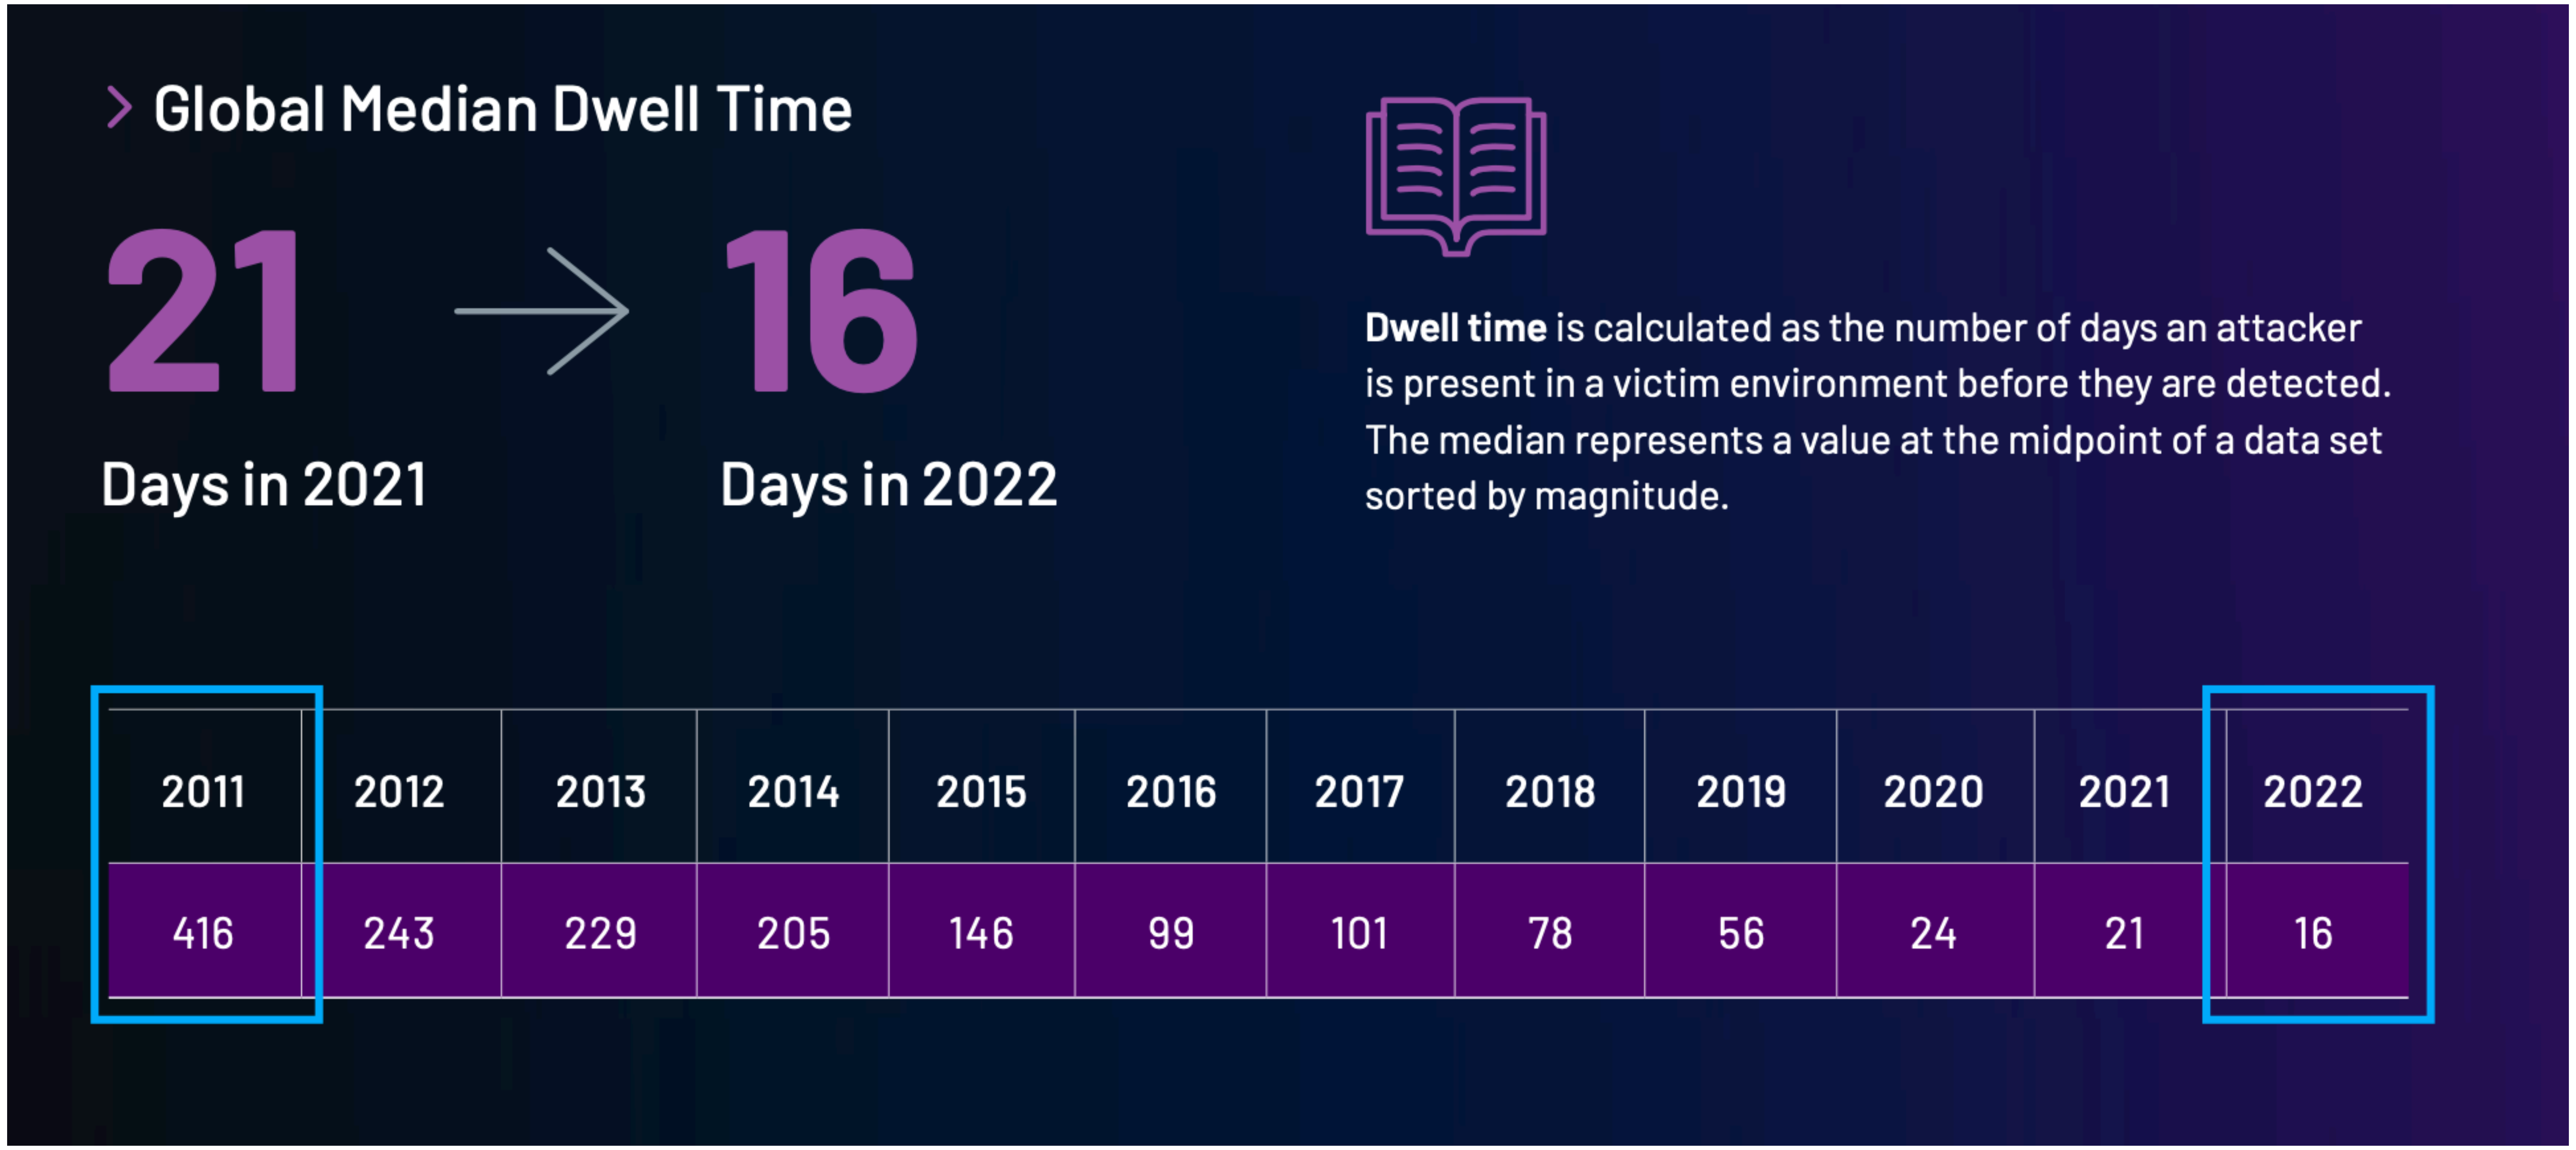

Global Median Dwell Time

21 → **16**

Days in 2021

Days in 2022

Dwell time is calculated as the number of days an attacker is present in a victim environment before they are detected. The median represents a value at the midpoint of a data set sorted by magnitude.

2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
416	243	229	205	146	99	101	78	56	24	21	16

Source: Mandiant M-Trends® 2023

When are attackers found

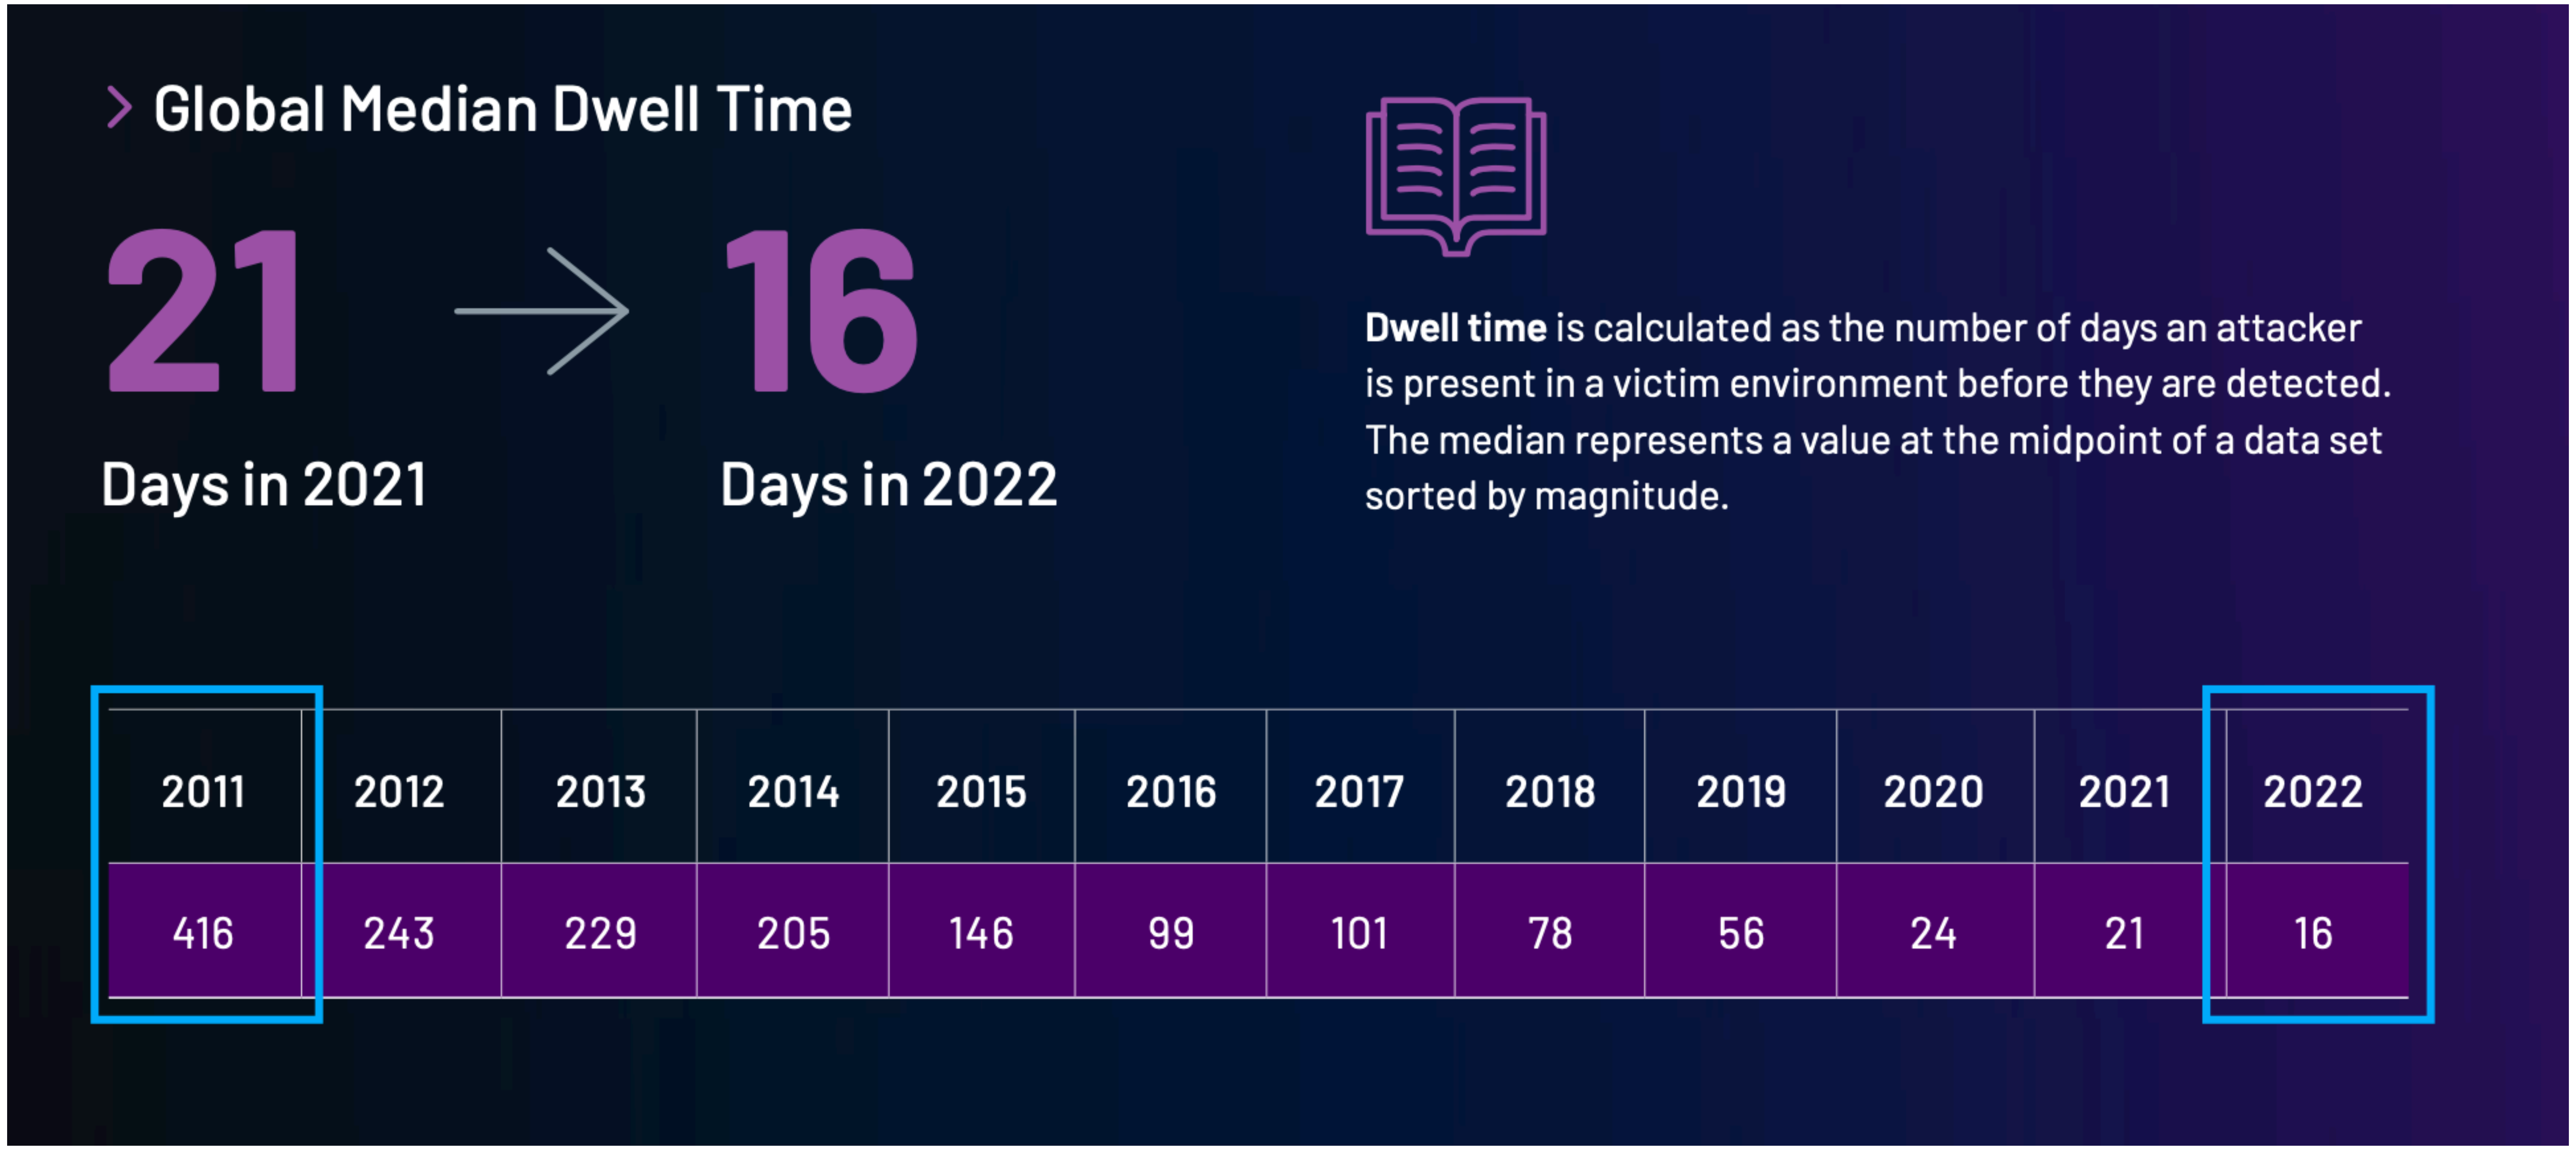

Global Median Dwell Time

21 → **16**

Days in 2021 Days in 2022

Dwell time is calculated as the number of days an attacker is present in a victim environment before they are detected. The median represents a value at the midpoint of a data set sorted by magnitude.

2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
416	243	229	205	146	99	101	78	56	24	21	16

Source: Mandiant M-Trends® 2023

As a defender you might just as well assume somebody is in your network already.

The rise of “threat hunting”

The rise of “threat hunting”

> “Threat hunting is the practice of proactively searching for cyber threats that are lurking undetected in a network. Cyber threat hunting digs deep to find malicious actors in your environment that have slipped past your initial […] security defenses.”

Crowdstrike (2023)

The rise of “threat hunting”

> “Threat hunting is the practice of proactively searching for cyber threats that are lurking undetected in a network. Cyber threat hunting digs deep to find malicious actors in your environment that have slipped past your initial […] security defenses.”

Crowdstrike (2023)

> “Threat hunters are incident responders and forensic investigators actively looking for new threats before traditional intrusion detection methods can find them.”

Bob Lee, SANS (2016)

The rise of “threat hunting”

“Threat hunting is the practice of proactively searching for cyber threats that are lurking undetected in a network. Cyber threat hunting digs deep to find malicious actors in your environment that have slipped past your initial […] security defenses.”

Crowdstrike (2023)

“Threat hunters are incident responders and forensic investigators actively looking for new threats before traditional intrusion detection methods can find them.”

Bob Lee, SANS (2016)

“Defenders must actively hunt intruders in their enterprise. […] Rather than hoping defenses will repel invaders, or that breaches will be caught by passive alerting mechanisms, […] defeating intruders requires actively detecting and responding to them.

Richard Bejtlich,”Become a Hunter”, Information Security Magazine (2011)

The rise of “threat hunting”

“Threat hunting is the practice of proactively searching for cyber threats that are lurking undetected in a network. Cyber threat hunting digs deep to find malicious actors in your environment that have slipped past your initial […] security defenses.”

Crowdstrike (2023)

“Threat hunters are incident responders and forensic investigators actively looking for new threats before traditional intrusion detection methods can find them.”

Bob Lee, SANS (2016)

“Defenders must actively hunt intruders in their enterprise. […] Rather than hoping defenses will repel invaders, or that breaches will be caught by passive alerting mechanisms, […] defeating intruders requires actively detecting and responding to them.

Richard Bejtlich,”Become a Hunter”, Information Security Magazine (2011)

Why wasn’t this detected earlier?

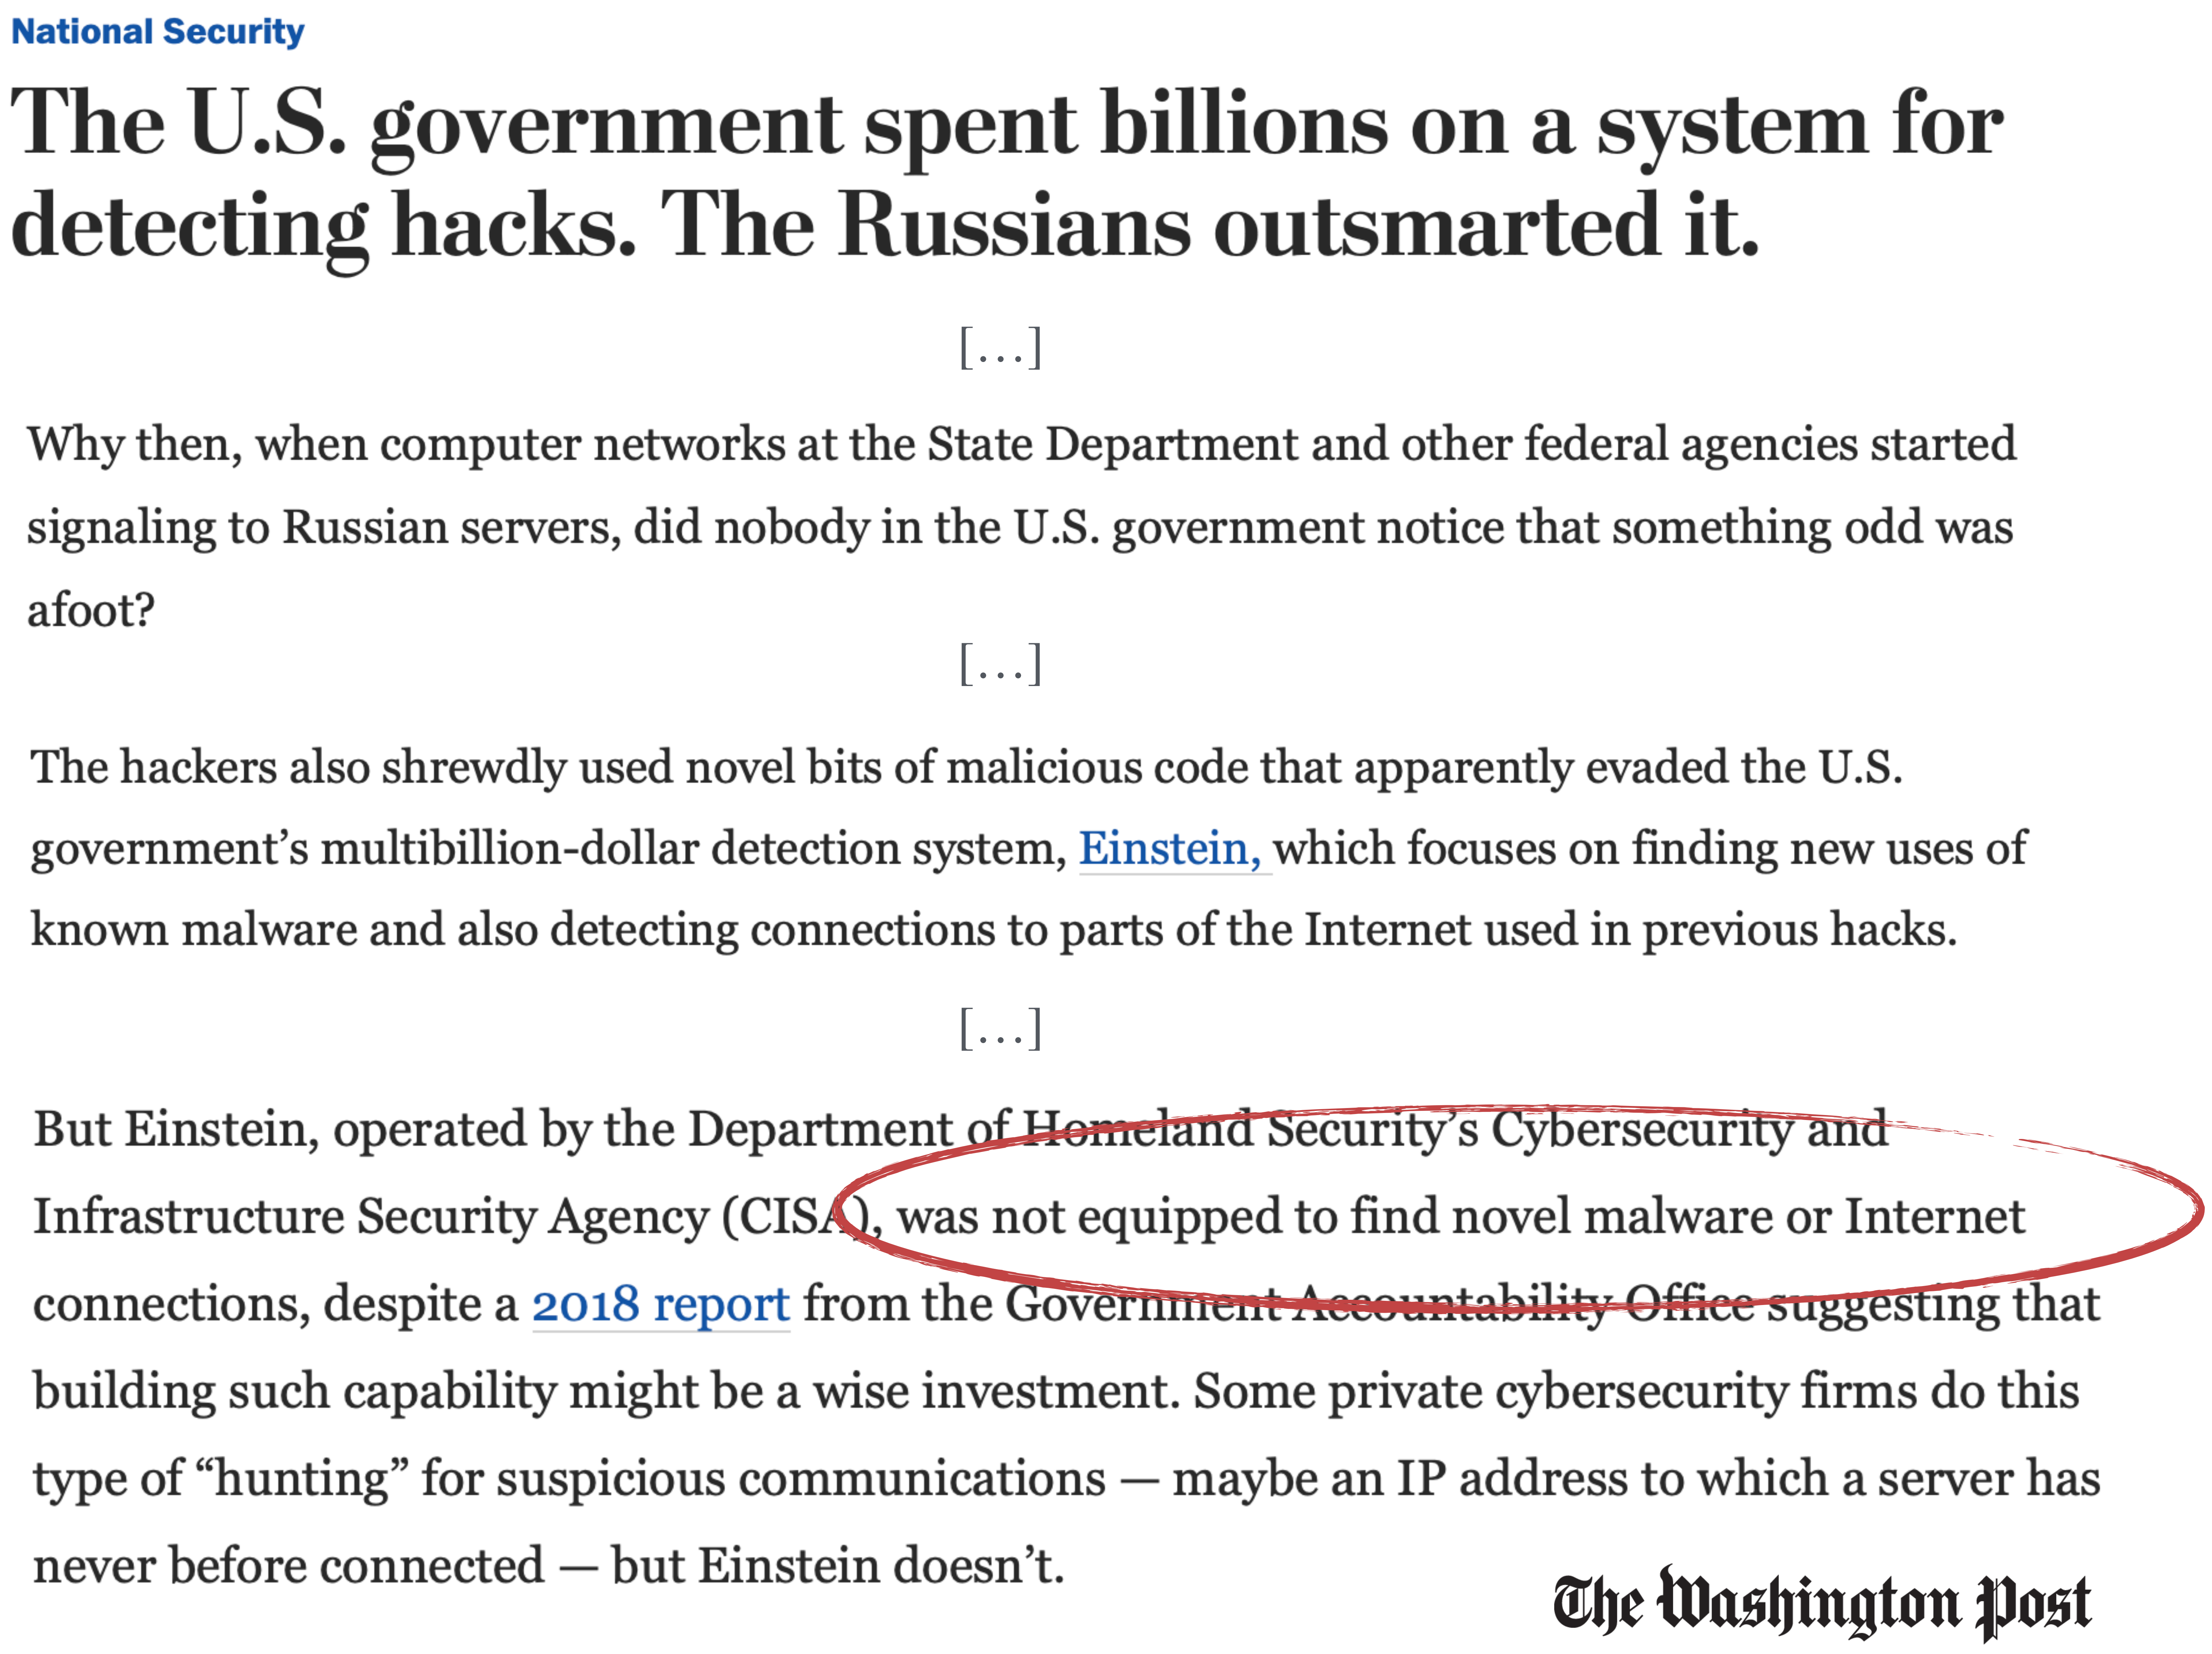

National Security

The U.S. government spent billions on a system for detecting hacks. The Russians outsmarted it.

[…]

Why then, when computer networks at the State Department and other federal agencies started signaling to Russian servers, did nobody in the U.S. government notice that something odd was afoot?

[…]

The hackers also shrewdly used novel bits of malicious code that apparently evaded the U.S. government’s multibillion-dollar detection system, Einstein, which focuses on finding new uses of known malware and also detecting connections to parts of the Internet used in previous hacks.

[…]

But Einstein, operated by the Department of Homeland Security’s Cybersecurity and Infrastructure Security Agency (CISA), was not equipped to find novel malware or Internet connections, despite a 2018 report from the Government Accountability Office suggesting that building such capability might be a wise investment. Some private cybersecurity firms do this type of “hunting” for suspicious communications — maybe an IP address to which a server has never before connected — but Einstein doesn’t.

The Washington Post

Why wasn’t this detected earlier?

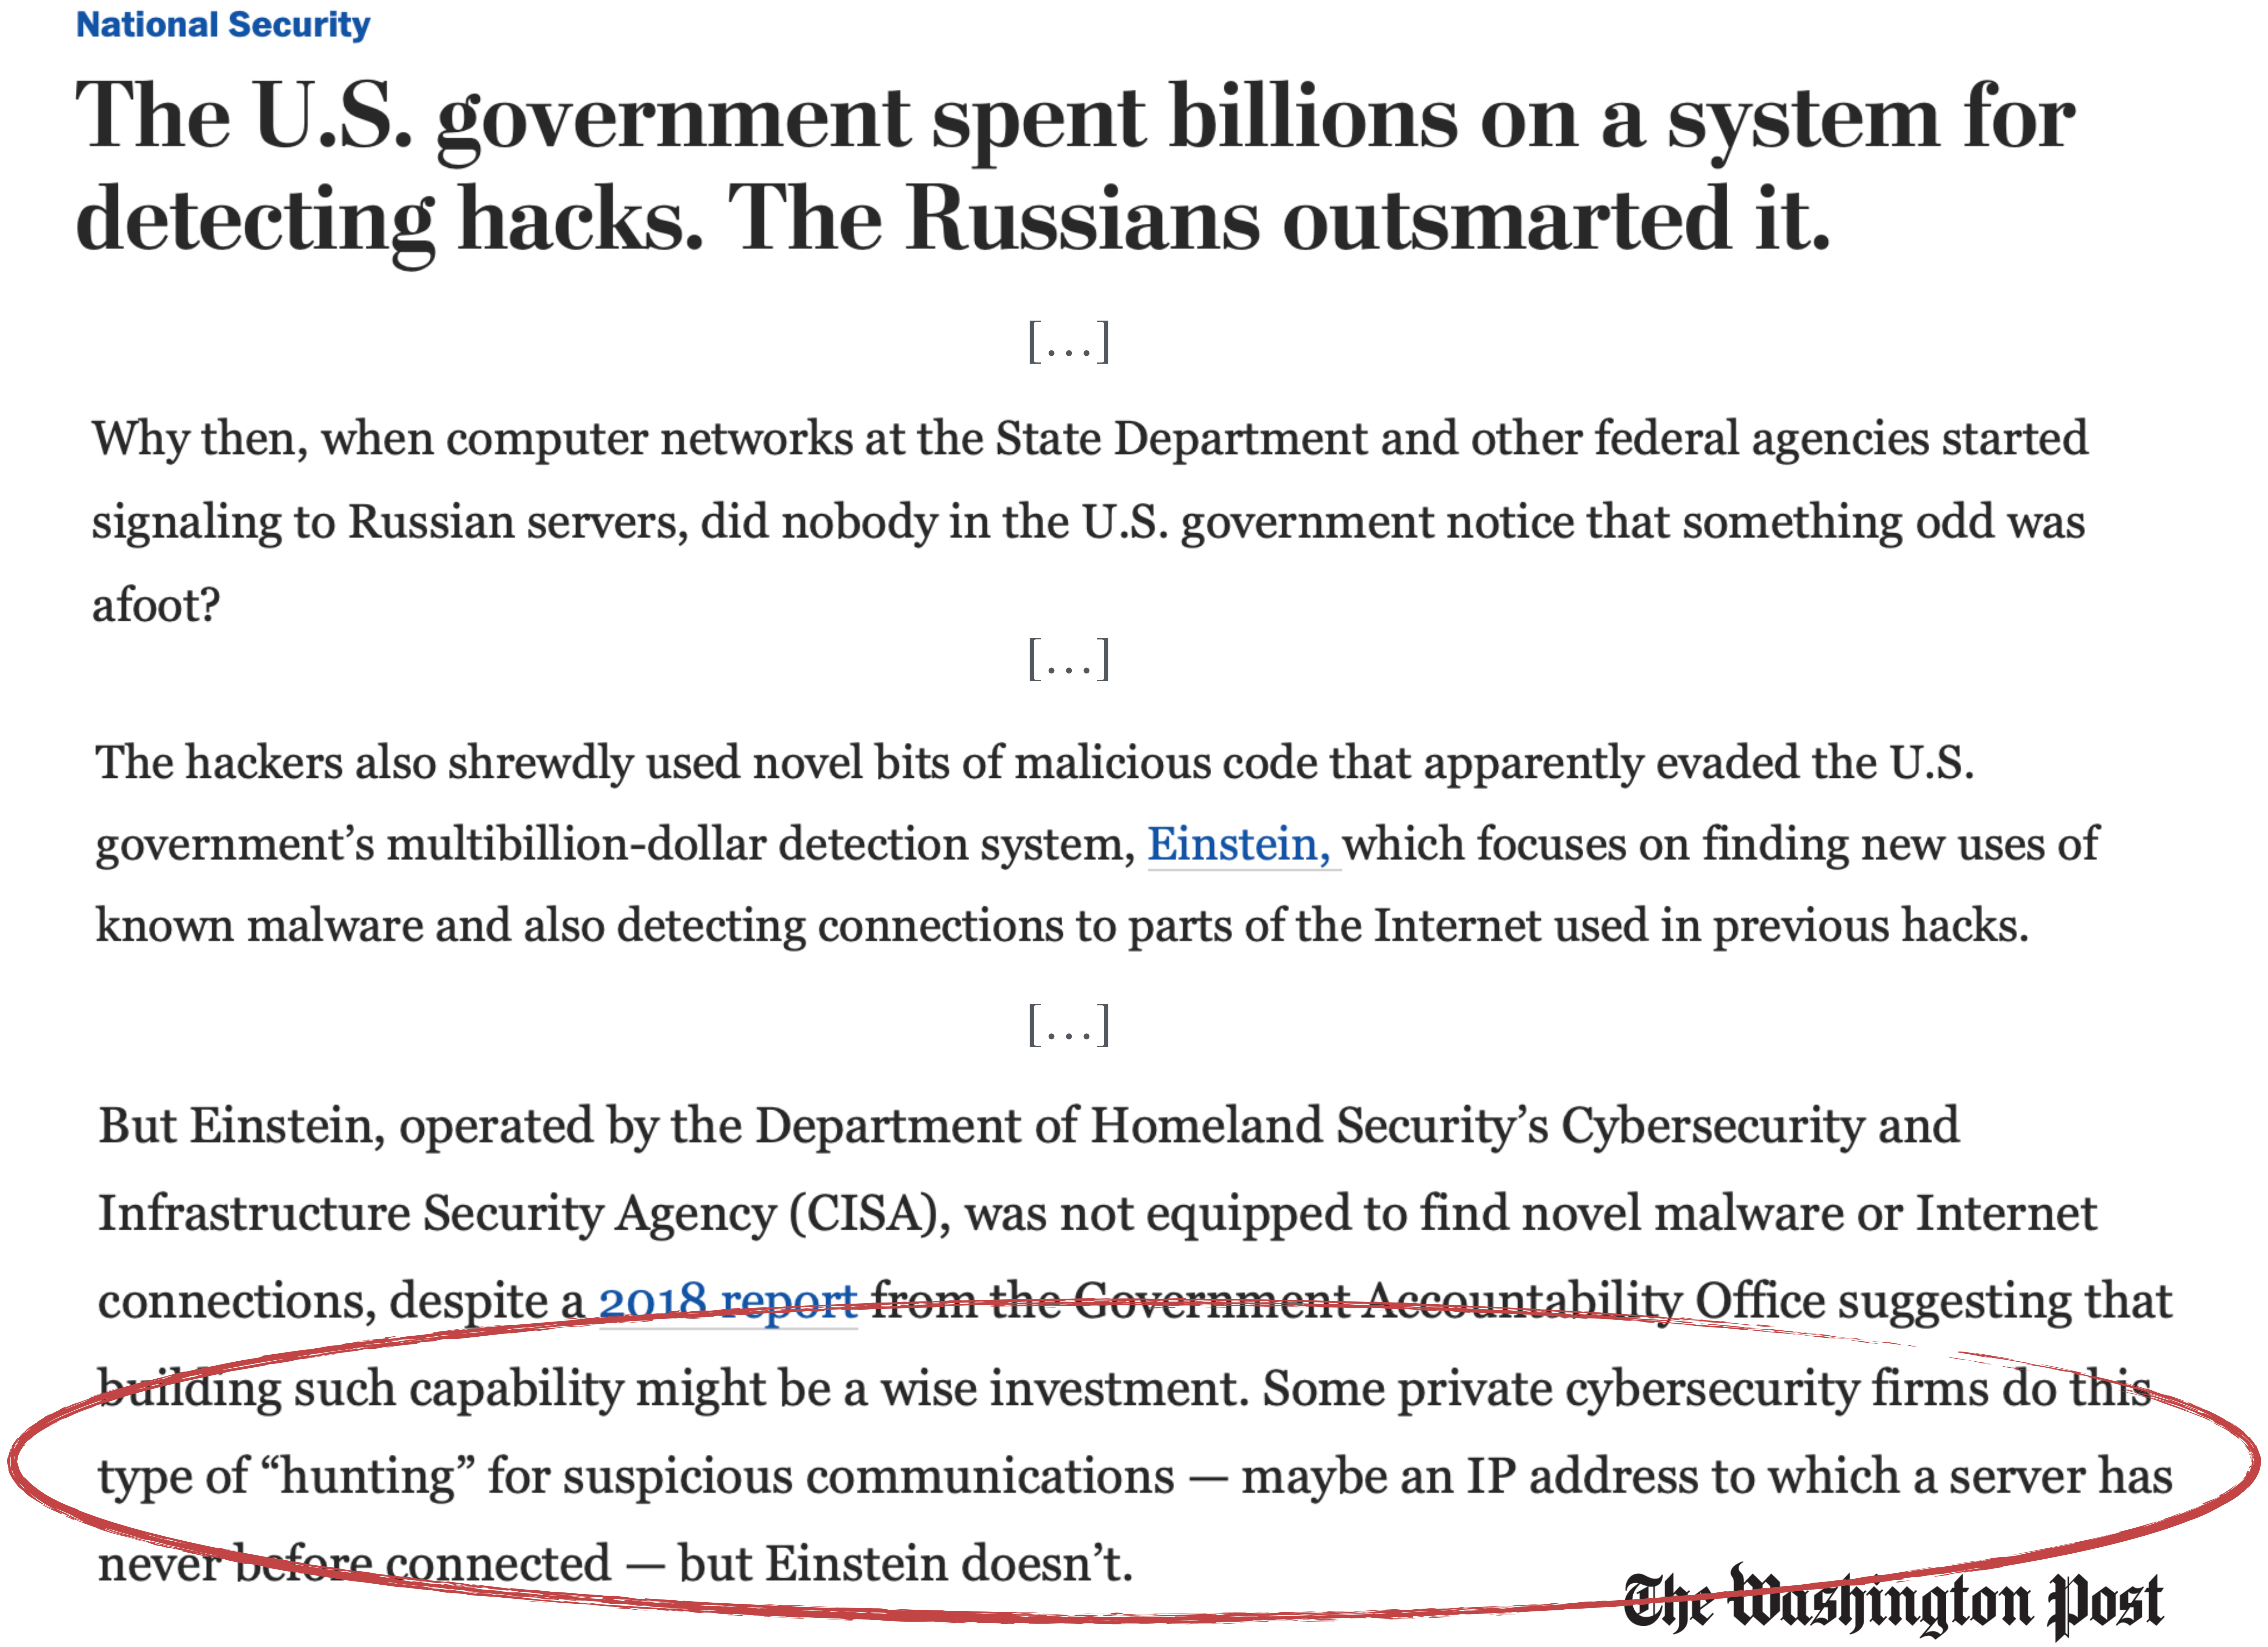

National Security

The U.S. government spent billions on a system for detecting hacks. The Russians outsmarted it.

[…]

Why then, when computer networks at the State Department and other federal agencies started signaling to Russian servers, did nobody in the U.S. government notice that something odd was afoot?

[…]

The hackers also shrewdly used novel bits of malicious code that apparently evaded the U.S. government’s multibillion-dollar detection system, Einstein, which focuses on finding new uses of known malware and also detecting connections to parts of the Internet used in previous hacks.

[…]

But Einstein, operated by the Department of Homeland Security’s Cybersecurity and Infrastructure Security Agency (CISA), was not equipped to find novel malware or Internet connections, despite a 2018 report from the Government Accountability Office suggesting that building such capability might be a wise investment. Some private cybersecurity firms do this type of “hunting” for suspicious communications — maybe an IP address to which a server has never before connected — but Einstein doesn’t.

The Washington Post

Threat hunting

Create visibility

Threat hunting

Create visibility

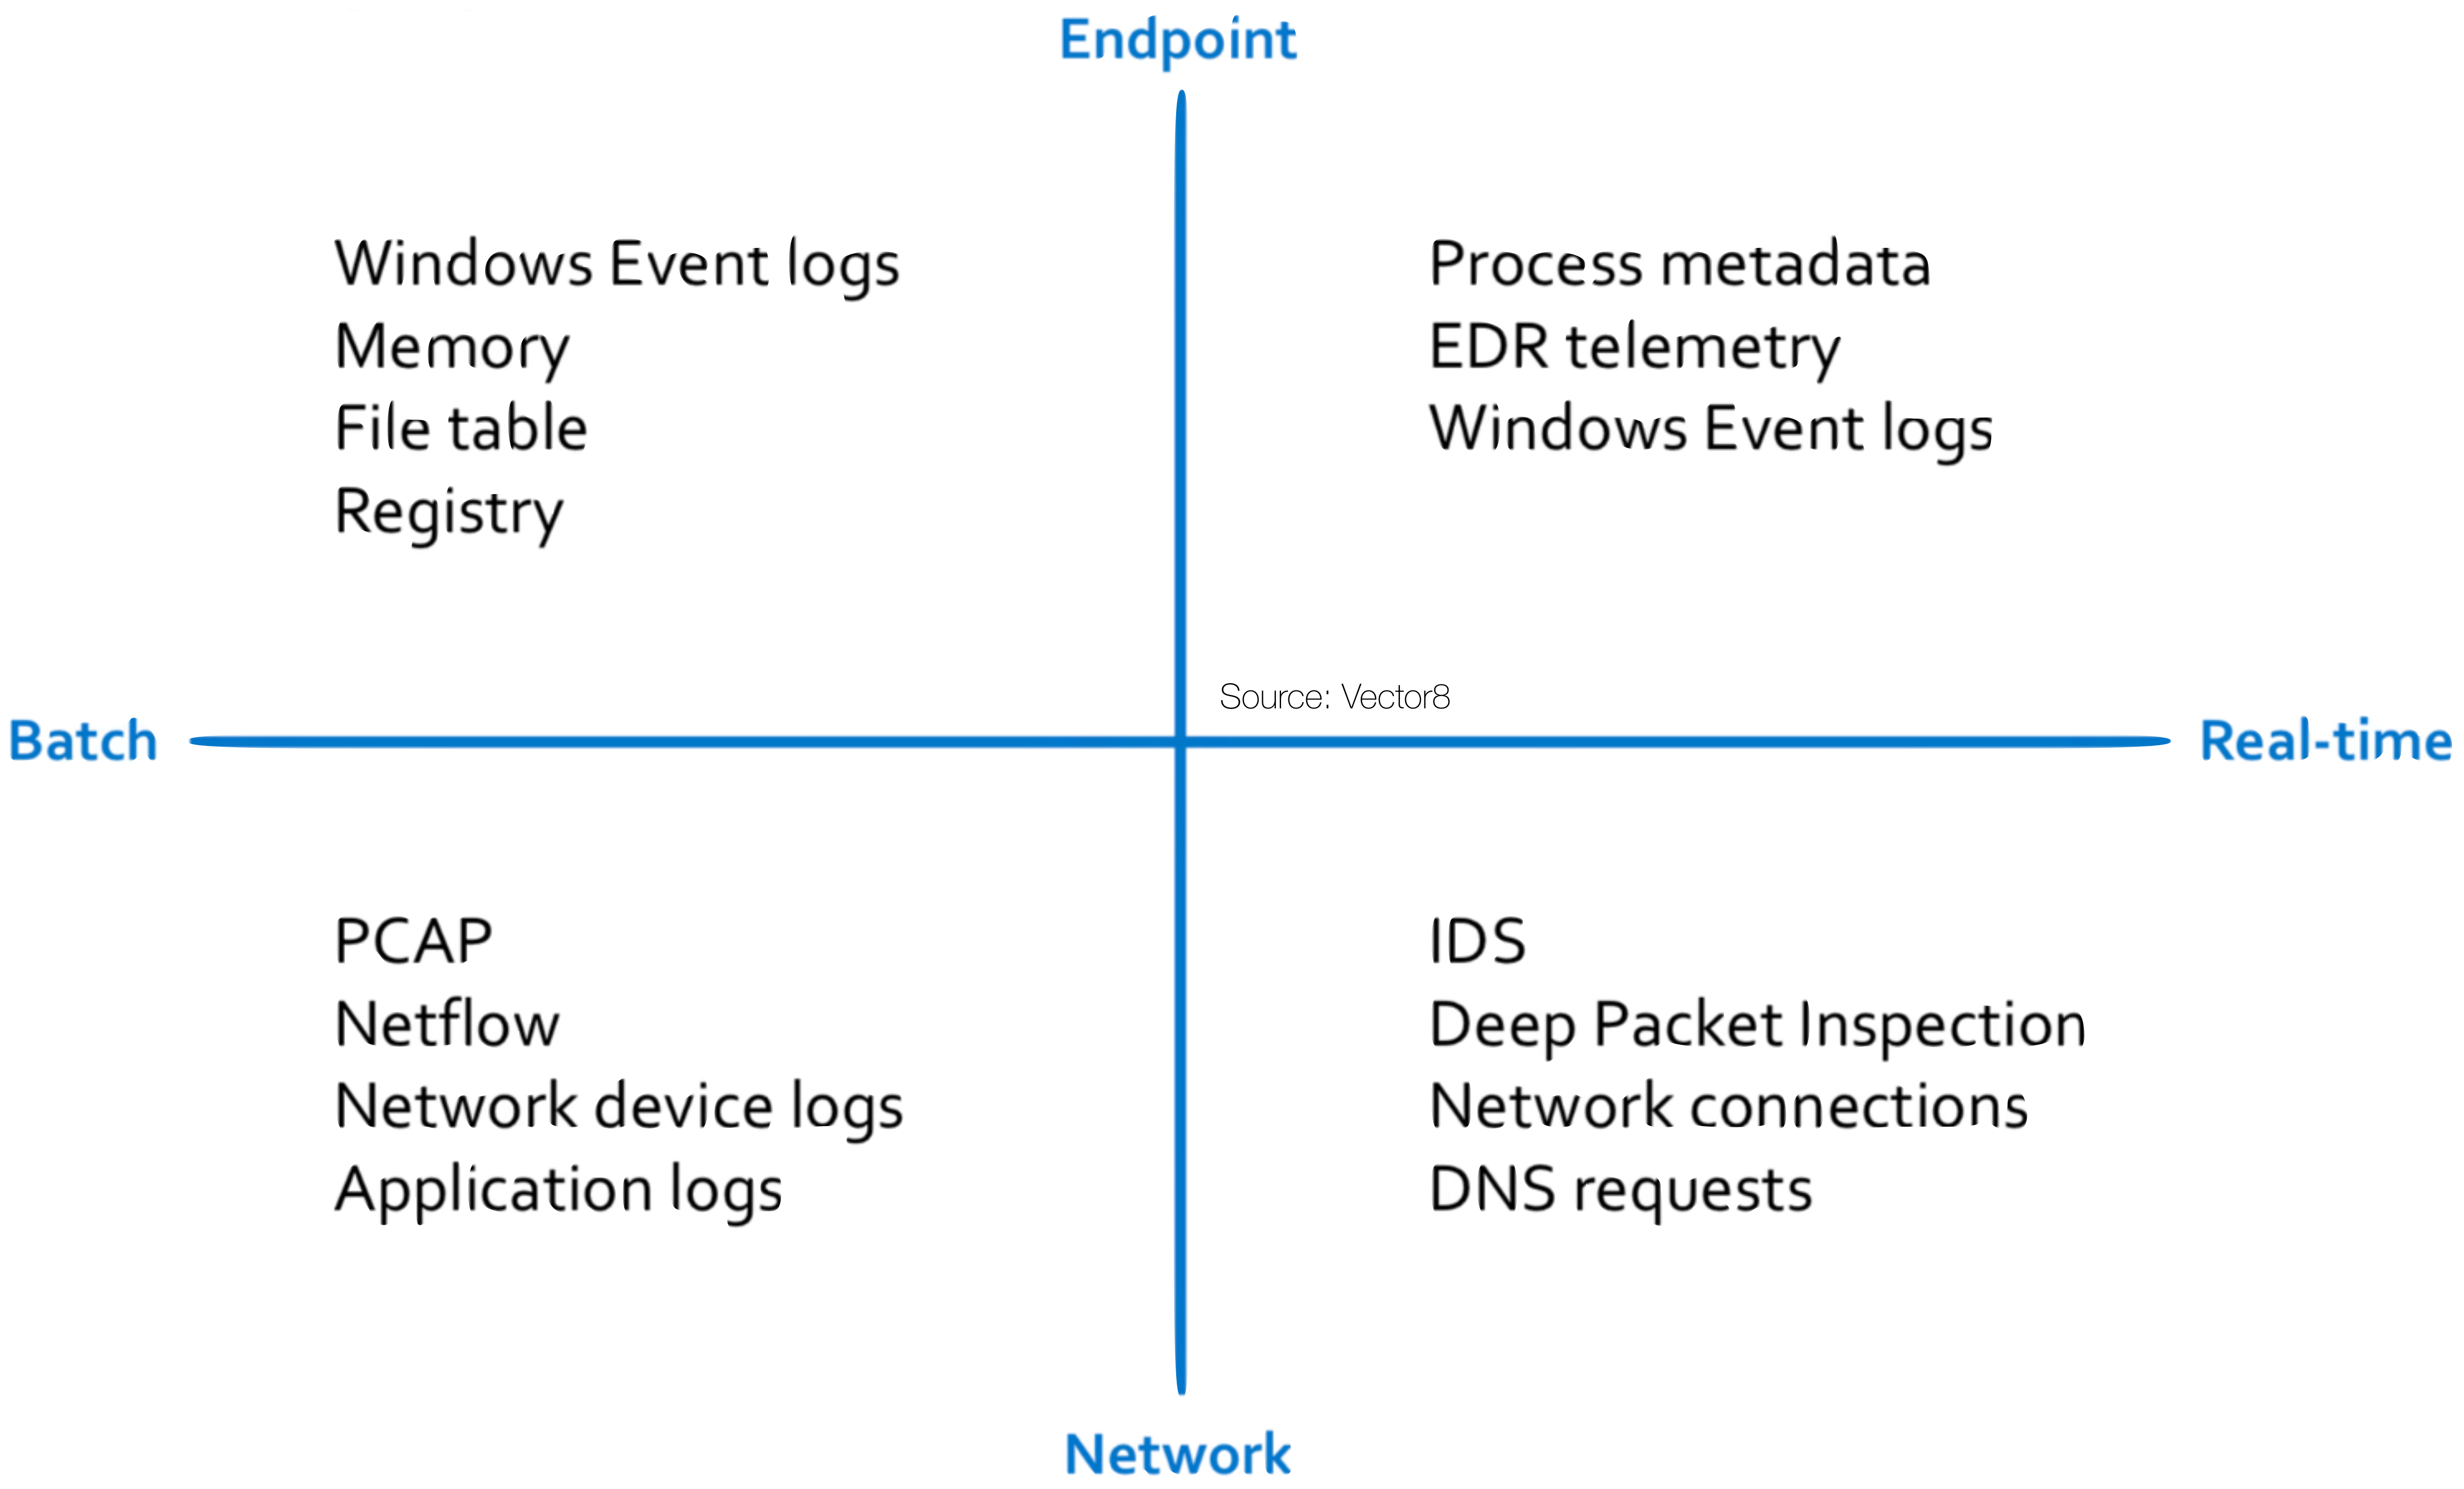

Endpoint

Windows Event logs
Memory
File table
Registry

Process metadata
EDR telemetry
Windows Event logs

Batch

Real-time

PCAP
Netflow
Network device logs
Application logs

IDS
Deep Packet Inspection
Network connections
DNS requests

Network

Source: Vector8

Threat hunting

Create visibility

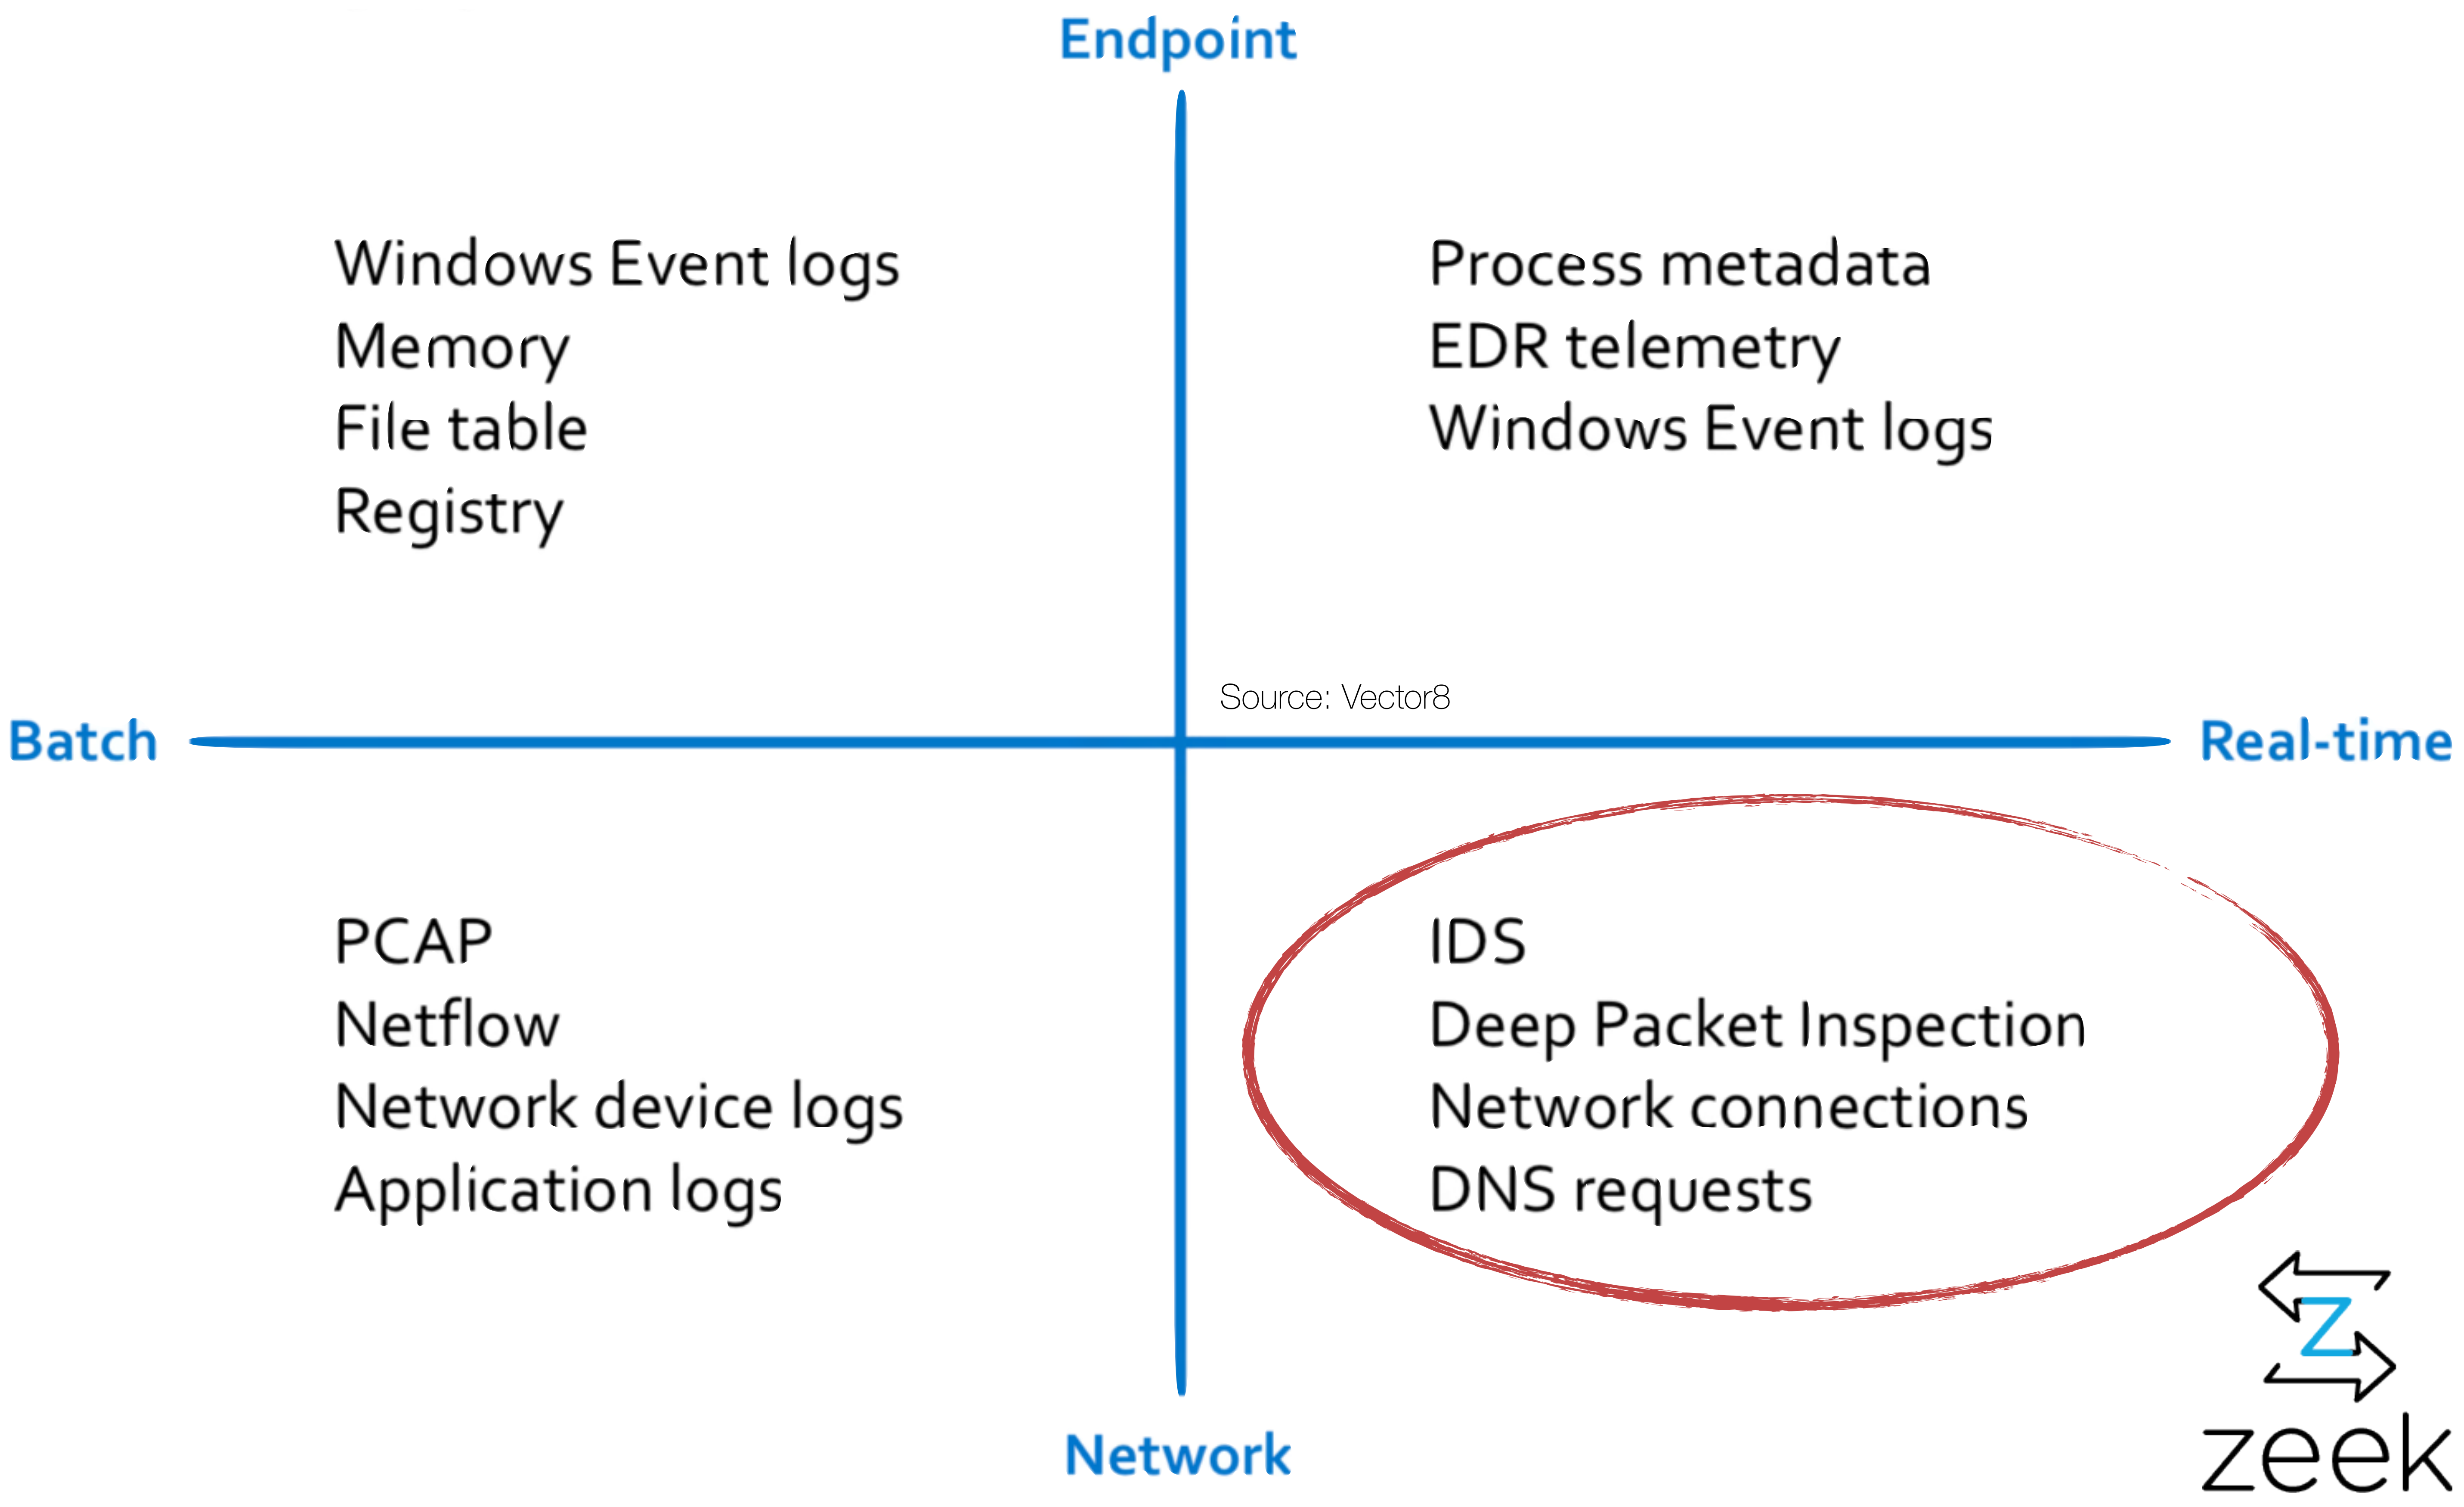

Windows Event logs
Memory
File table
Registry

Process metadata
EDR telemetry
Windows Event logs

PCAP
Netflow
Network device logs
Application logs

IDS
Deep Packet Inspection
Network connections
DNS requests

Source: Vector8

Threat hunting

Create visibility

elastic splunk>

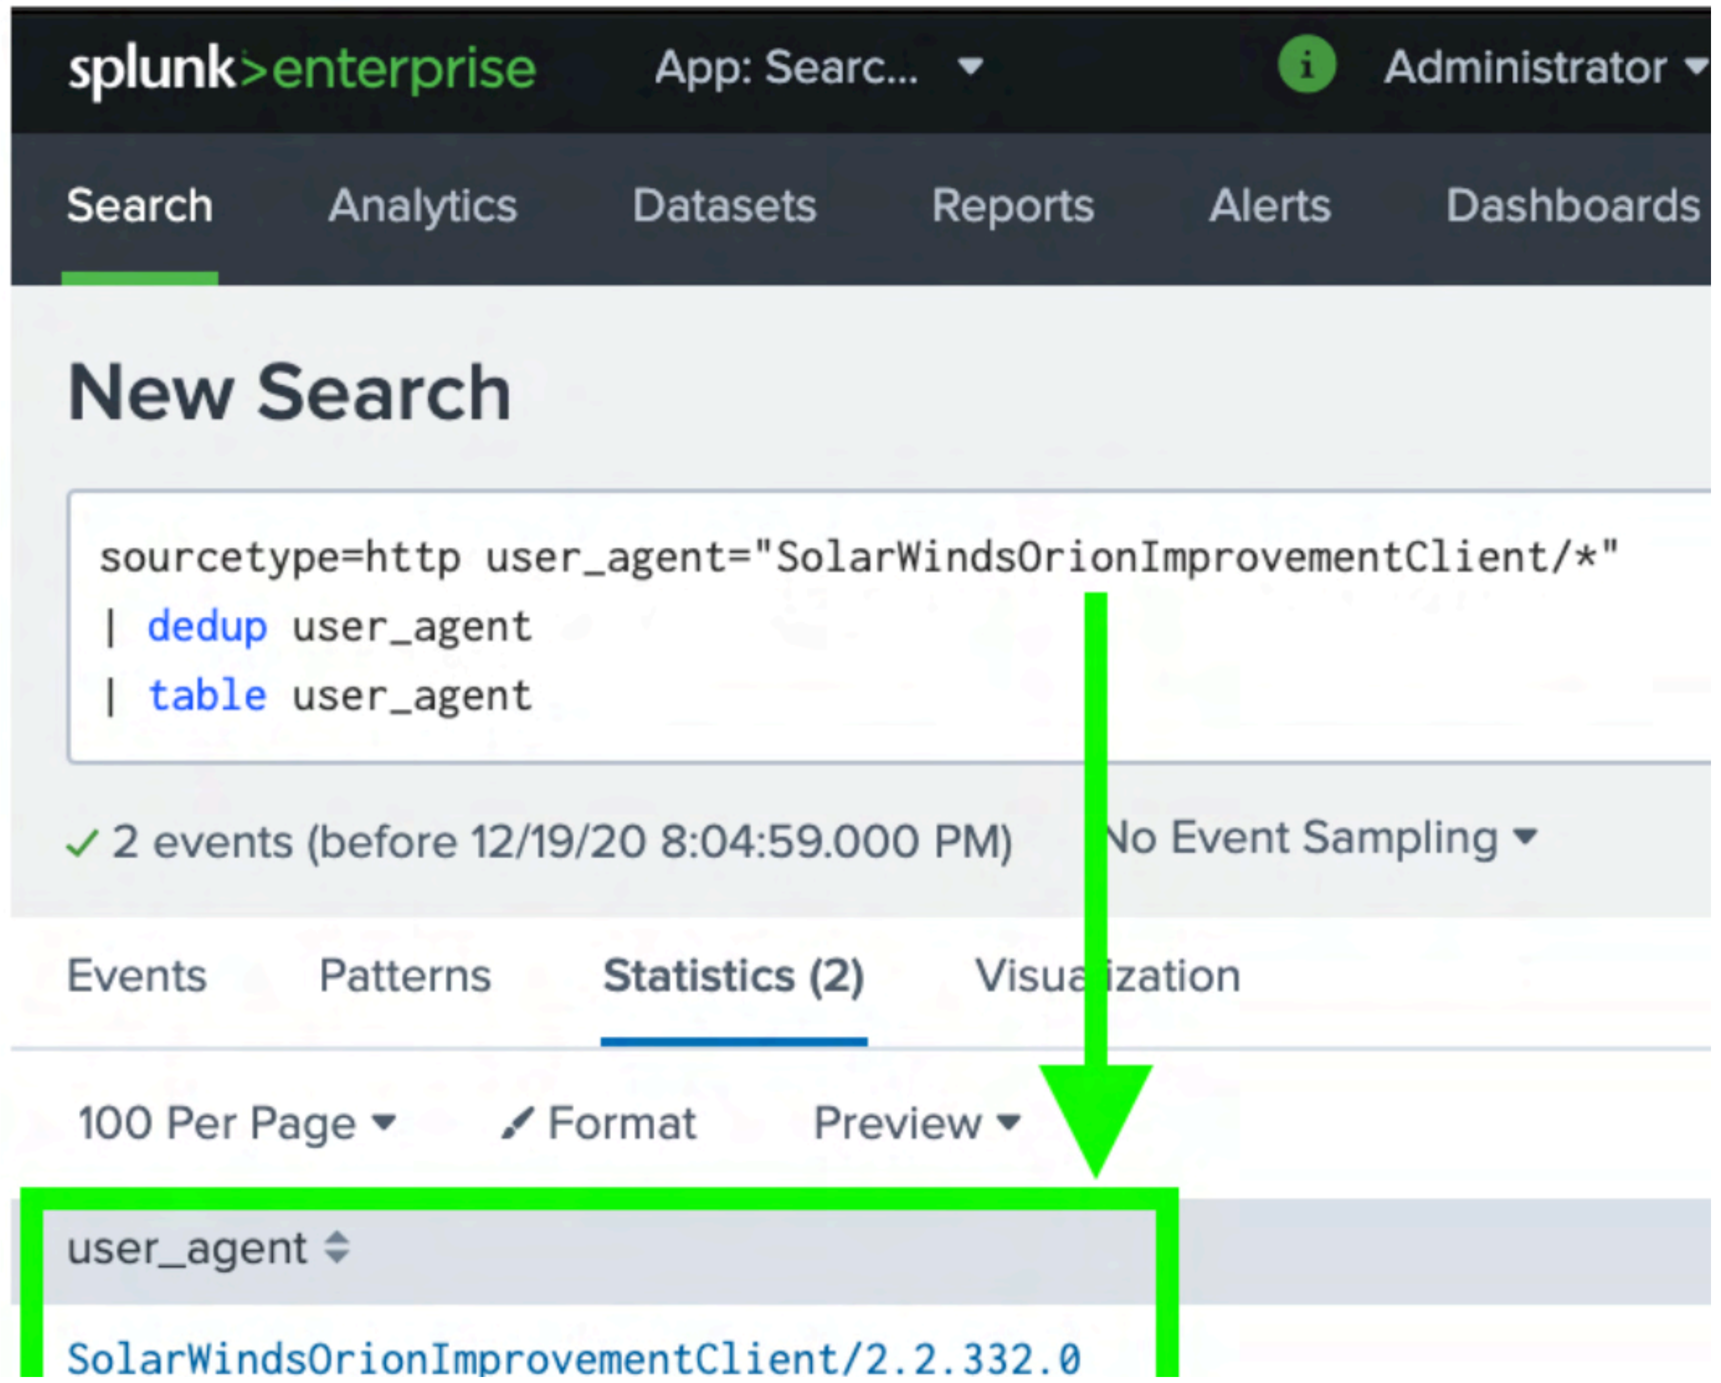

New Search

```
sourcetype=http user_agent="SolarWindsOrionImprovementClient/*"
| dedup user_agent
| table user_agent
```

✓ 2 events (before 12/19/20 8:04:59.000 PM)

user_agent
SolarWindsOrionImprovementClient/2.2.332.0

Threat hunting

Create visibility

Which IP did that box reach out to last week?

How many people received that email?

Who opened the suspicious attachment?

What DNS requests did the system issue?

When did we first see that CoC traffic?

Which systems did the person access?

Which services do normally run on a system?

Was the session encrypted?

What's the server name of that HTTPS endpoint?

Did the certificate check out ok?

Did they try to connect to our LDAP server?

Has somebody modified the file?

Threat hunting

Create visibility

elastic splunk>

Threat hunting

Create visibility

elastic splunk>

Let security team actively search for threats

Threat hunting

Create visibility

Let security team actively search for threats

Common types of “hunts”

- Hypothesis: “What if an attacker wanted to do *that*”?
- Trigger: “Something’s fishy …”
- Retrospective: “Where we hit by the same as FireEye”?

Threat hunting needs a highly skilled team

Threat hunting needs a highly skilled team

Threat hunting maturity model

Source: Sqrrl

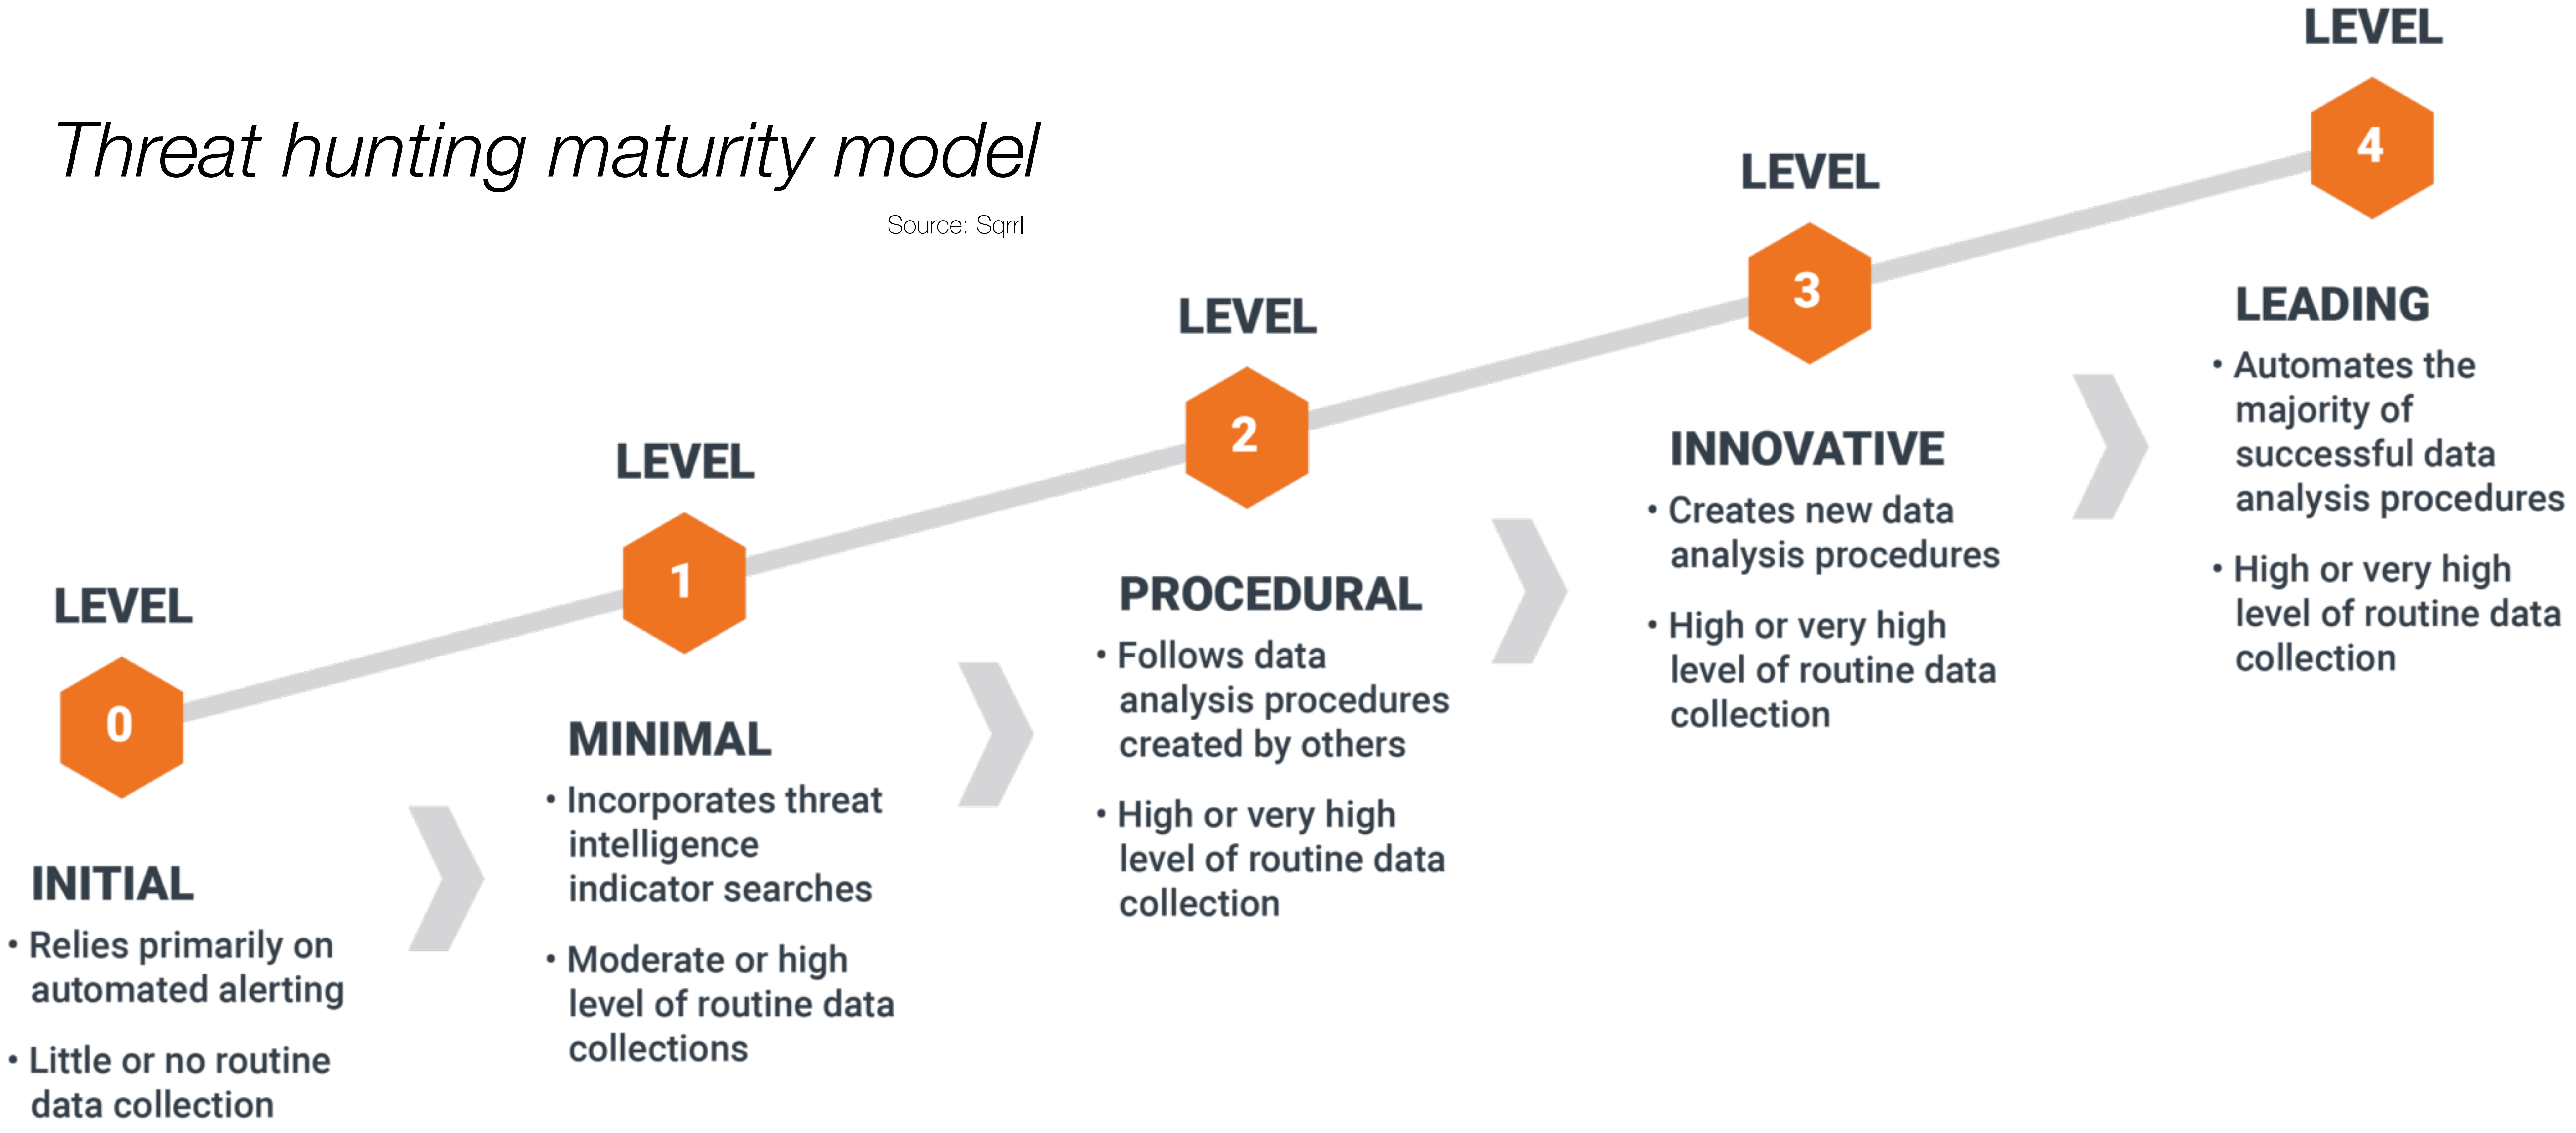

LEVEL 0

INITIAL

- Relies primarily on automated alerting
- Little or no routine data collection

LEVEL 1

MINIMAL

- Incorporates threat intelligence indicator searches
- Moderate or high level of routine data collections

LEVEL 2

PROCEDURAL

- Follows data analysis procedures created by others
- High or very high level of routine data collection

LEVEL 3

INNOVATIVE

- Creates new data analysis procedures
- High or very high level of routine data collection

LEVEL 4

LEADING

- Automates the majority of successful data analysis procedures
- High or very high level of routine data collection

Threat hunting needs a highly skilled team

Threat hunting maturity model

Source: Sqrrl

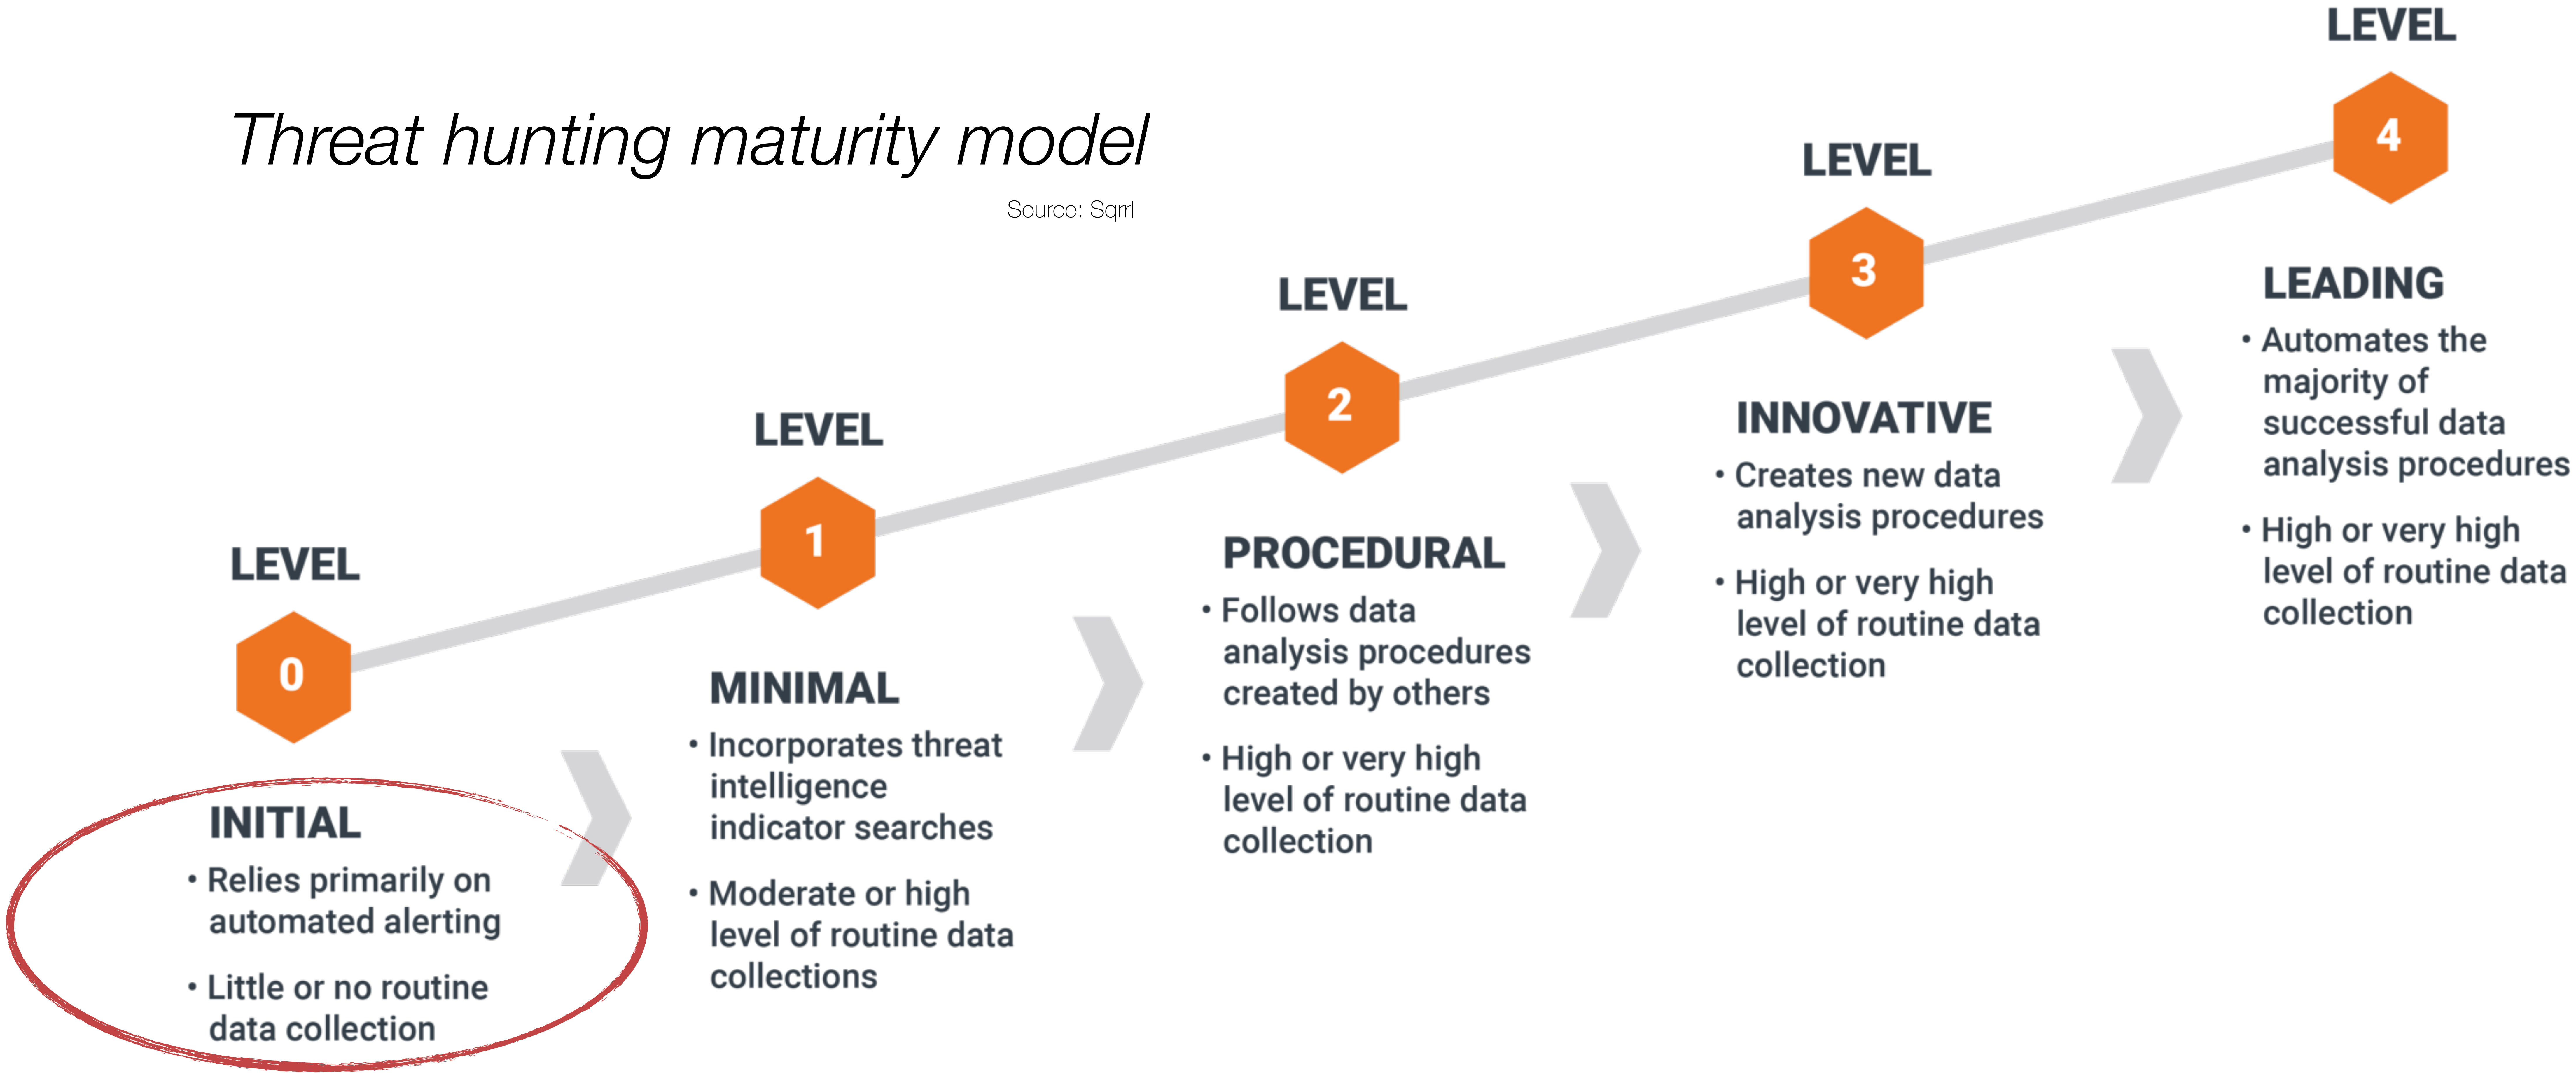

LEVEL 0

INITIAL
- Relies primarily on automated alerting
- Little or no routine data collection

LEVEL 1

MINIMAL
- Incorporates threat intelligence indicator searches
- Moderate or high level of routine data collections

LEVEL 2

PROCEDURAL
- Follows data analysis procedures created by others
- High or very high level of routine data collection

LEVEL 3

INNOVATIVE
- Creates new data analysis procedures
- High or very high level of routine data collection

LEVEL 4

LEADING
- Automates the majority of successful data analysis procedures
- High or very high level of routine data collection

Classic IDS deployment

Threat hunting needs a highly skilled team

Threat hunting maturity model

Source: Sqrrl

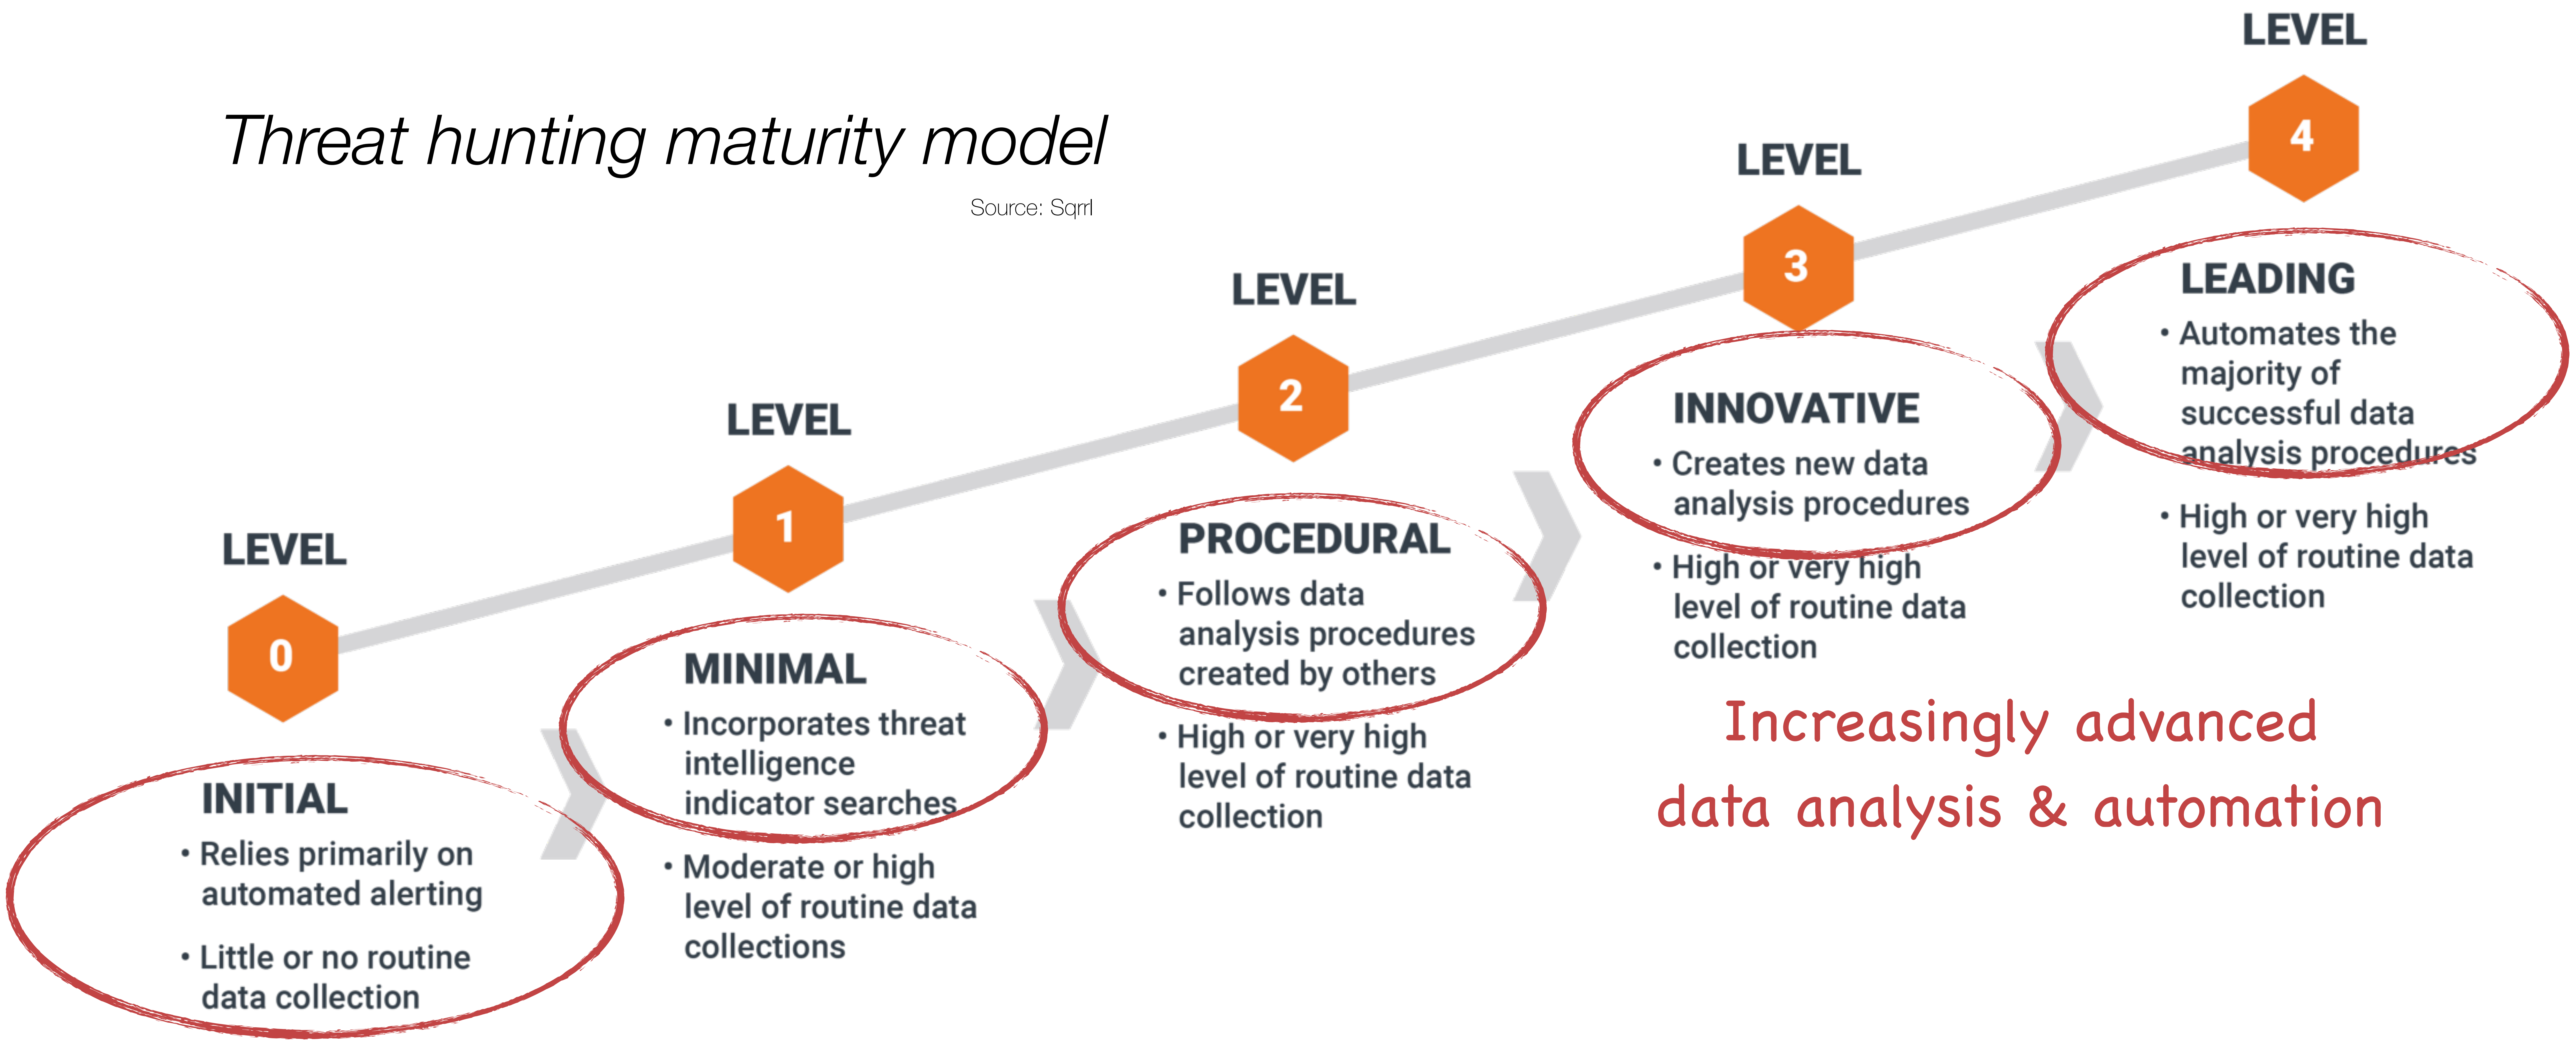

LEVEL 0

INITIAL

- Relies primarily on automated alerting
- Little or no routine data collection

LEVEL 1

MINIMAL

- Incorporates threat intelligence indicator searches
- Moderate or high level of routine data collections

LEVEL 2

PROCEDURAL

- Follows data analysis procedures created by others
- High or very high level of routine data collection

LEVEL 3

INNOVATIVE

- Creates new data analysis procedures
- High or very high level of routine data collection

LEVEL 4

LEADING

- Automates the majority of successful data analysis procedures
- High or very high level of routine data collection

Increasingly advanced data analysis & automation

Classic IDS deployment

Threat hunting needs a highly skilled team

Threat hunting maturity model

Source: Sqrrl

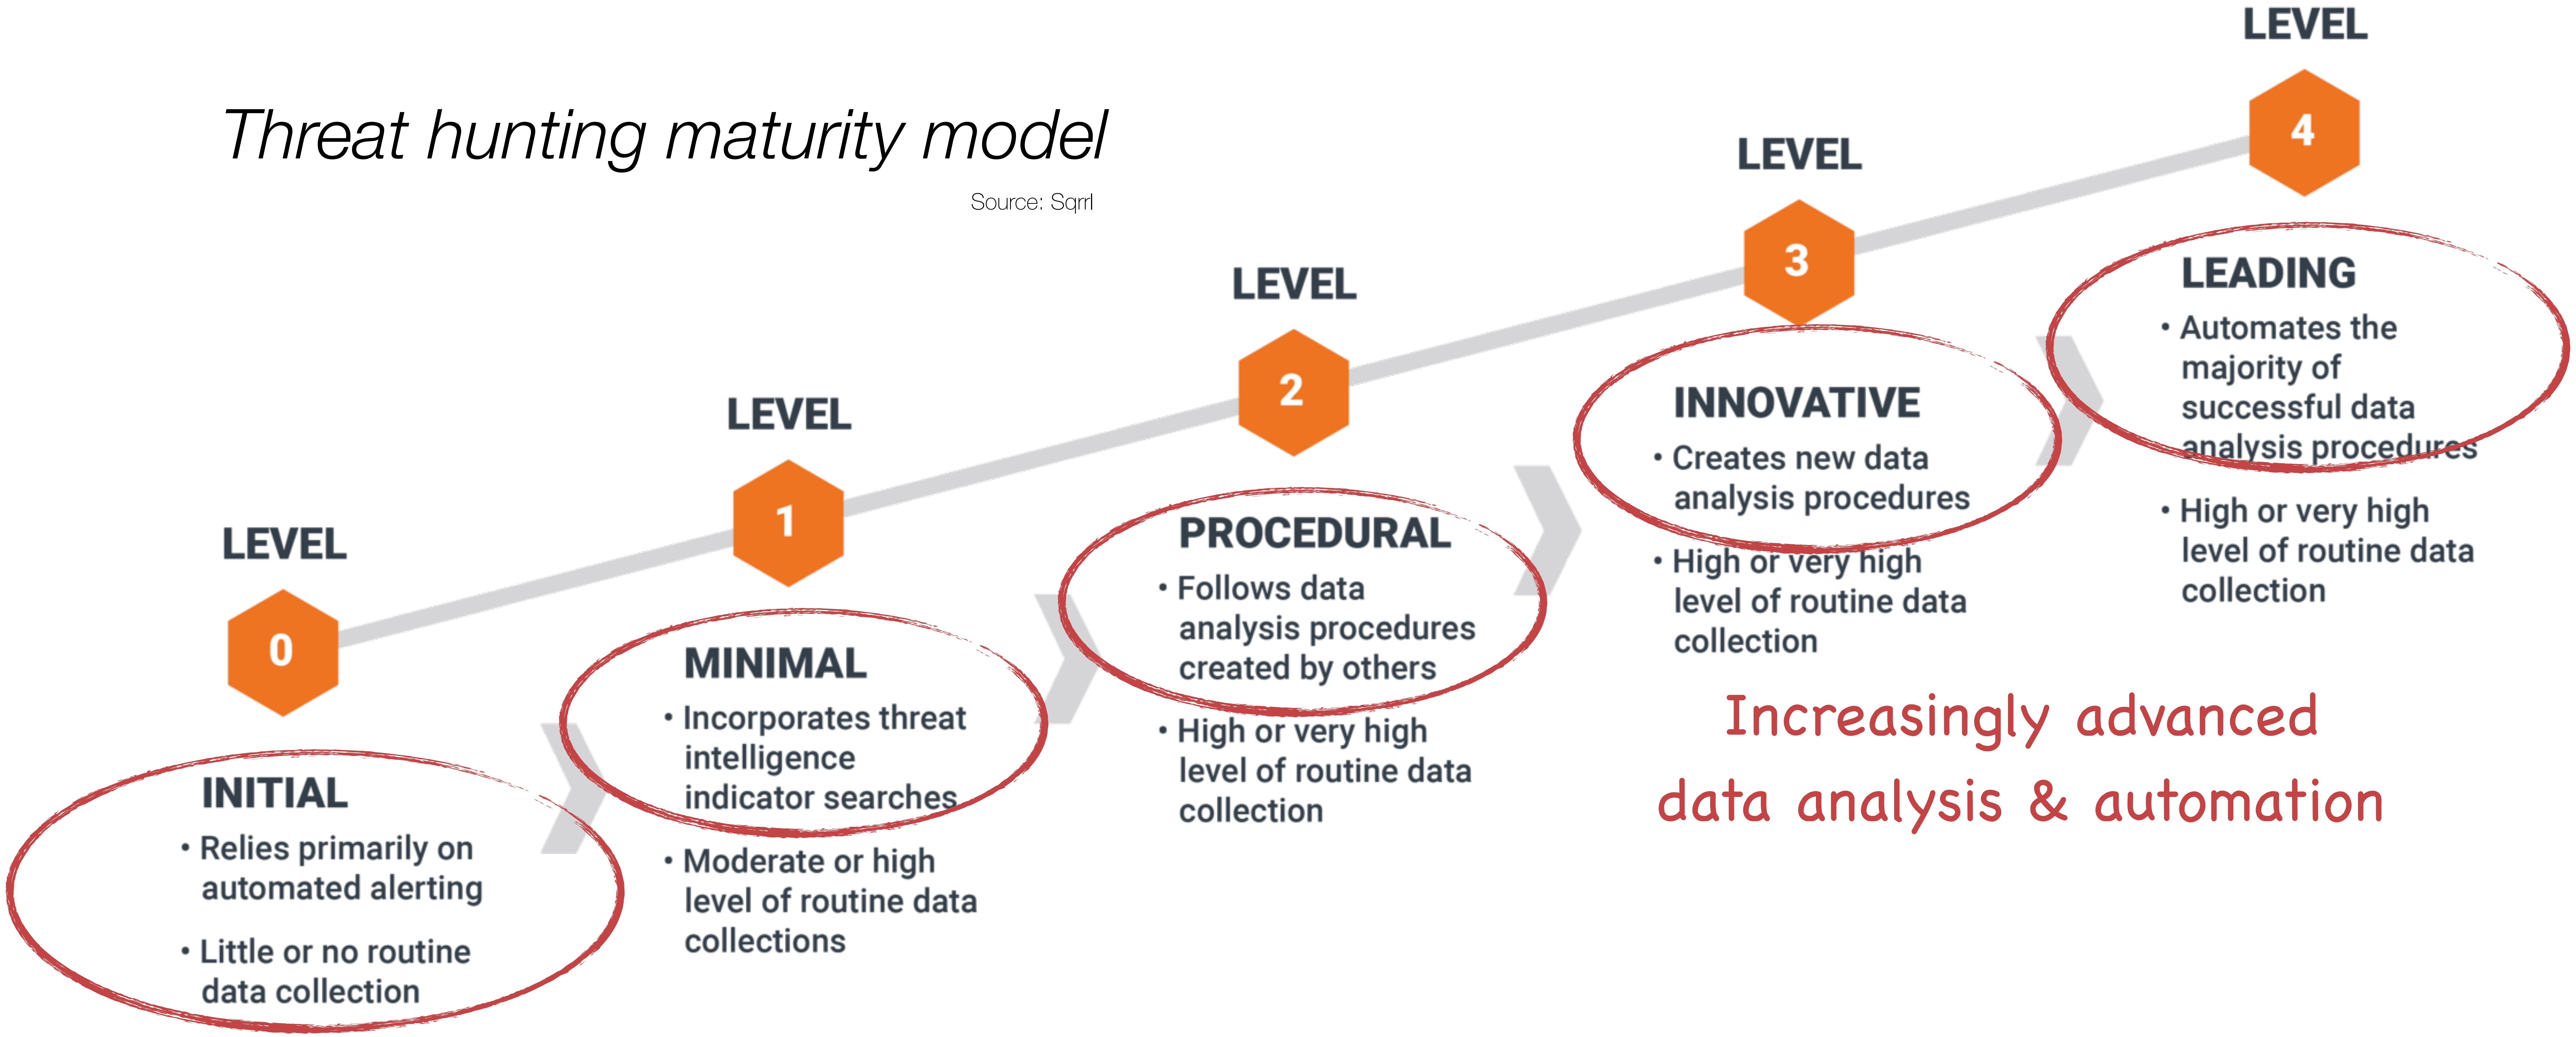

LEVEL 0

INITIAL

- Relies primarily on automated alerting
- Little or no routine data collection

LEVEL 1

MINIMAL

- Incorporates threat intelligence indicator searches
- Moderate or high level of routine data collections

LEVEL 2

PROCEDURAL

- Follows data analysis procedures created by others
- High or very high level of routine data collection

LEVEL 3

INNOVATIVE

- Creates new data analysis procedures
- High or very high level of routine data collection

LEVEL 4

LEADING

- Automates the majority of successful data analysis procedures
- High or very high level of routine data collection

Increasingly advanced data analysis & automation

Classic IDS de

Opportunity: Deploy AI to support the human analysts.

Beyond detection: A New Role for AI

Support the hunters

Narrow classifiers as triggers

Activity known to remain quite stable

- Service availability
- SSL certificates
- Executables on a server

Individual features with characteristic distributions

- URL parameters
- DNS lookups
- Communication timing (e.g., interactive logins)

Variations of known attacks

- Pre-canned attack tools
- Phishing emails

Example: Typo squatting

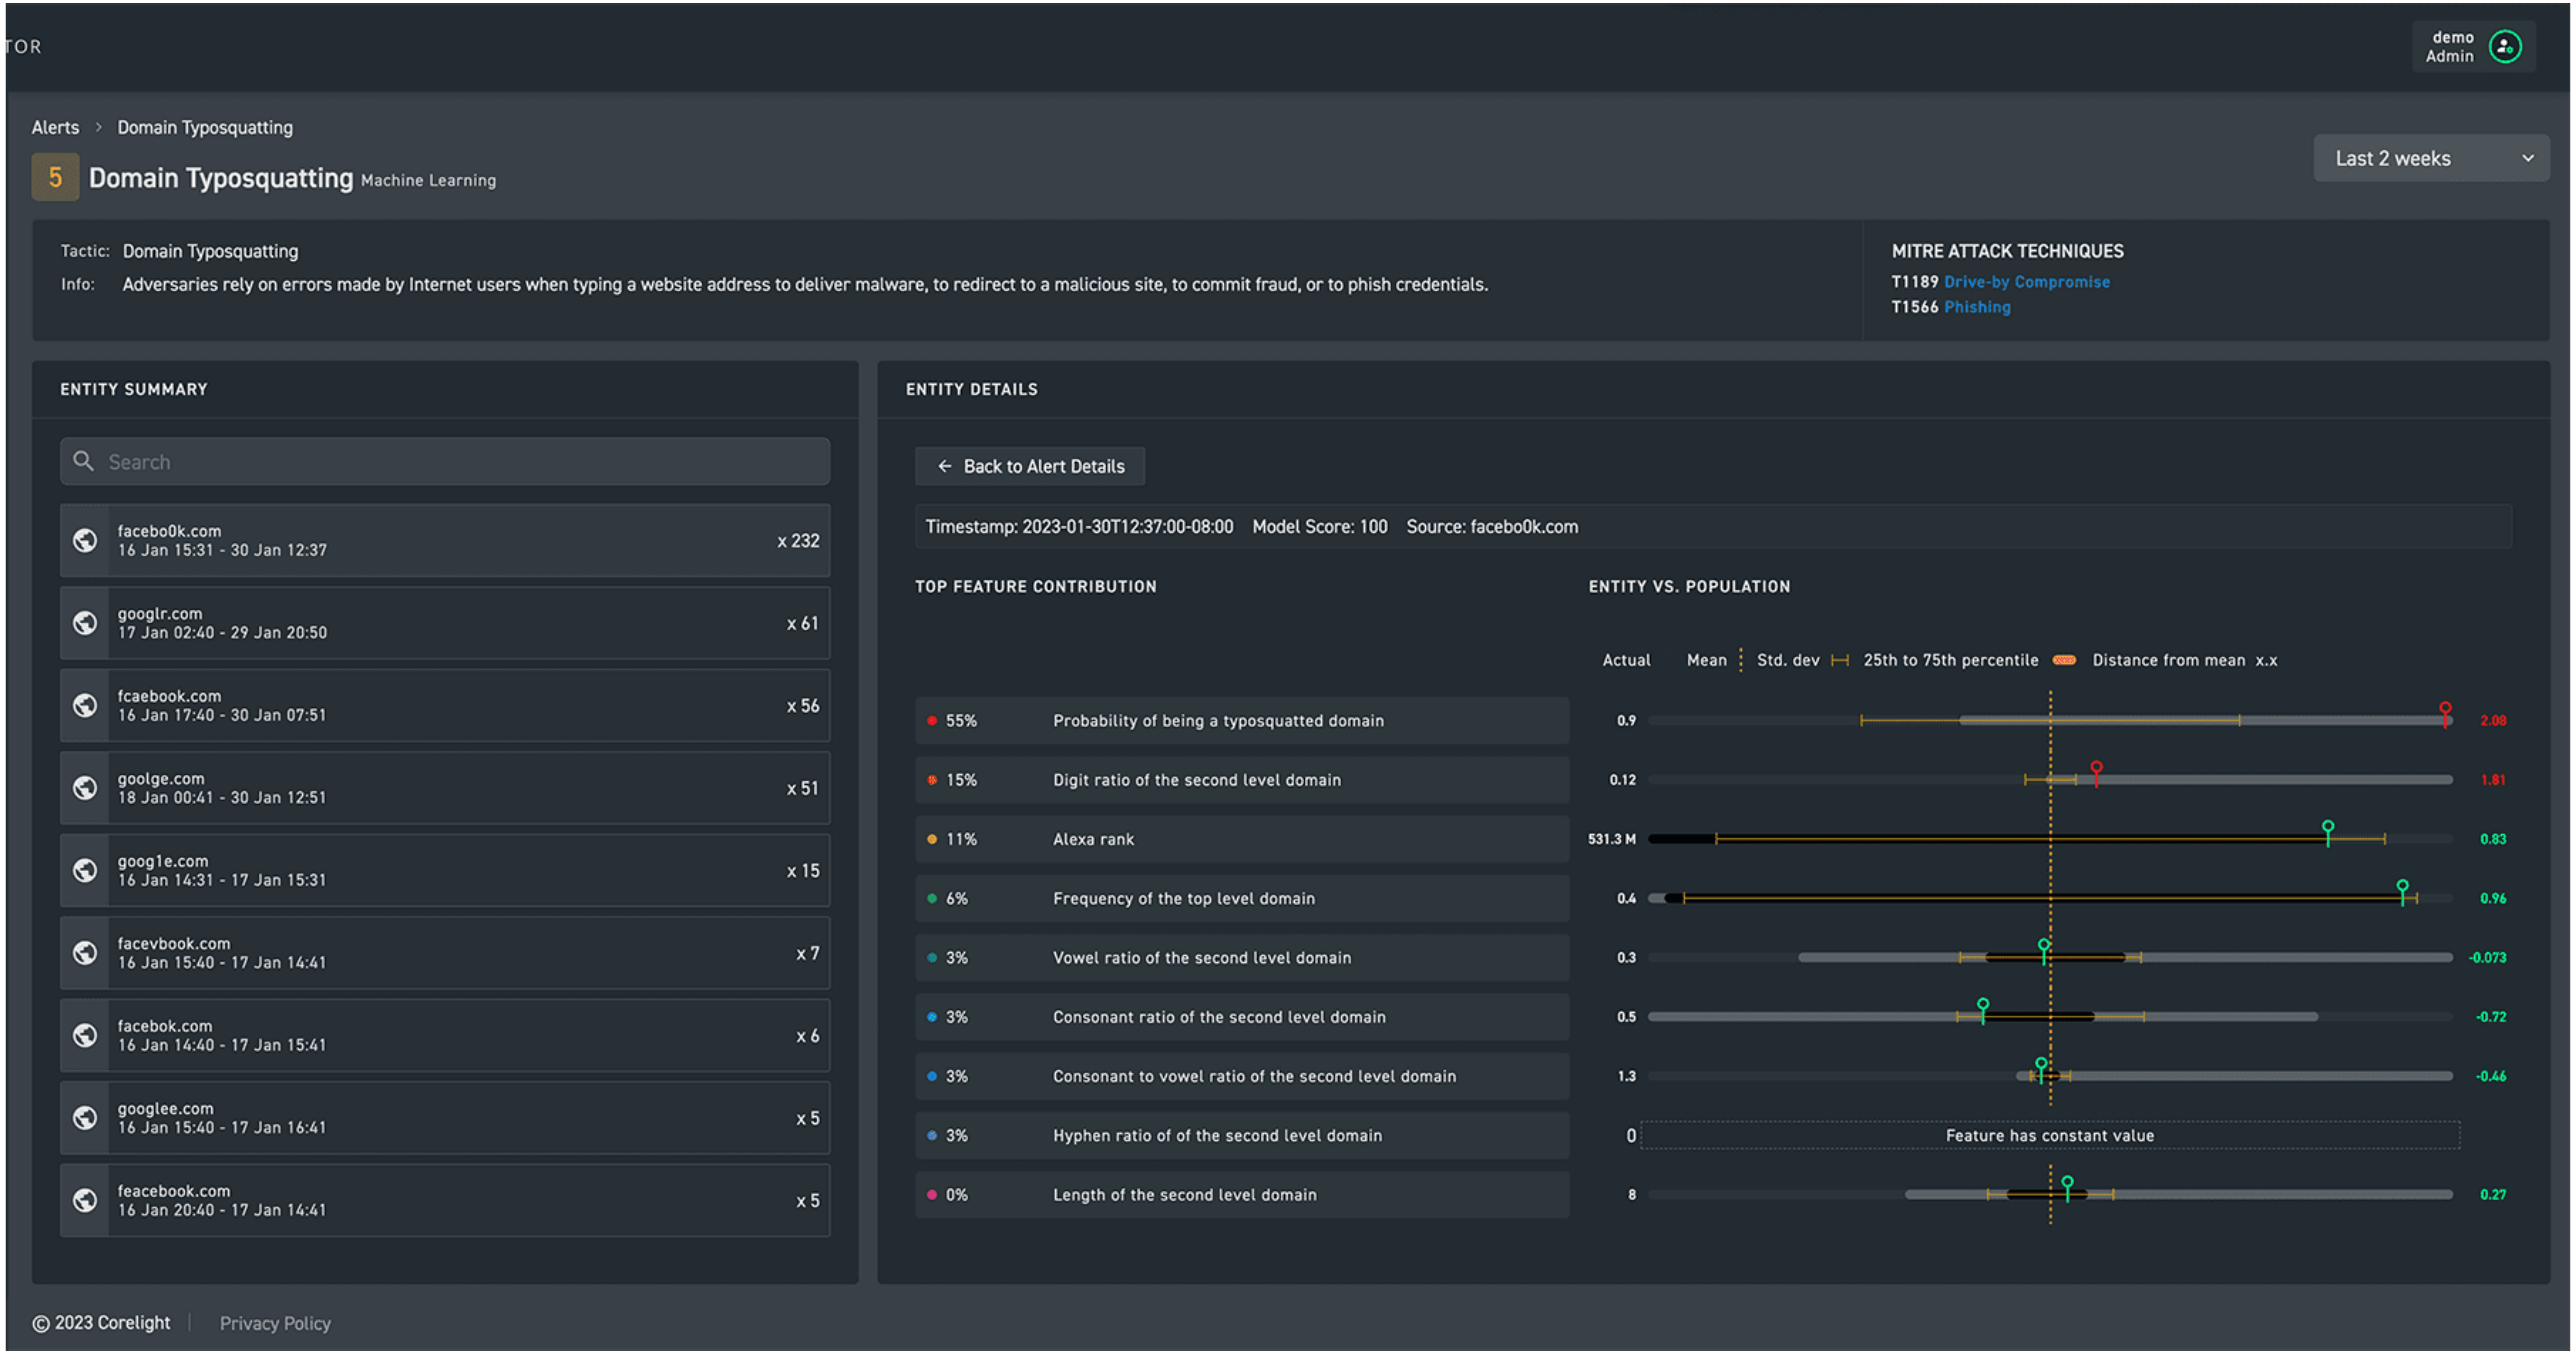

Alerts > Domain Typosquatting

5 Domain Typosquatting Machine Learning

Tactic: Domain Typosquatting

Info: Adversaries rely on errors made by Internet users when typing a website address to deliver malware, to redirect to a malicious site, to commit fraud, or to phish credentials.

MITRE ATTACK TECHNIQUES

T1189 Drive-by Compromise

T1566 Phishing

ENTITY SUMMARY

Domain	Time range	Count
facebo0k.com	16 Jan 15:31 - 30 Jan 12:37	x 232
googlr.com	17 Jan 02:40 - 29 Jan 20:50	x 61
fcaebook.com	16 Jan 17:40 - 30 Jan 07:51	x 56
goolge.com	18 Jan 00:41 - 30 Jan 12:51	x 51
goog1e.com	16 Jan 14:31 - 17 Jan 15:31	x 15
facevbook.com	16 Jan 15:40 - 17 Jan 14:41	x 7
facebok.com	16 Jan 14:40 - 17 Jan 15:41	x 6
googlee.com	16 Jan 15:40 - 17 Jan 16:41	x 5
feacebook.com	16 Jan 20:40 - 17 Jan 14:41	x 5

ENTITY DETAILS

Timestamp: 2023-01-30T12:37:00-08:00 Model Score: 100 Source: facebo0k.com

TOP FEATURE CONTRIBUTION / ENTITY VS. POPULATION

Actual | Mean ⋮ | Std. dev ├┤ | 25th to 75th percentile | Distance from mean x.x

TOP FEATURE CONTRIBUTION		Actual	Distance from mean
55%	Probability of being a typosquatted domain	0.9	2.08
15%	Digit ratio of the second level domain	0.12	1.81
11%	Alexa rank	531.3 M	0.83
6%	Frequency of the top level domain	0.4	0.96
3%	Vowel ratio of the second level domain	0.3	-0.073
3%	Consonant ratio of the second level domain	0.5	-0.72
3%	Consonant to vowel ratio of the second level domain	1.3	-0.46
3%	Hyphen ratio of of the second level domain	0	Feature has constant value
0%	Length of the second level domain	8	0.27

© 2023 Corelight | Privacy Policy

Source: Corelight

Generative AI: Guiding the analyst

Lower the bar for effective threat hunting

Leverage expertise of more advanced organizations

Generative AI: Guiding the analyst

Lower the bar for effective threat hunting

Leverage expertise of more advanced organizations

Threat hunting questions

- What does this trigger mean?
- Where can I find out more about this?
- How likely is this malicious? How to confirm?
- Is my host normally be doing *that*?
- What entities I should I focus on?
- What hunts are my peers doing these days?

Incident response & triage

- How bad is it?
- What are the next steps now?
- What do I need to do to clean up?
- How do find this next time?
- Write report: findings, impact, mitigation

Baby steps: Explaining Suricata rules

Baby steps: Explaining Suricata rules

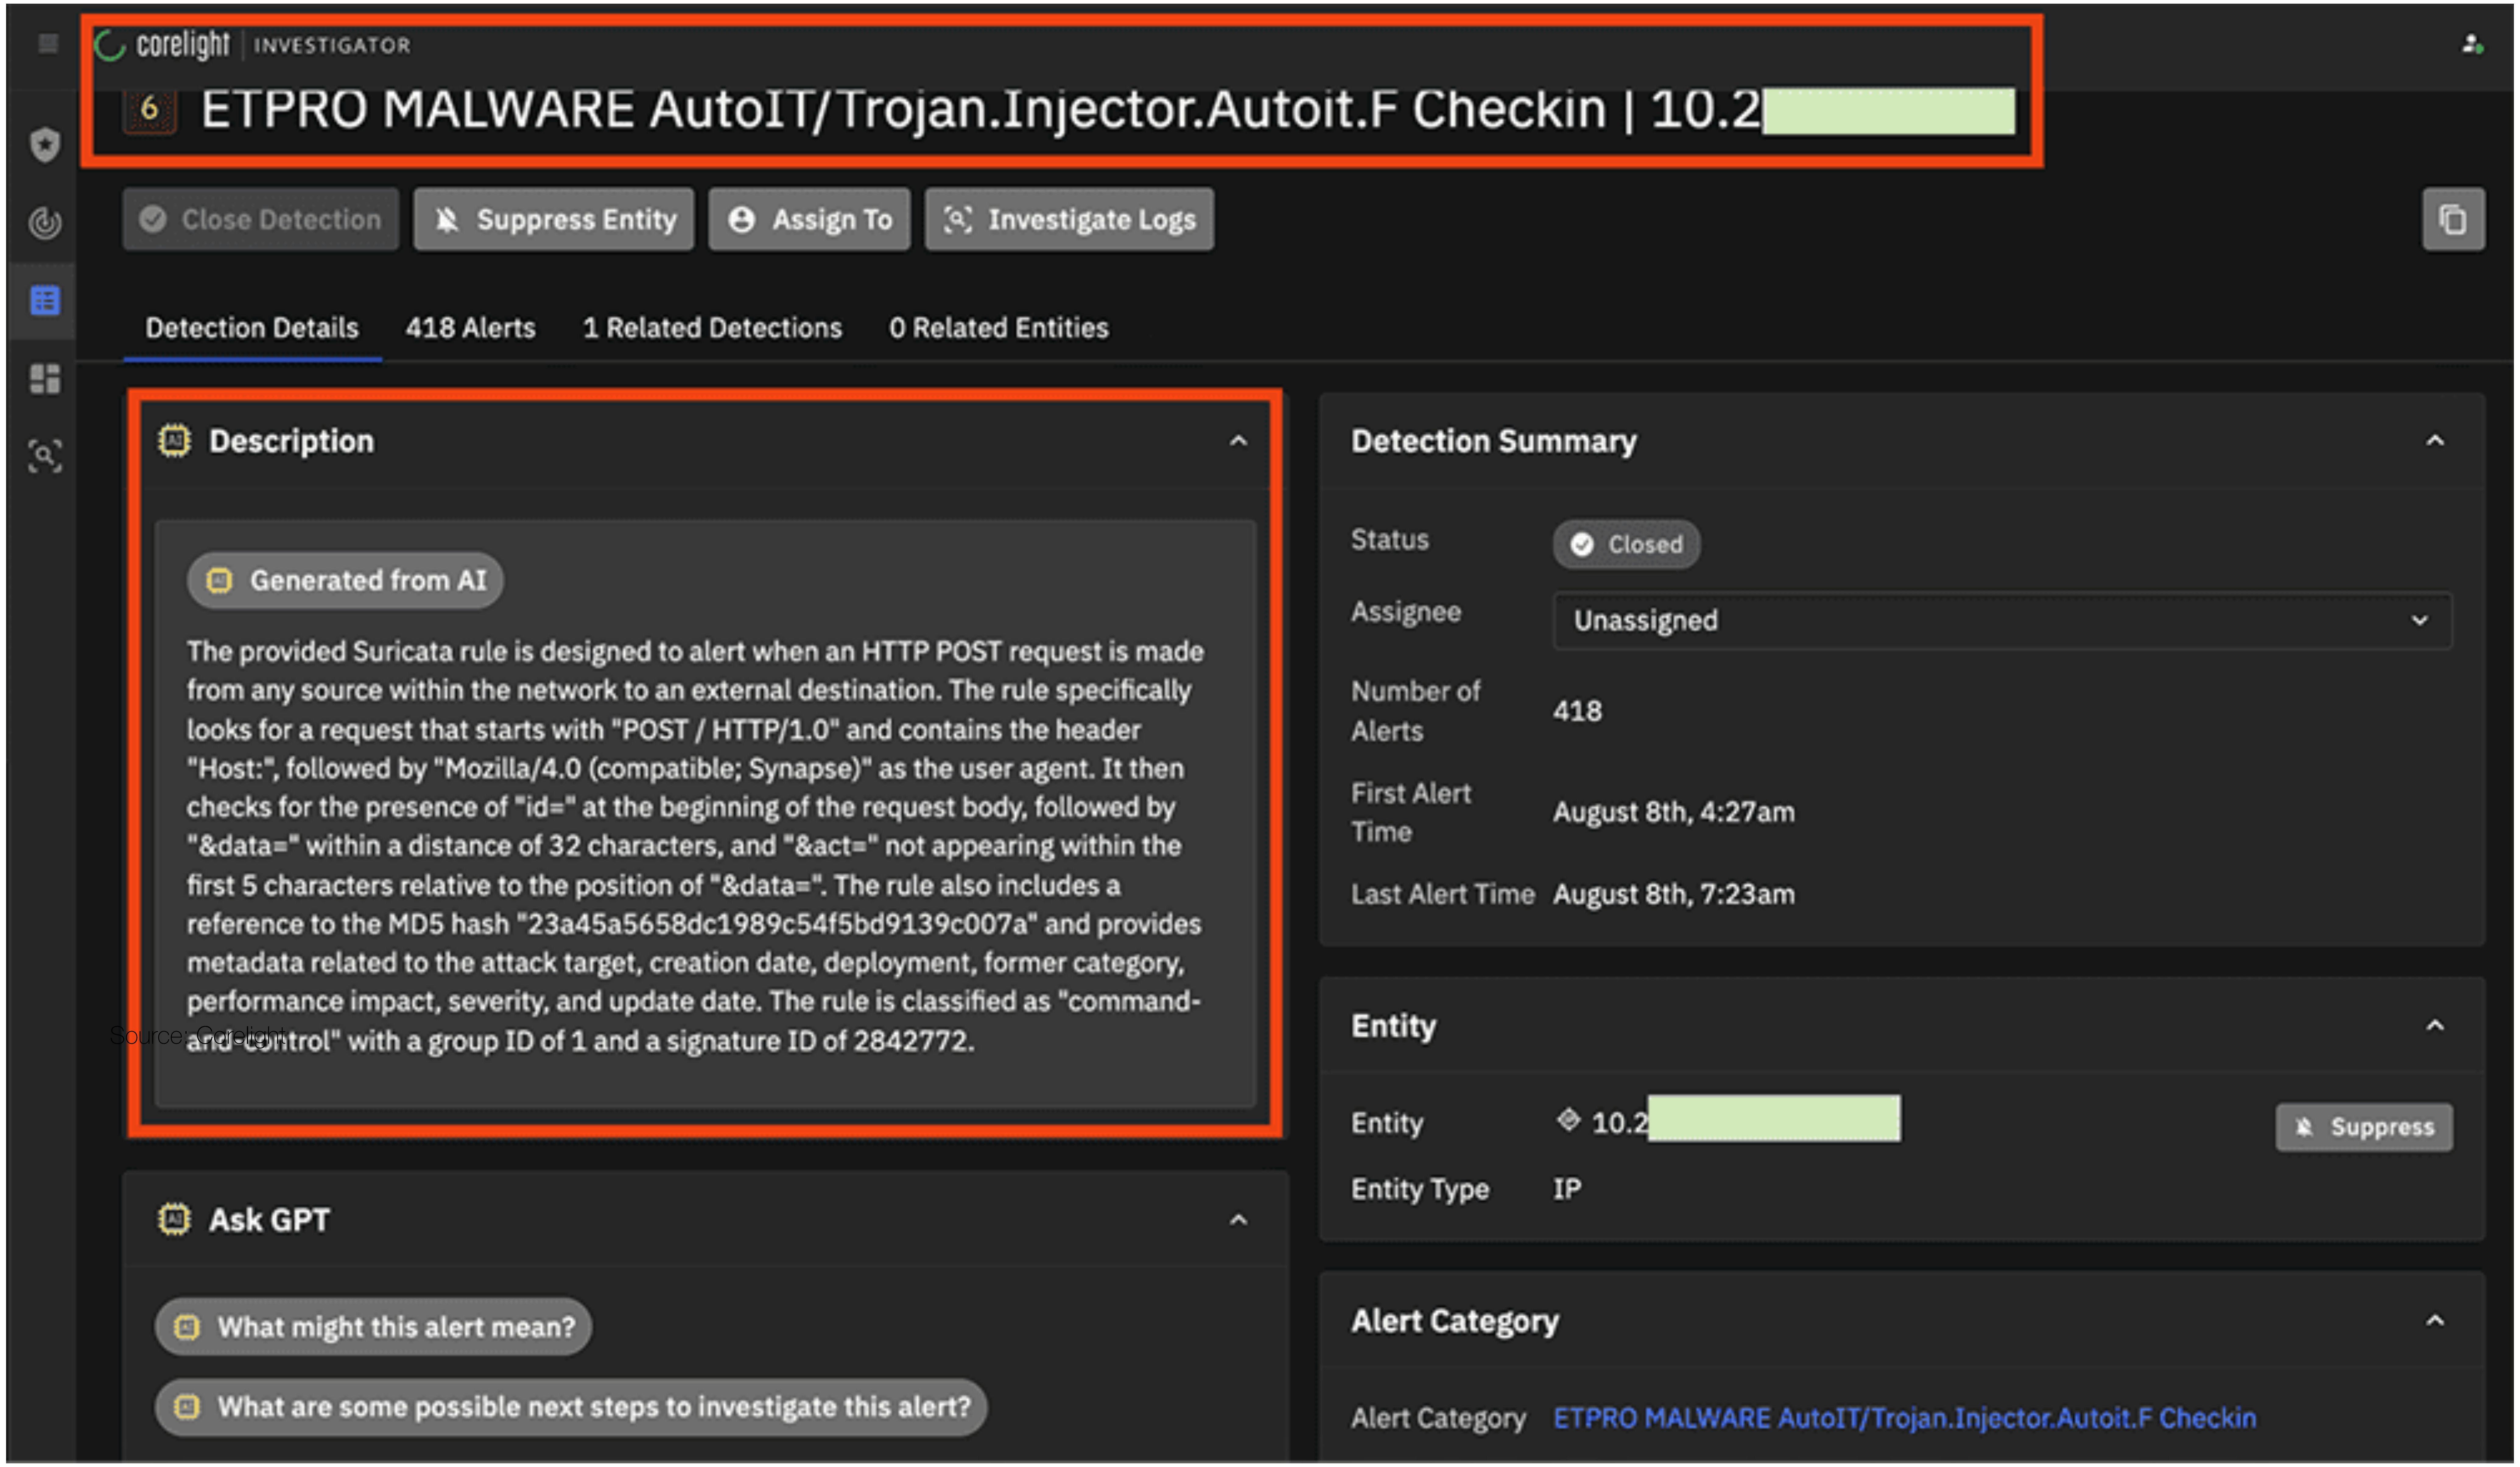

6 ETPRO MALWARE AutoIT/Trojan.Injector.Autoit.F Checkin | 10.2[redacted]

Description

Generated from AI

The provided Suricata rule is designed to alert when an HTTP POST request is made from any source within the network to an external destination. The rule specifically looks for a request that starts with "POST / HTTP/1.0" and contains the header "Host:", followed by "Mozilla/4.0 (compatible; Synapse)" as the user agent. It then checks for the presence of "id=" at the beginning of the request body, followed by "&data=" within a distance of 32 characters, and "&act=" not appearing within the first 5 characters relative to the position of "&data=". The rule also includes a reference to the MD5 hash "23a45a5658dc1989c54f5bd9139c007a" and provides metadata related to the attack target, creation date, deployment, former category, performance impact, severity, and update date. The rule is classified as "command-and-control" with a group ID of 1 and a signature ID of 2842772.

Source: Corelight

Detection Summary

Status	Closed
Assignee	Unassigned
Number of Alerts	418
First Alert Time	August 8th, 4:27am
Last Alert Time	August 8th, 7:23am

Entity

Entity	10.2[redacted]
Entity Type	IP

Alert Category

Alert Category	ETPRO MALWARE AutoIT/Trojan.Injector.Autoit.F Checkin

Baby steps: Explaining Suricata rules

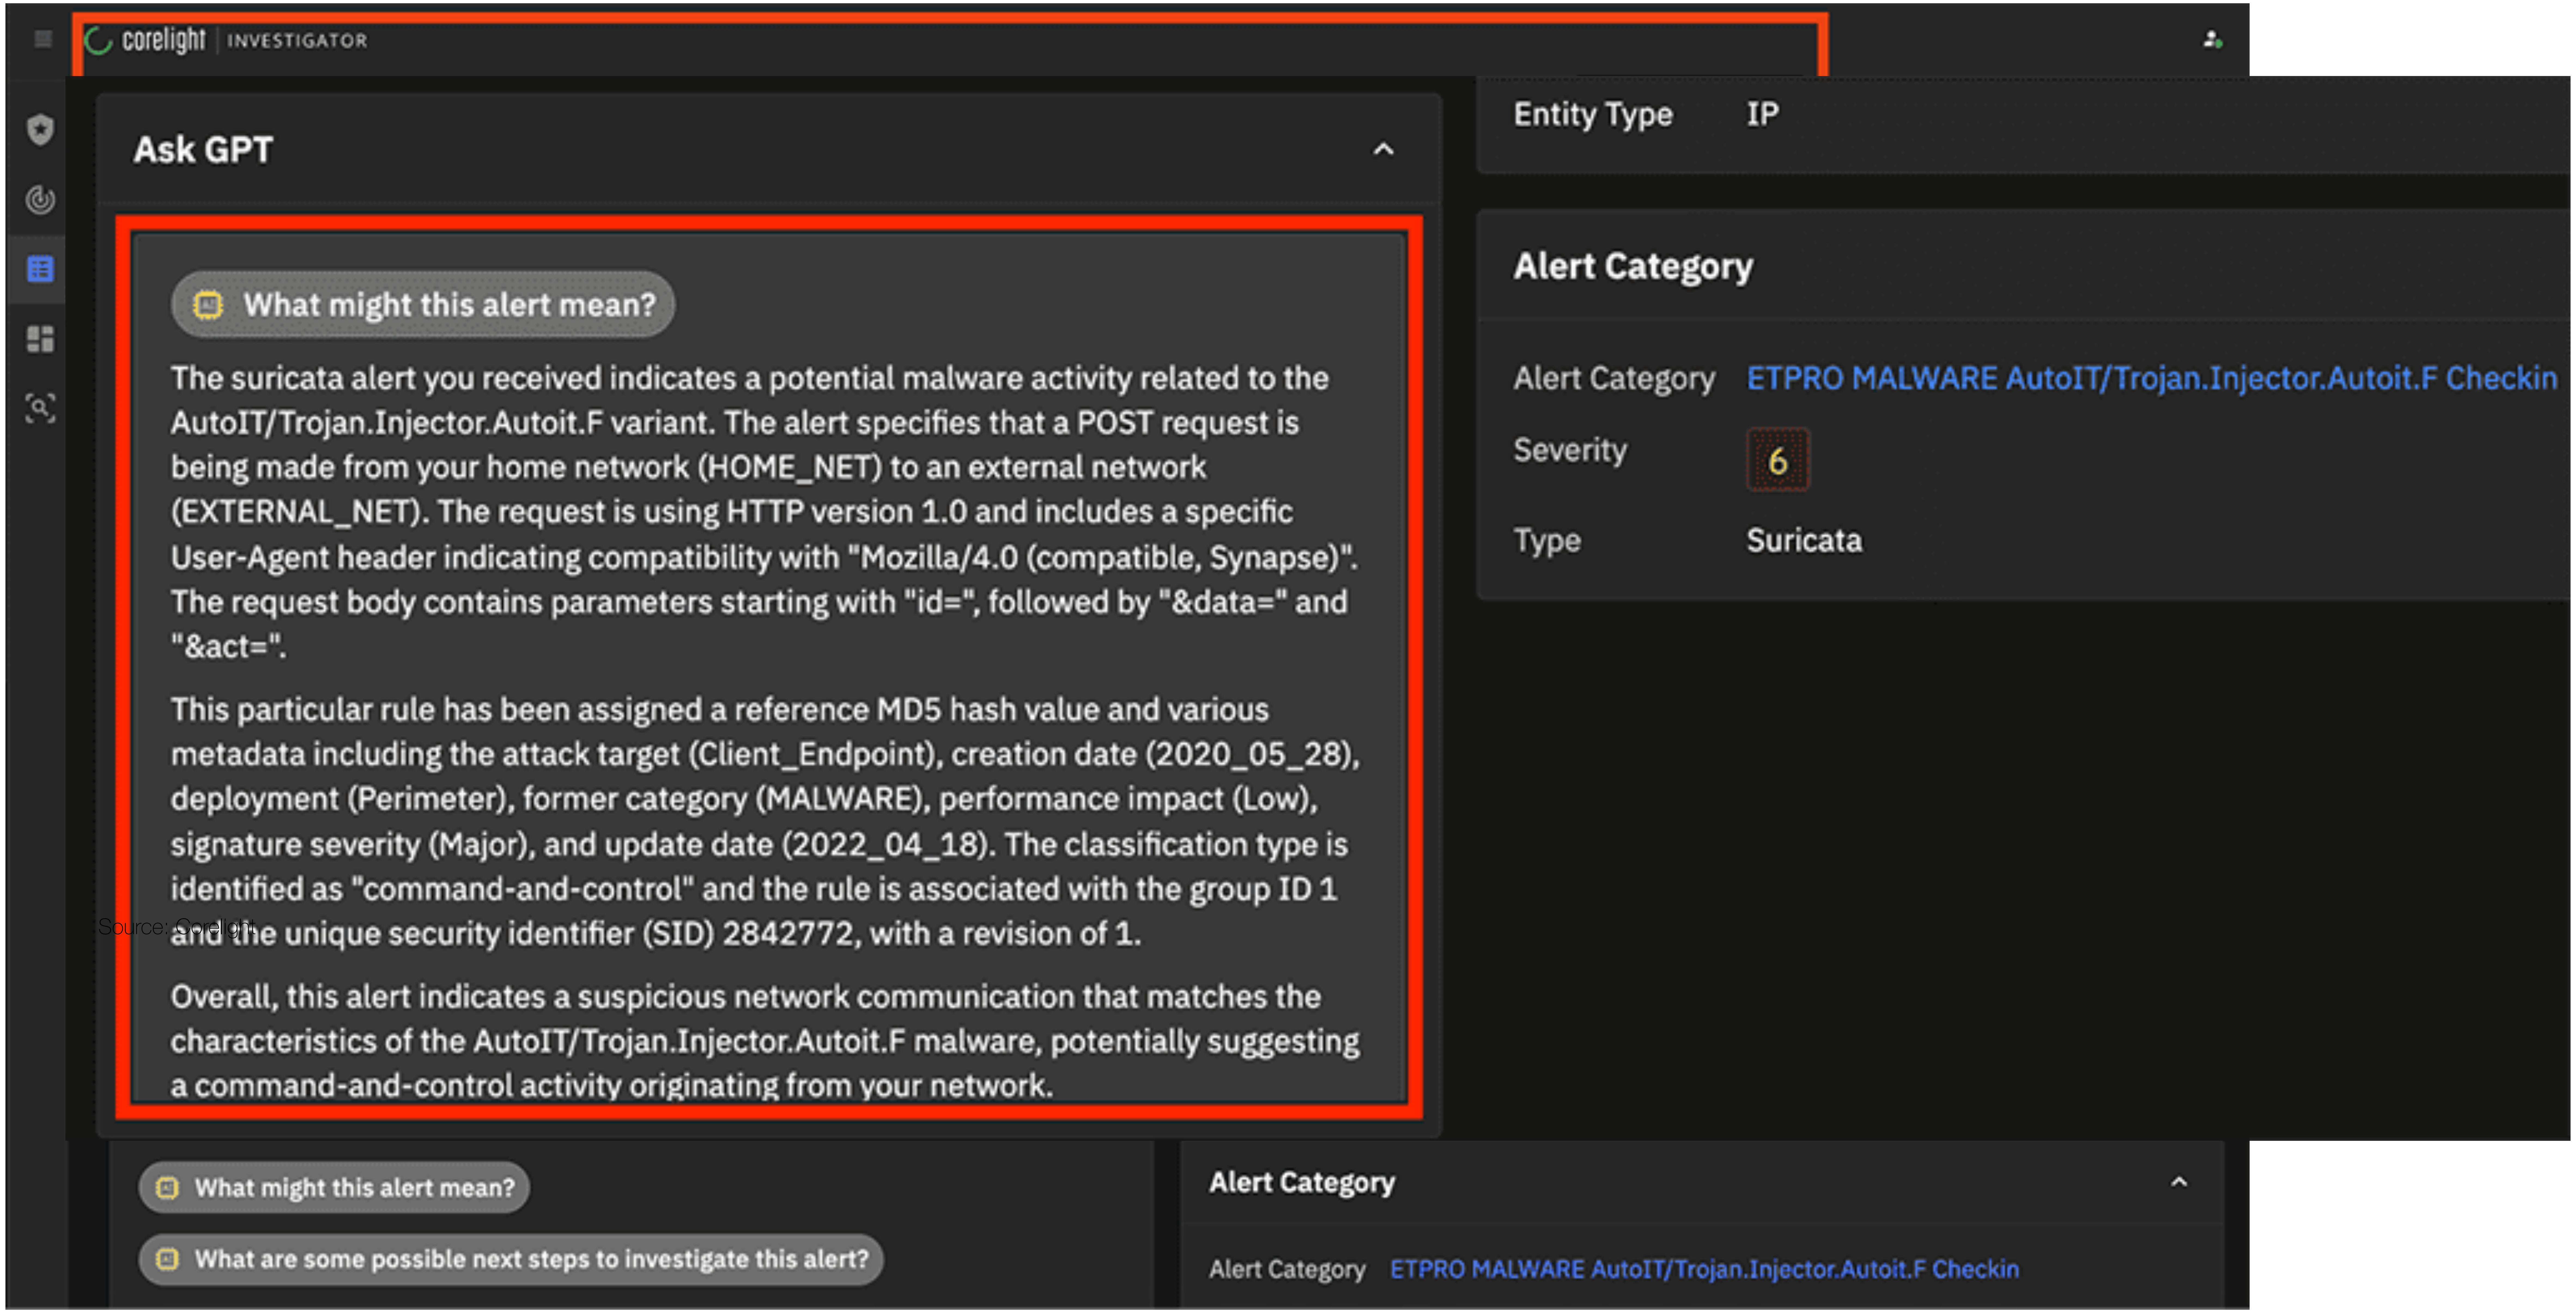

The suricata alert you received indicates a potential malware activity related to the AutoIT/Trojan.Injector.Autoit.F variant. The alert specifies that a POST request is being made from your home network (HOME_NET) to an external network (EXTERNAL_NET). The request is using HTTP version 1.0 and includes a specific User-Agent header indicating compatibility with "Mozilla/4.0 (compatible, Synapse)". The request body contains parameters starting with "id=", followed by "&data=" and "&act=".

This particular rule has been assigned a reference MD5 hash value and various metadata including the attack target (Client_Endpoint), creation date (2020_05_28), deployment (Perimeter), former category (MALWARE), performance impact (Low), signature severity (Major), and update date (2022_04_18). The classification type is identified as "command-and-control" and the rule is associated with the group ID 1 and the unique security identifier (SID) 2842772, with a revision of 1.

Overall, this alert indicates a suspicious network communication that matches the characteristics of the AutoIT/Trojan.Injector.Autoit.F malware, potentially suggesting a command-and-control activity originating from your network.

Entity Type IP

Alert Category

Alert Category ETPRO MALWARE AutoIT/Trojan.Injector.Autoit.F Checkin

Severity 6

Type Suricata

What might this alert mean?
What are some possible next steps to investigate this alert?
Alert Category

Alert Category ETPRO MALWARE AutoIT/Trojan.Injector.Autoit.F Checkin

Source: Corelight

Baby steps: Explaining Suricata rules

corelight | INVESTIGATOR
Ask GPT
Request More
Suppre

ID of 2842772.

What are some possible next steps to investigate this alert?

To investigate the Suricata alert "ETPRO MALWARE AutoIT/Trojan.Injector.Autoit.F Checkin," a Cyber Security Analyst can utilize various Zeek logs to gather additional information. The following Zeek logs are relevant for this investigation:

1. http.log: This log captures HTTP traffic and contains details such as client IP, server IP, timestamps, HTTP method, URI, host, user-agent, and more. It can provide information about the request made during the alert trigger.
2. files.log: This log records file-related activities, including file transfers. It can help identify any potentially malicious files involved in the alert.
3. conn.log: This log provides information about network connections, including source and destination IP addresses, ports, timestamps, connection duration, and flags. It can assist in understanding the network flow associated with the alert.

When investigating the Suricata alert using Zeek logs, the following steps can be followed:

1. Identify the alert in Suricata: Note the time, source IP address, destination IP address, and any other relevant details from the alert.
2. Search the conn.log: Look for connections involving the source IP address and destination IP address

What might this alert mean?

What are some possible next steps to investigate this alert?

Entity Type IP

Entity

Entity	10.2[redacted]
Entity Type	IP

Alert Category

Alert Category	ETPRO MALWARE AutoIT/Trojan.Injector.Autoit.F Checkin
Severity	6
Type	Suricata

Alert Category

Alert Category ETPRO MALWARE AutoIT/Trojan.Injector.Autoit.F Checkin

Source: Corelight

No surprise: Many such AI solutions emerging

No surprise: Many such AI solutions emerging

Google

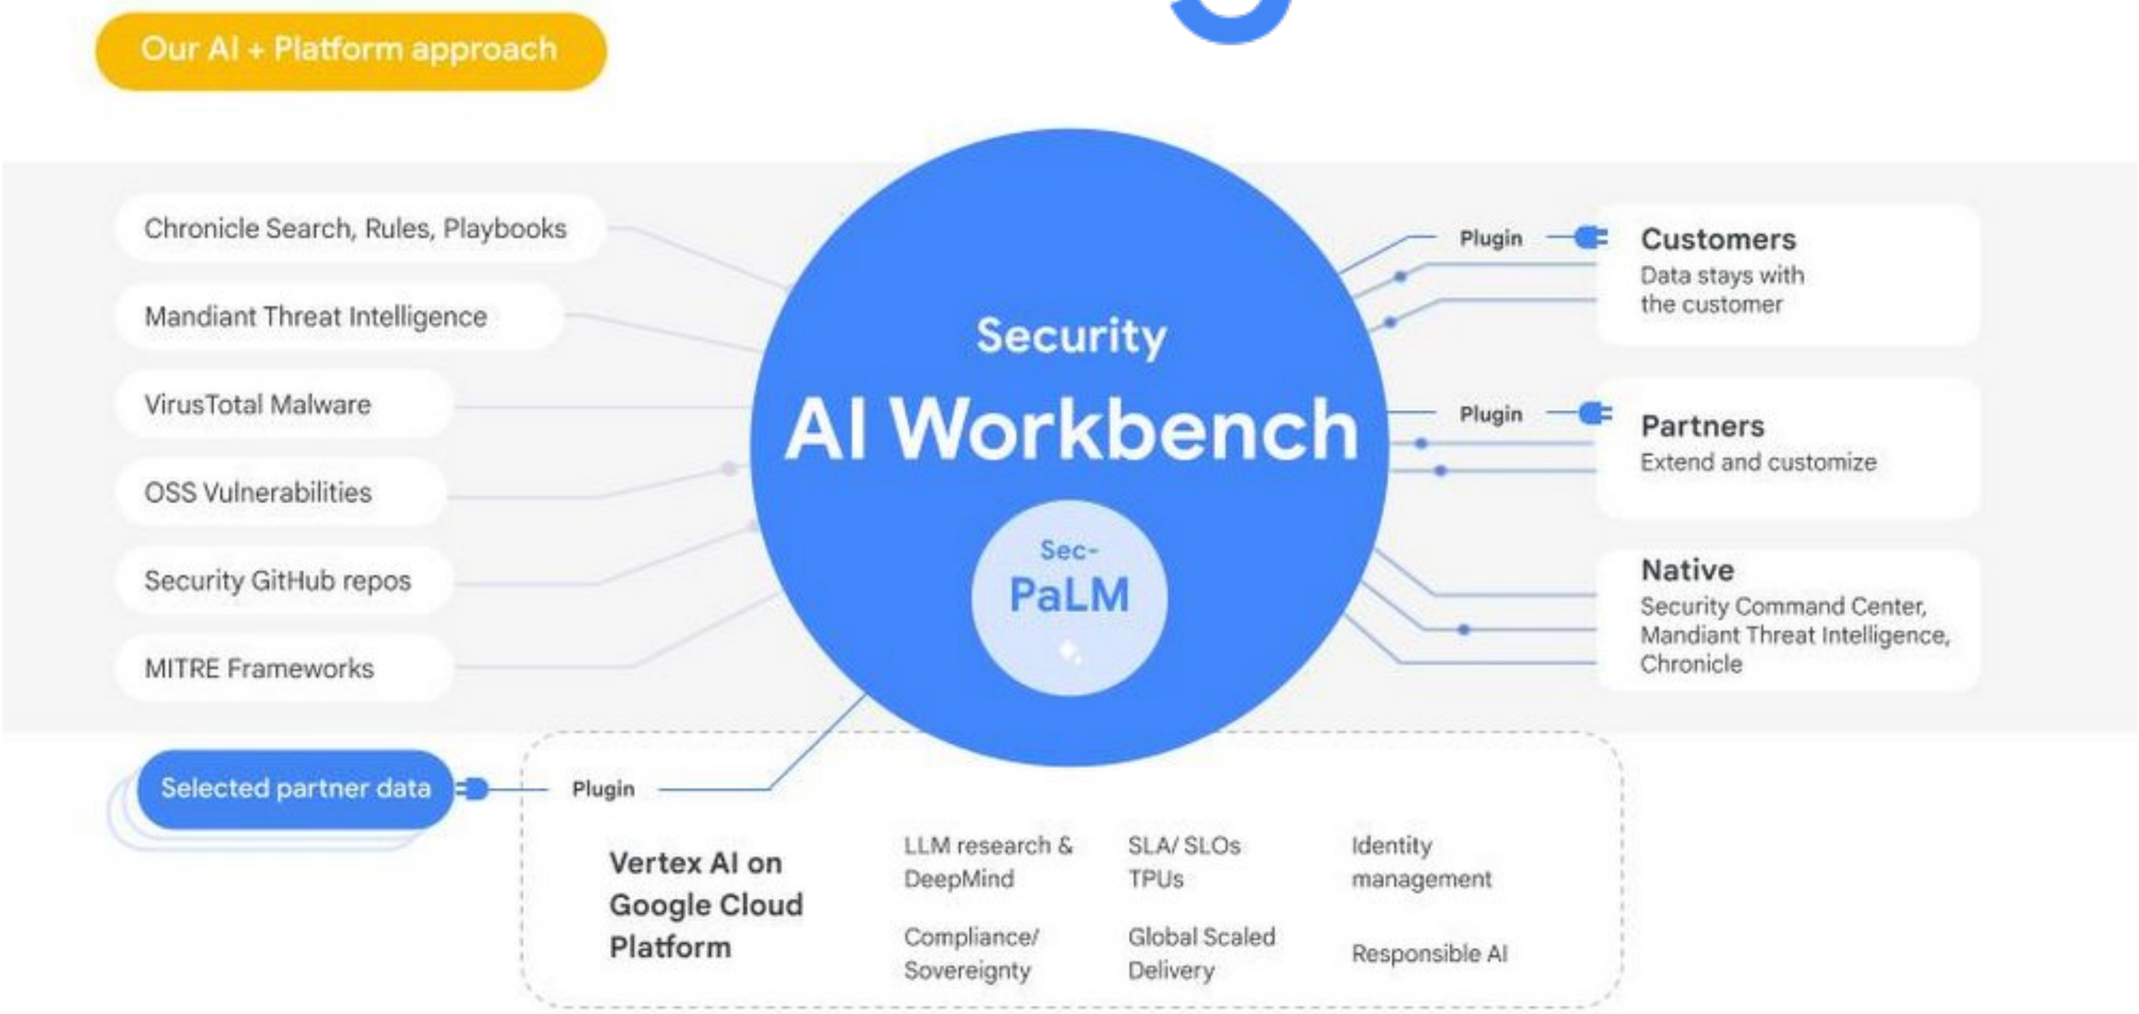
Our AI + Platform approach

Chronicle Search, Rules, Playbooks
Mandiant Threat Intelligence
VirusTotal Malware
OSS Vulnerabilities
Security GitHub repos
MITRE Frameworks

Security AI Workbench
Sec-PaLM

Plugin
Customers
Data stays with the customer

Plugin
Partners
Extend and customize

Native
Security Command Center, Mandiant Threat Intelligence, Chronicle

Selected partner data
Plugin

Vertex AI on Google Cloud Platform
LLM research & DeepMind
SLA/ SLOs TPUs
Identity management
Compliance/ Sovereignty
Global Scaled Delivery
Responsible AI

Charlotte

CROWDSTRIKE

Crowdstrike’s Charlotte

Crowdstrike's Charlotte

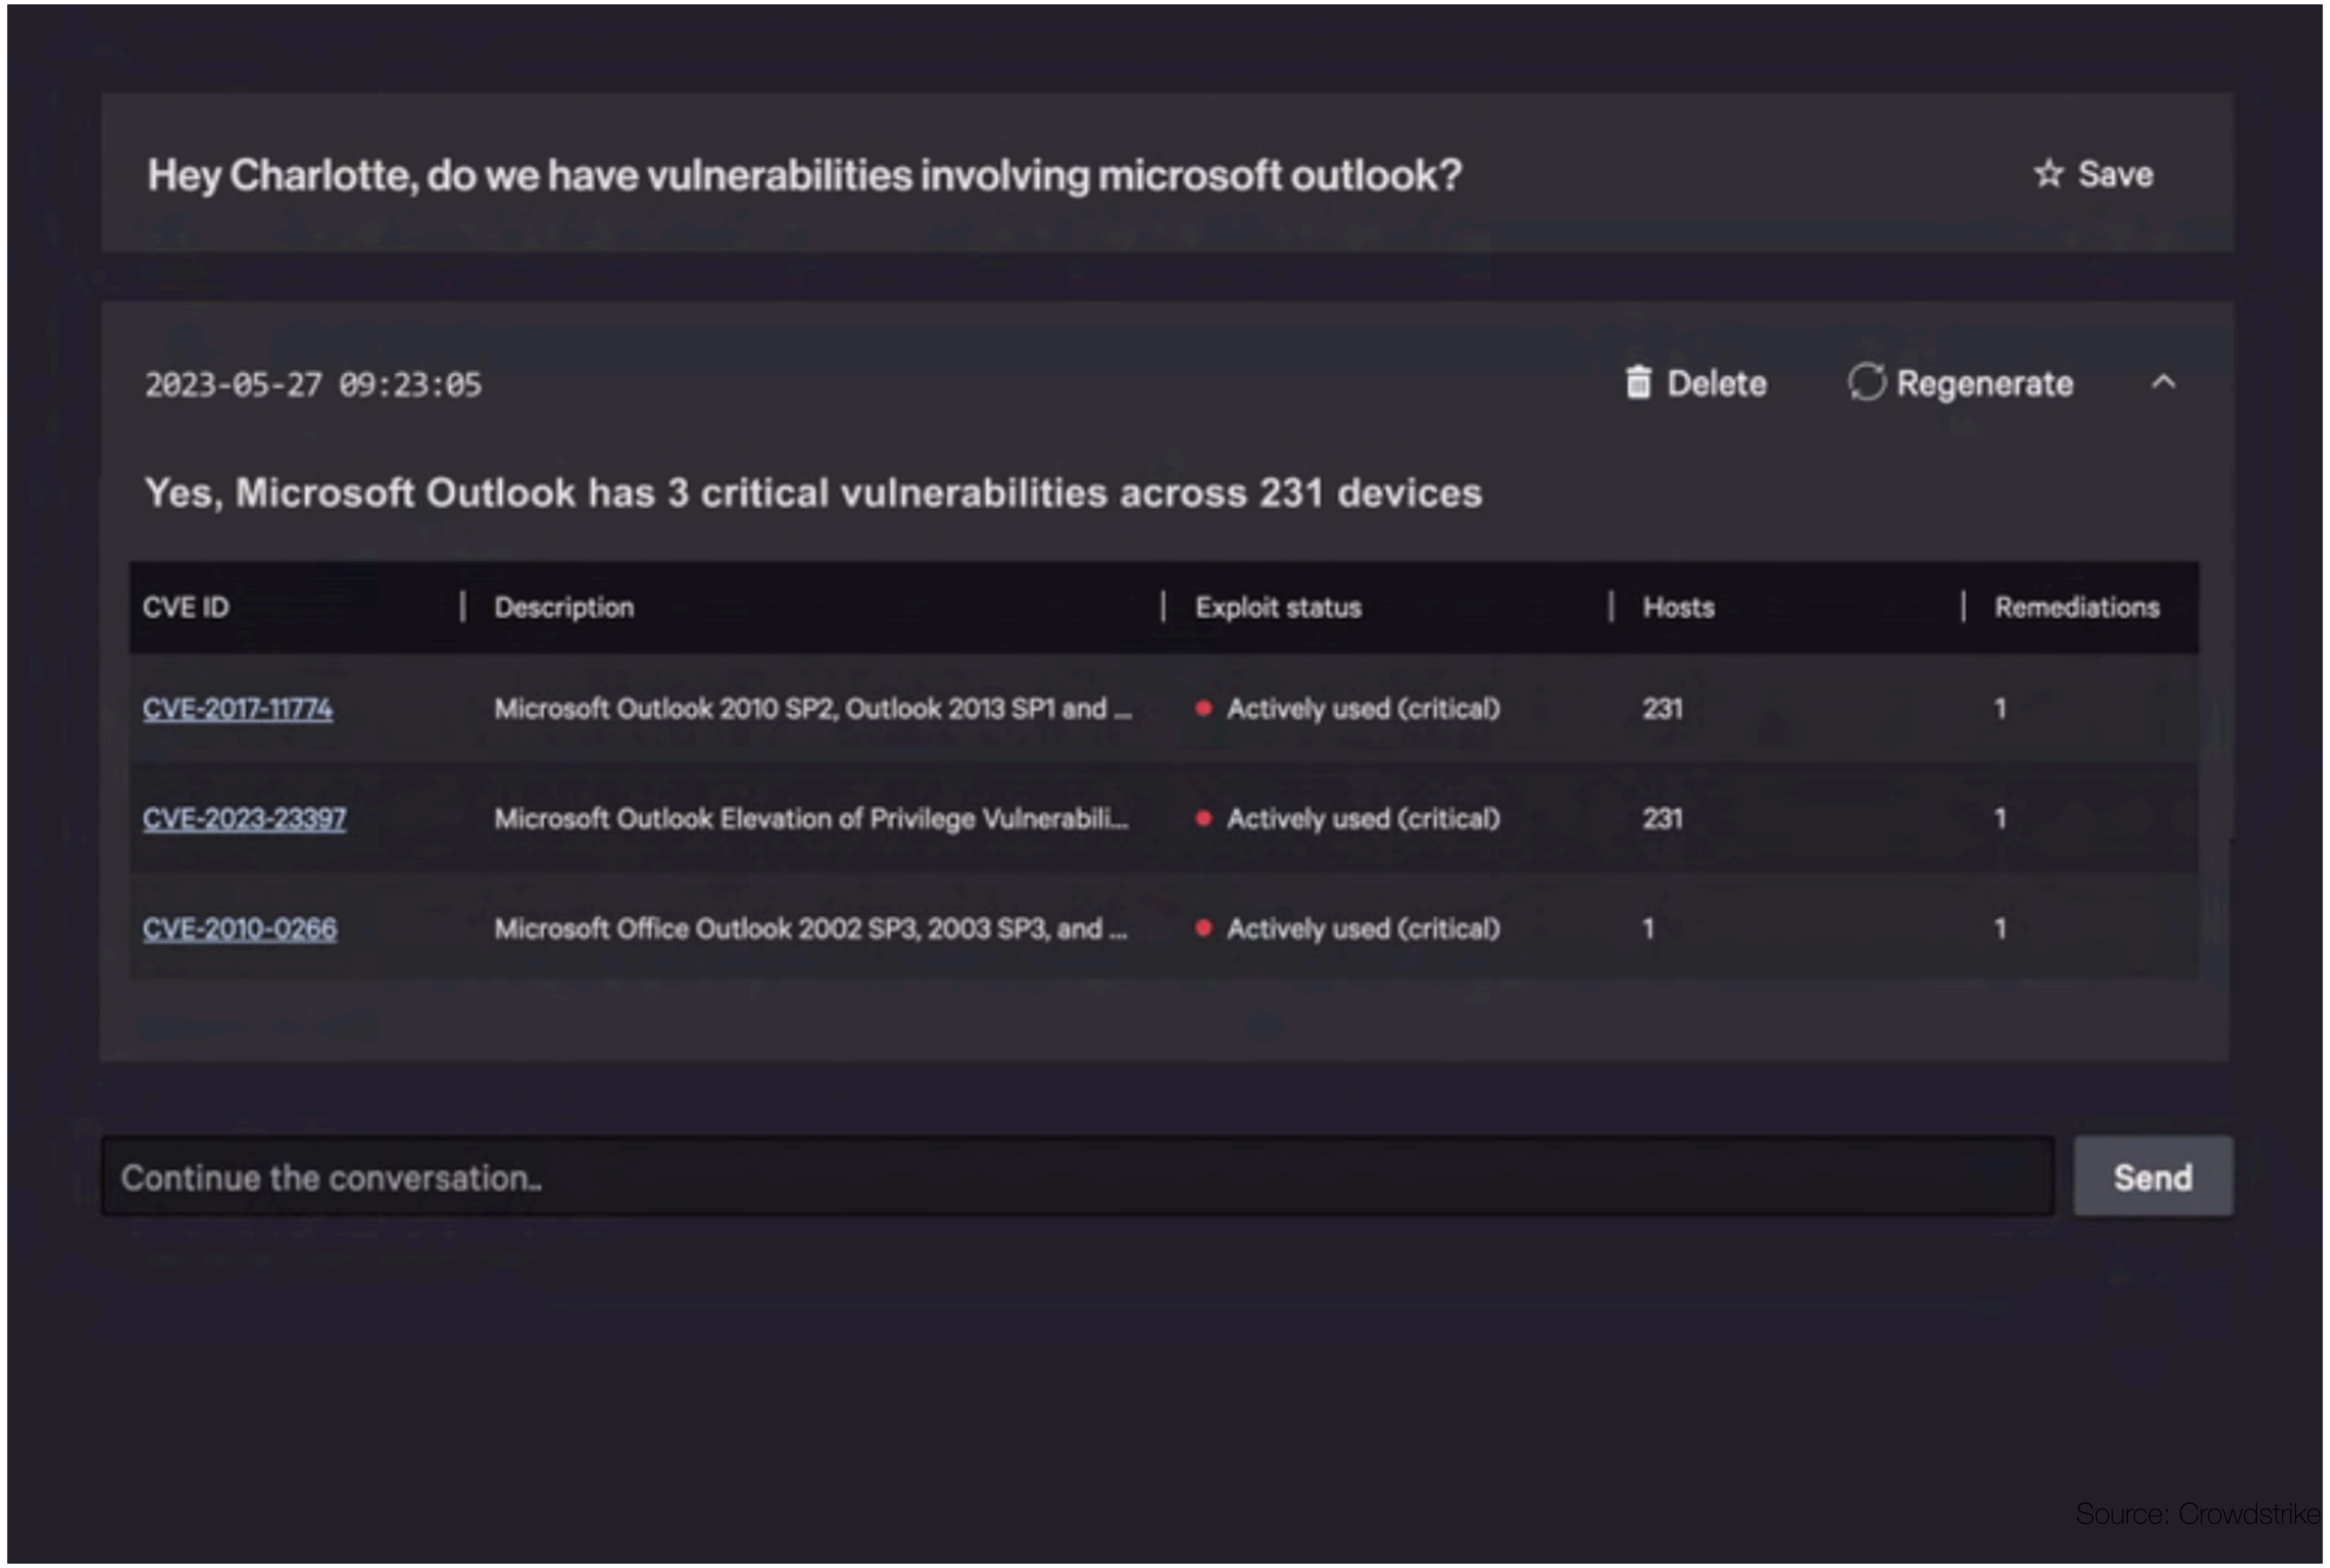

Hey Charlotte, do we have vulnerabilities involving microsoft outlook?

2023-05-27 09:23:05

Yes, Microsoft Outlook has 3 critical vulnerabilities across 231 devices

CVE ID	Description	Exploit status	Hosts	Remediations
CVE-2017-11774	Microsoft Outlook 2010 SP2, Outlook 2013 SP1 and ...	● Actively used (critical)	231	1
CVE-2023-23397	Microsoft Outlook Elevation of Privilege Vulnerabili...	● Actively used (critical)	231	1
CVE-2010-0266	Microsoft Office Outlook 2002 SP3, 2003 SP3, and ...	● Actively used (critical)	1	1

Source: Crowdstrike

Crowdstrike's Charlotte

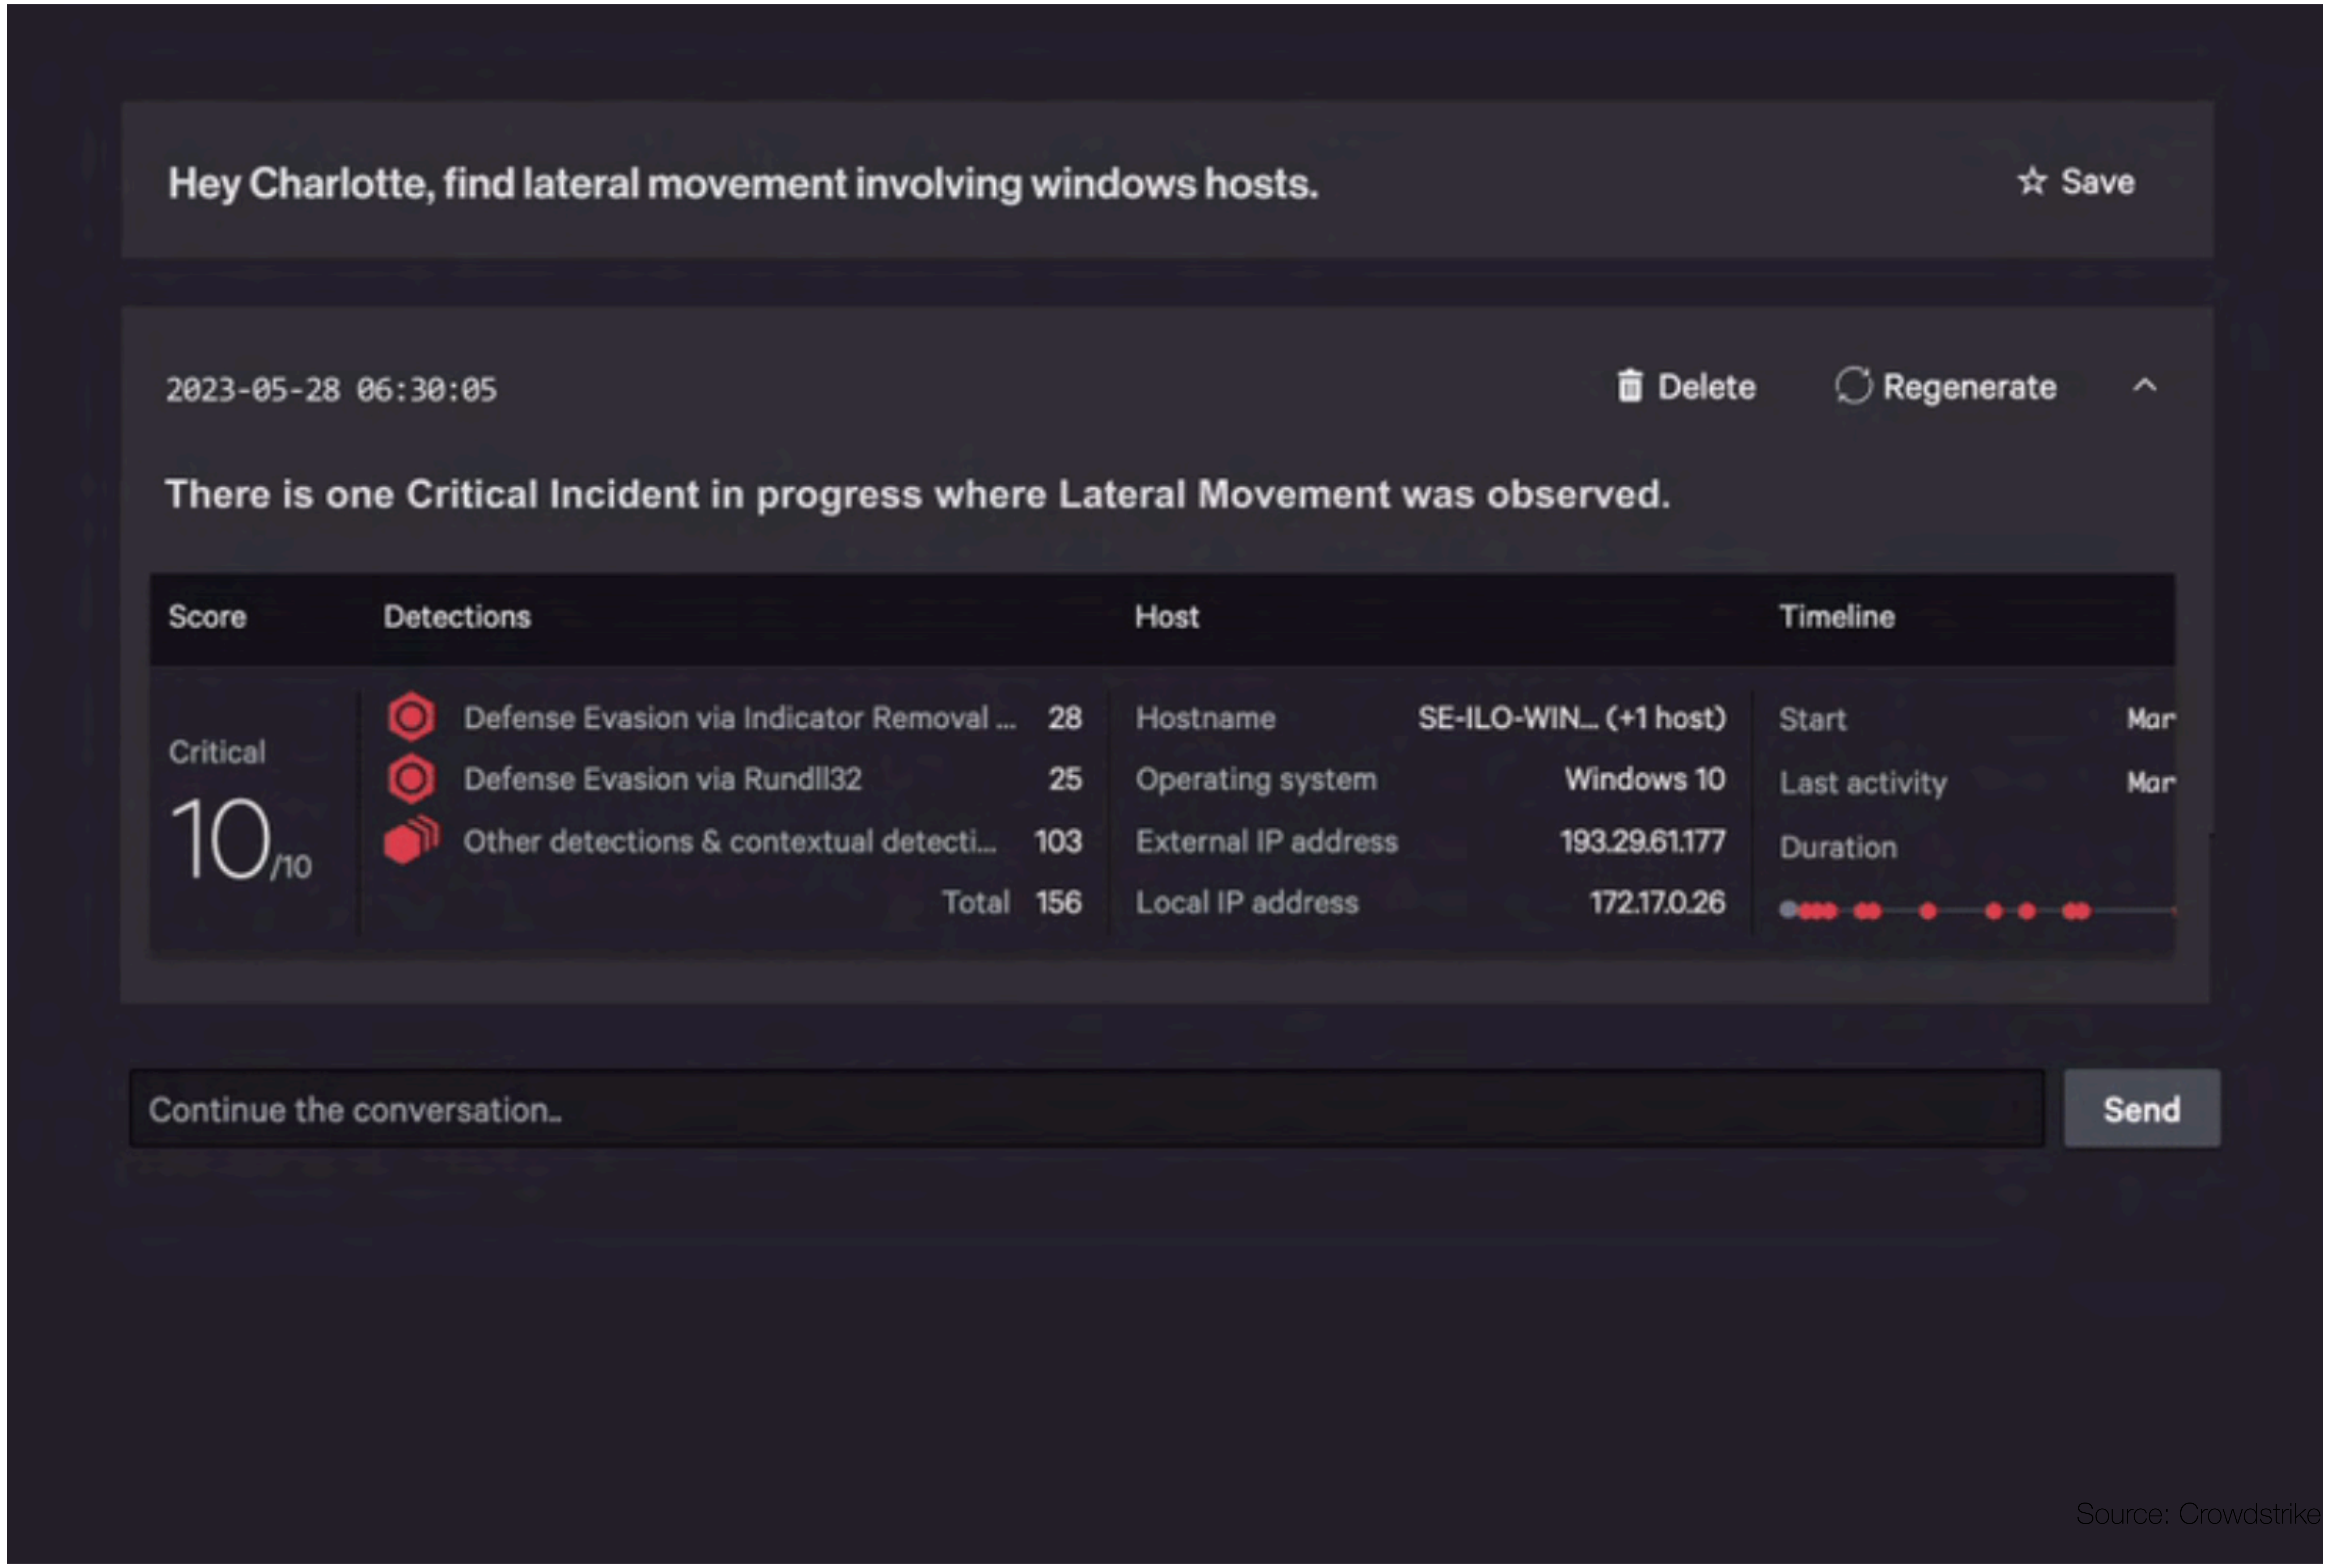

Hey Charlotte, find lateral movement involving windows hosts.

2023-05-28 06:30:05

There is one Critical Incident in progress where Lateral Movement was observed.

Score	Detections		Host		Timeline	
Critical 10/10	Defense Evasion via Indicator Removal ...	28	Hostname	SE-ILO-WIN... (+1 host)	Start	Mar
	Defense Evasion via Rundll32	25	Operating system	Windows 10	Last activity	Mar
	Other detections & contextual detecti...	103	External IP address	193.29.61.177	Duration	
	Total	156	Local IP address	172.17.0.26		

Source: Crowdstrike

Crowdstrike's Charlotte

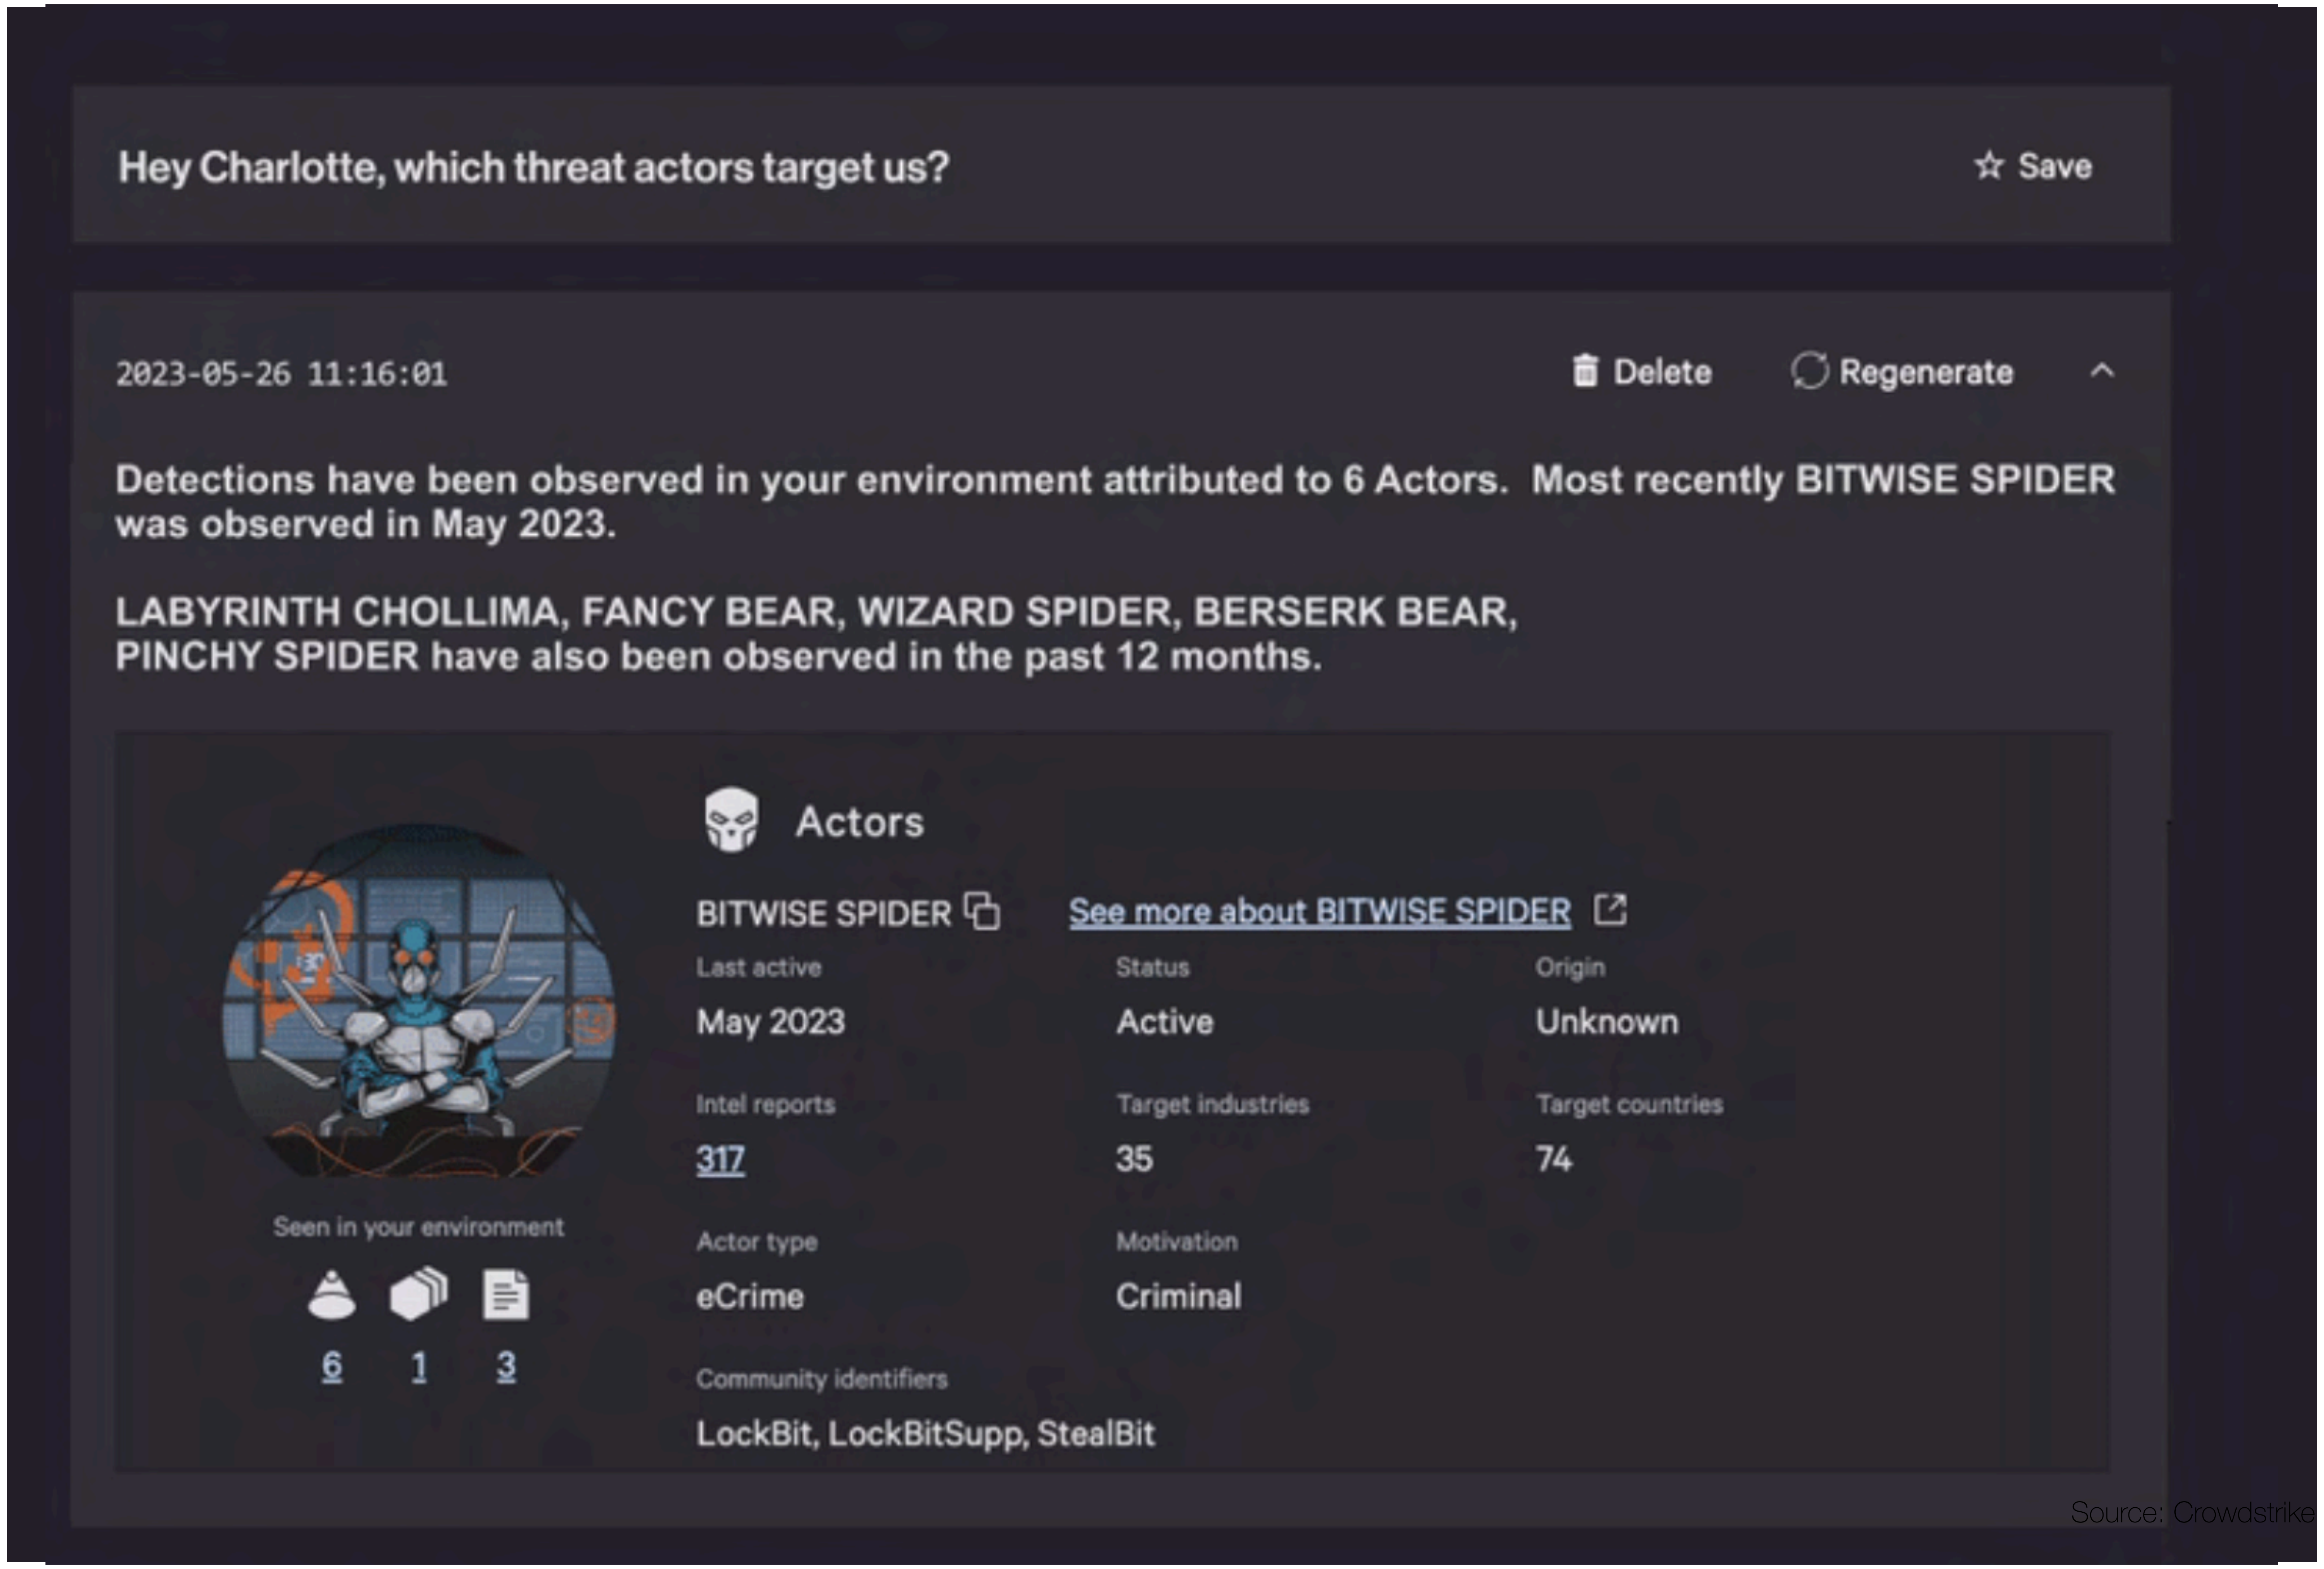

Hey Charlotte, which threat actors target us?

2023-05-26 11:16:01

Detections have been observed in your environment attributed to 6 Actors. Most recently BITWISE SPIDER was observed in May 2023.

LABYRINTH CHOLLIMA, FANCY BEAR, WIZARD SPIDER, BERSERK BEAR, PINCHY SPIDER have also been observed in the past 12 months.

BITWISE SPIDER

Last active	Status	Origin
May 2023	Active	Unknown
Intel reports	**Target industries**	**Target countries**
317	35	74
Actor type	**Motivation**	
eCrime	Criminal	

Community identifiers
LockBit, LockBitSupp, StealBit

Seen in your environment
6 1 3

Source: Crowdstrike

Corelight Threat Hunting Guide

Corelight Threat Hunting Guide

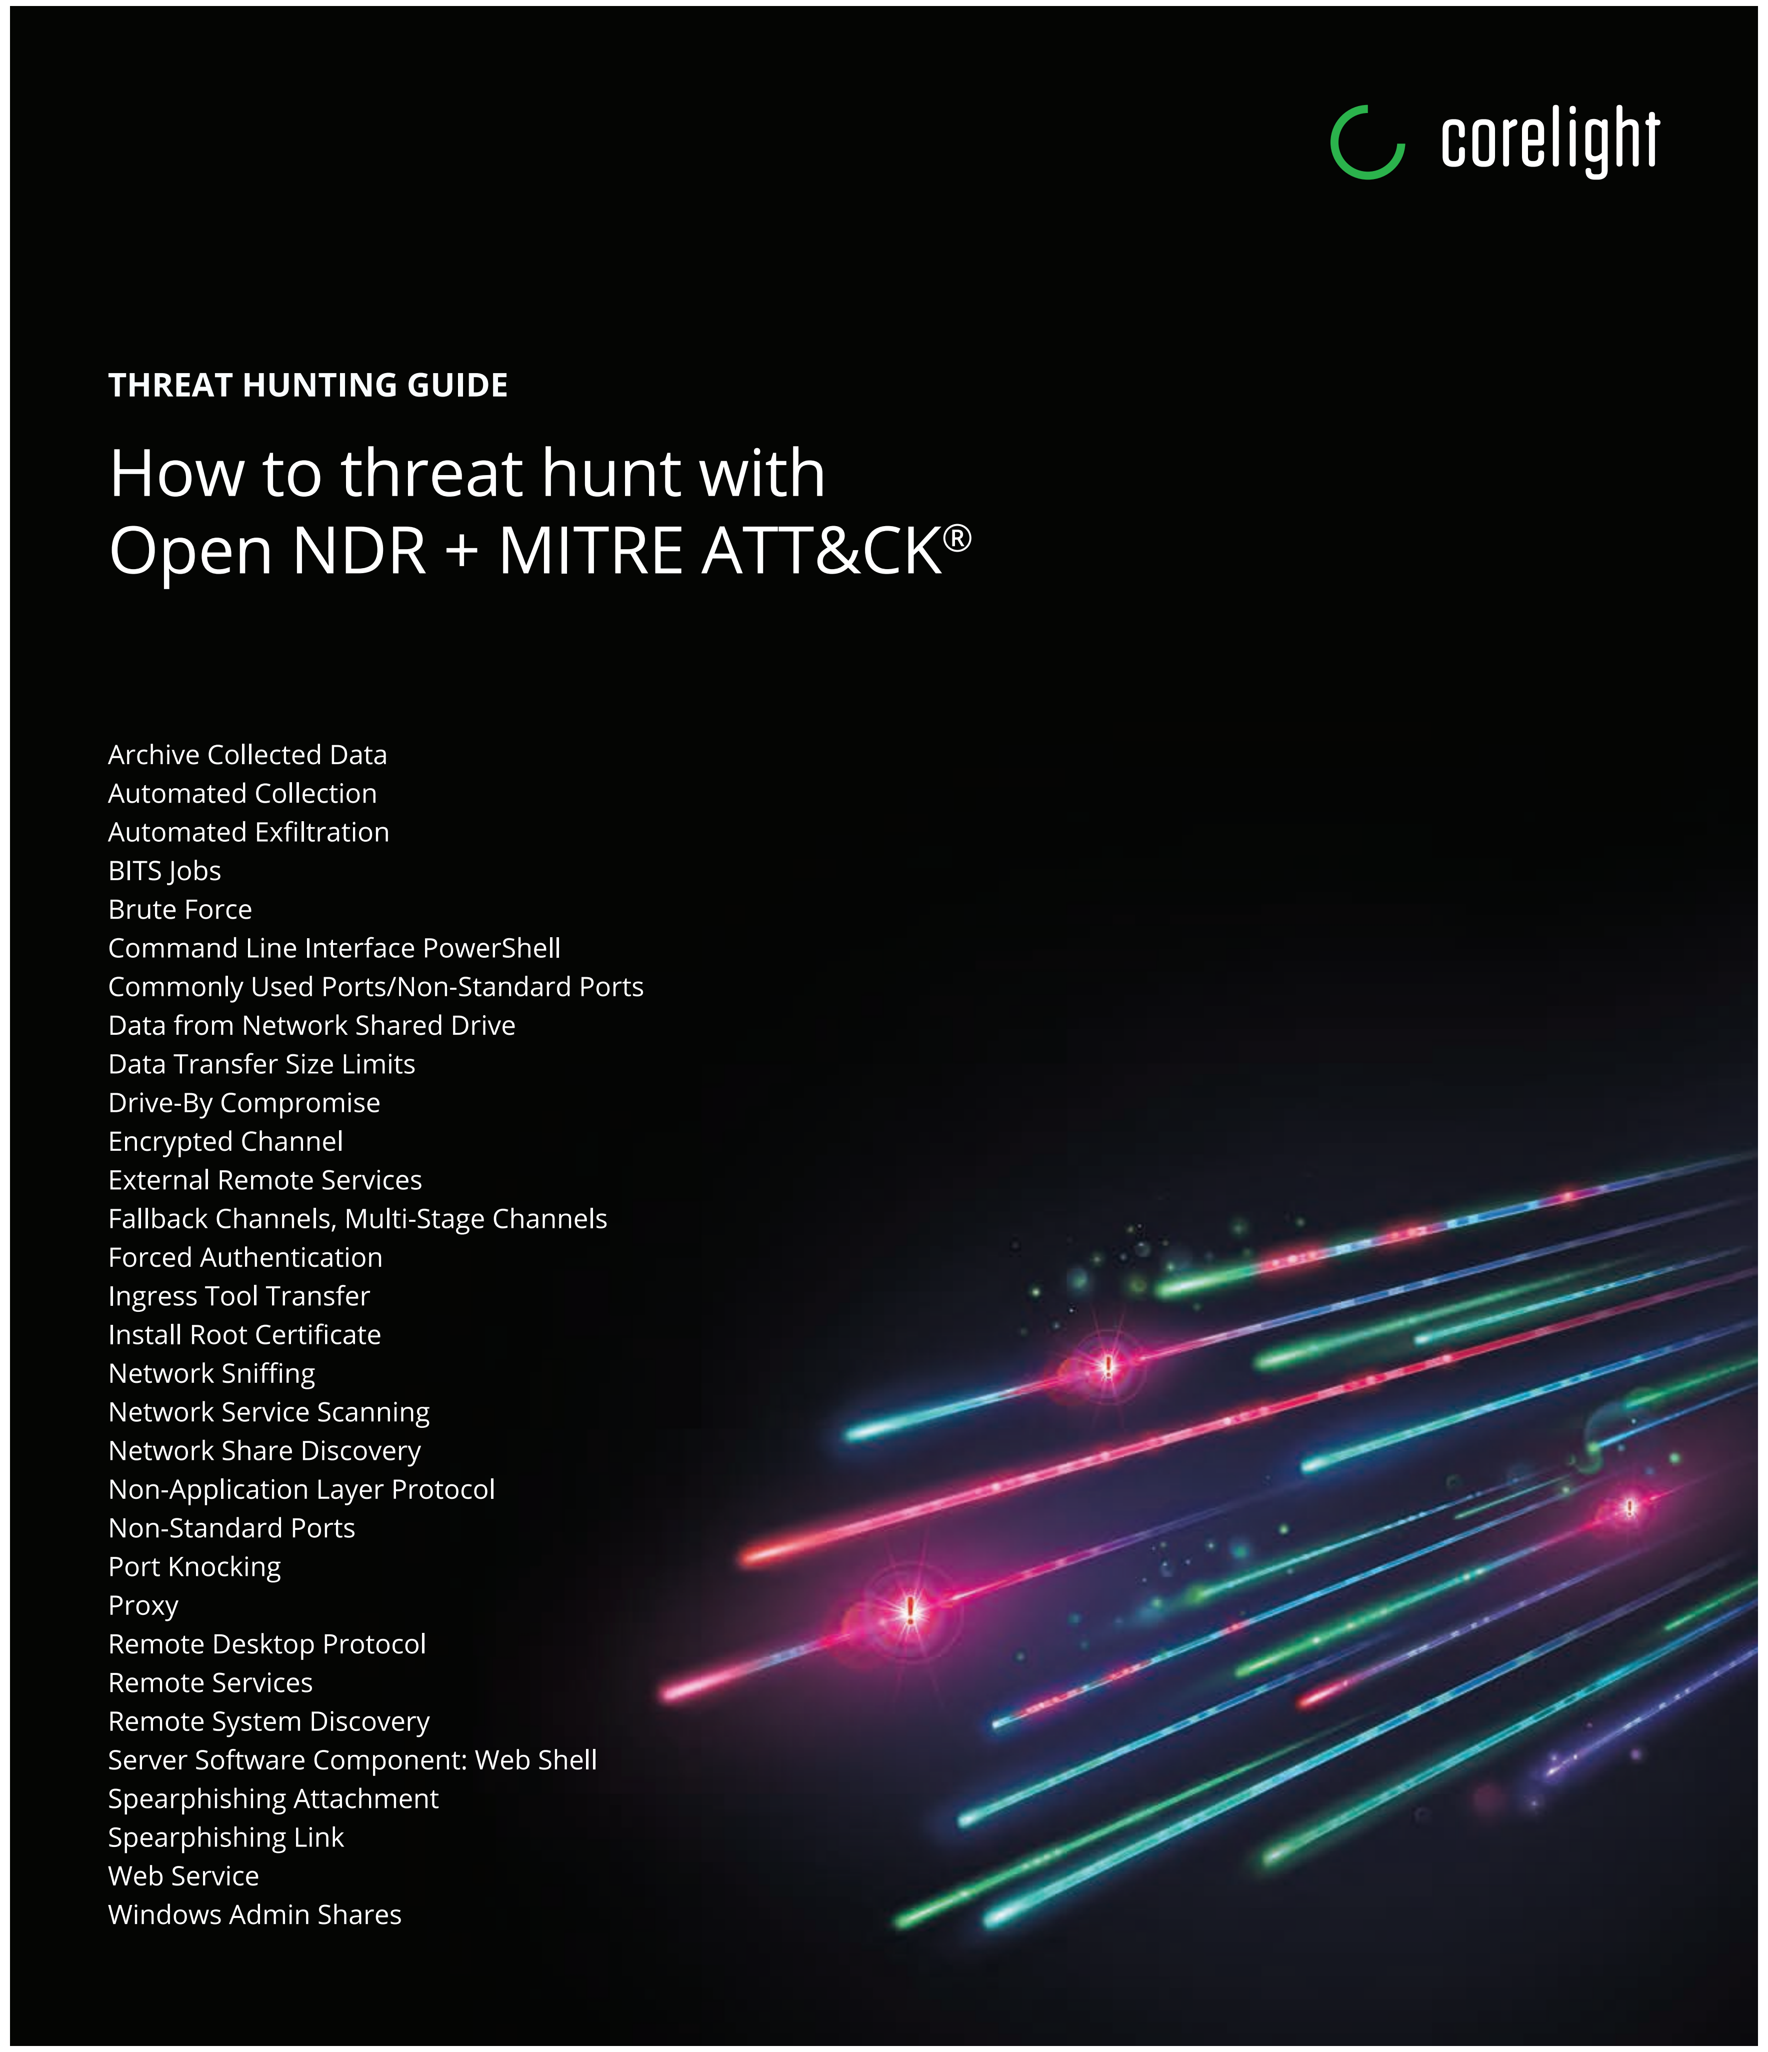

corelight

THREAT HUNTING GUIDE

How to threat hunt with
Open NDR + MITRE ATT&CK®

- Archive Collected Data
- Automated Collection
- Automated Exfiltration
- BITS Jobs
- Brute Force
- Command Line Interface PowerShell
- Commonly Used Ports/Non-Standard Ports
- Data from Network Shared Drive
- Data Transfer Size Limits
- Drive-By Compromise
- Encrypted Channel
- External Remote Services
- Fallback Channels, Multi-Stage Channels
- Forced Authentication
- Ingress Tool Transfer
- Install Root Certificate
- Network Sniffing
- Network Service Scanning
- Network Share Discovery
- Non-Application Layer Protocol
- Non-Standard Ports
- Port Knocking
- Proxy
- Remote Desktop Protocol
- Remote Services
- Remote System Discovery
- Server Software Component: Web Shell
- Spearphishing Attachment
- Spearphishing Link
- Web Service
- Windows Admin Shares

Source: Corelight

Corelight Threat Hunting Guide

THREAT HUNTING GUIDE

How to threat hunt with
Open NDR + MITRE ATT&CK®

Archive Collected Data
Automated Collection
Automated Exfiltration
BITS Jobs
Brute Force
Command Line Interface PowerShell
Commonly Used Ports/Non-Standard Ports
Data from Network Shared Drive
Data Transfer Size Limits
Drive-By Compromise
Encrypted Channel
External Remote Services
Fallback Channels, Multi-Stage Channels
Forced Authentication
Ingress Tool Transfer
Install Root Certificate
Network Sniffing
Network Service Scanning
Network Share Discovery
Non-Application Layer Protocol
Non-Standard Ports
Port Knocking
Proxy
Remote Desktop Protocol
Remote Services
Remote System Discovery
Server Software Component: Web Shell
Spearphishing Attachment
Spearphishing Link
Web Service
Windows Admin Shares

Source: Corelight

EXFILTRATION

Automated Exfiltration
If an attacker is using an automated means of exfiltration, data artifacts are captured in the Corelight data.

To look for exfiltration in your network, you can use the Zeek package developed to calculate Producer/Consumer Ratio (PCR). PCR values indicate whether flows are consumptive (download) versus productive (upload). PCR values range from -1 (consumptive) to +1 (productive). To hunt for exfiltration using this package:

1. Install and enable the PCR package.
2. Generate a table of id.orig_h, id.resp_h, id.resp_p, and pcr from the conn log.
3. Use local_orig is false or local_resp is true to filter the results.
4. Reduce the results by filtering where pcr <= 0.
5. For each host generating flows where pcr >= 0, consider whether that host is expected to transmit data, inside or outside the network.

Another option is to use a SIEM to calculate the PCR using the information available in the Corelight conn log. The following query creates a table organized by host that contains the originating and responding bytes and a PCR value.

```
index=corelight sourcetype=corelight_conn | stats sum(orig_bytes) as Total_orig_bytes, sum(resp_bytes) as Total_resp_bytes by id.orig_h id.resp_h | eval PCR=(Total_orig_bytes-Total_resp_bytes)/(Total_orig_bytes+Total_resp_bytes) | fields id.orig_h id.resp_h Total_orig_bytes Total_resp_bytes PCR
```

Data Transfer Size Limits
An attacker may attempt to transfer data or files by "chunking" them into smaller pieces, to avoid hard-coded data transfer limits or thresholds. We will present two methods to hunt for this technique.

The first method analyzes data leaving the network based on source and destination pairs and requires a data aggregation/visualization platform (unless you enjoy AWKing and GREPing through data):

1. Generate a table from the conn log including the id.orig_h, id.resp_h, id.resp_p, and sum(orig_bytes).
2. Sort the results by the largest sum (orig_bytes).
3. Examine each host and determine if there is a legitimate reason for uploads to that destination.

The second method analyzes the frequency, and sizes, of outbound transfers from each source:

1. Generate a table from the conn log including id.orig_h, id.resp_h, id.resp_p, and count(orig_bytes).
2. Sort the results by the largest count(orig_bytes).
3. Examine the results and determine the reason for all the connections with the same amount of data flowing from the source to the destination.

Corelight Threat Hunting Guide

corelight

THREAT HUNTING GUIDE

How to threat hunt with
Open NDR + MITRE ATT&CK®

Archive Collected Data
Automated Collection
Automated Exfiltration
BITS Jobs
Brute Force
Command Line Interface PowerShell
Commonly Used Ports/Non-Standard Ports
Data from Network Shared Drive
Data Transfer Size Limits
Drive-By Compromise
Encrypted Channel
External Remote Services
Fallback Channels, Multi-Stage Channels
Forced Authentication
Ingress Tool Transfer
Install Root Certificate
Network Sniffing
Network Service Scanning
Network Share Discovery
Non-Application Layer Protocol
Non-Standard Ports
Port Knocking
Proxy
Remote Desktop Protocol
Remote Services
Remote System Discovery
Server Software Component: Web Shell
Spearphishing Attachment
Spearphishing Link
Web Service
Windows Admin Shares

Source: Corelight

EXFILTRATION

Automated Exfiltration
If an attacker is using an automated means of exfiltration, data artifacts are captured in the Corelight data.

To look for exfiltration in your network, you can use the Zeek package developed to calculate Producer/Consumer Ratio (PCR). PCR values indicate whether flows are consumptive (download) versus productive (upload). PCR values range from -1 (consumptive) to +1 (productive). To hunt for exfiltration using this package:

1. Install and enable the PCR package.
2. Generate a table of id.orig_h, id.resp_h, id.resp_p, and pcr from the conn log.
3. Use local_orig is false or local_resp is true to filter the results.
4. Reduce the results by filtering where pcr <= 0.
5. For each host generating flows where pcr >= 0, consider whether that host is expected to transmit data, inside or outside the network.

Another option is to use a SIEM to calculate the PCR using the information available in the Corelight conn log. The following query creates a table organized by host that contains the originating and responding bytes and a PCR value.

```
index=corelight sourcetype=corelight_conn | stats sum(orig_bytes) as Total_orig_bytes, sum(resp_bytes) as Total_resp_
bytes by id.orig_h id.resp_h | eval PCR=(Total_orig_bytes-Total_resp_bytes)/(Total_orig_bytes+Total_resp_bytes) | fields
id.orig_h id.resp_h Total_orig_bytes Total_resp_bytes PCR
```

Data Transfer Size Limits
An attacker may attempt to transfer data or files by "chunking" them into smaller pieces, to avoid hard-coded data transfer limits or thresholds. We will present two methods to hunt for this technique.

The first method analyzes data leaving the network based on source and destination pairs and requires a data aggregation/visualization platform (unless you enjoy AWKing and GREPing through data):

1. Generate a table from the conn log including the id.orig_h, id.resp_h, id.resp_p, and sum(orig_bytes).
2. Sort the results by the largest sum (orig_bytes).
3. Examine each host and determine if there is a legitimate reason for uploads to that destination.

The second method analyzes the frequency, and sizes, of outbound transfers from each source:

1. Generate a table from the conn log including id.orig_h, id.resp_h, id.resp_p, and count(orig_bytes).
2. Sort the results by the largest count(orig_bytes).
3. Examine the results and determine the reason for all the connections with the same amount of data flowing from the source to the destination.

There are lots of workflows here that could be largely automated.

Threat hunting maturity model

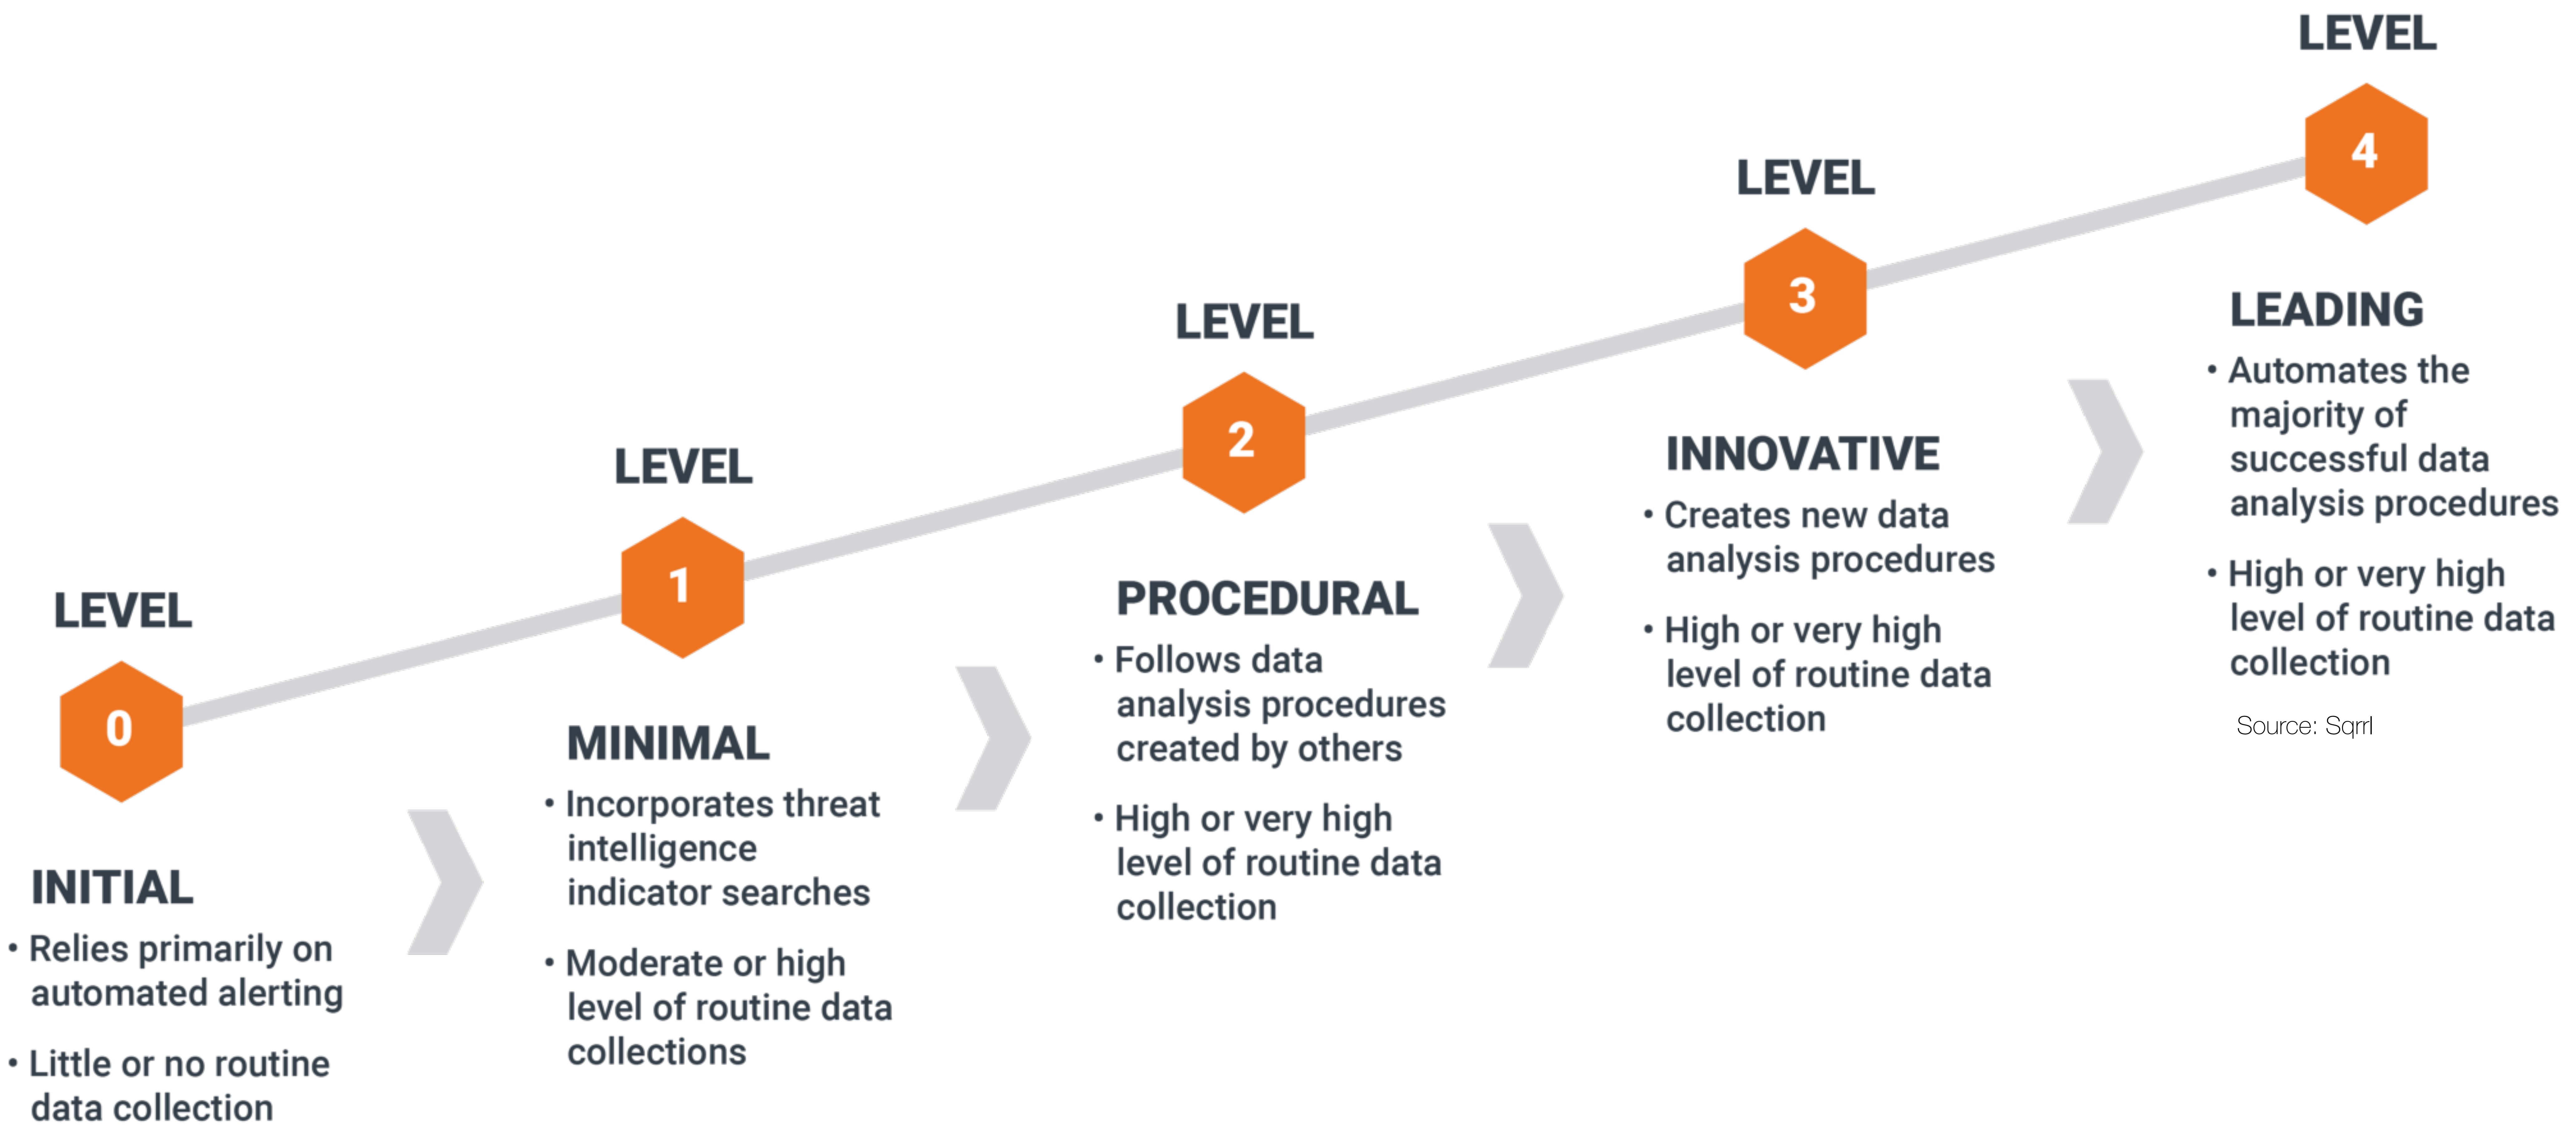

LEVEL

0

INITIAL

- Relies primarily on automated alerting
- Little or no routine data collection

LEVEL

1

MINIMAL

- Incorporates threat intelligence indicator searches
- Moderate or high level of routine data collections

LEVEL

2

PROCEDURAL

- Follows data analysis procedures created by others
- High or very high level of routine data collection

LEVEL

3

INNOVATIVE

- Creates new data analysis procedures
- High or very high level of routine data collection

LEVEL

4

LEADING

- Automates the majority of successful data analysis procedures
- High or very high level of routine data collection

Source: Sqrrl

Threat hunting maturity model

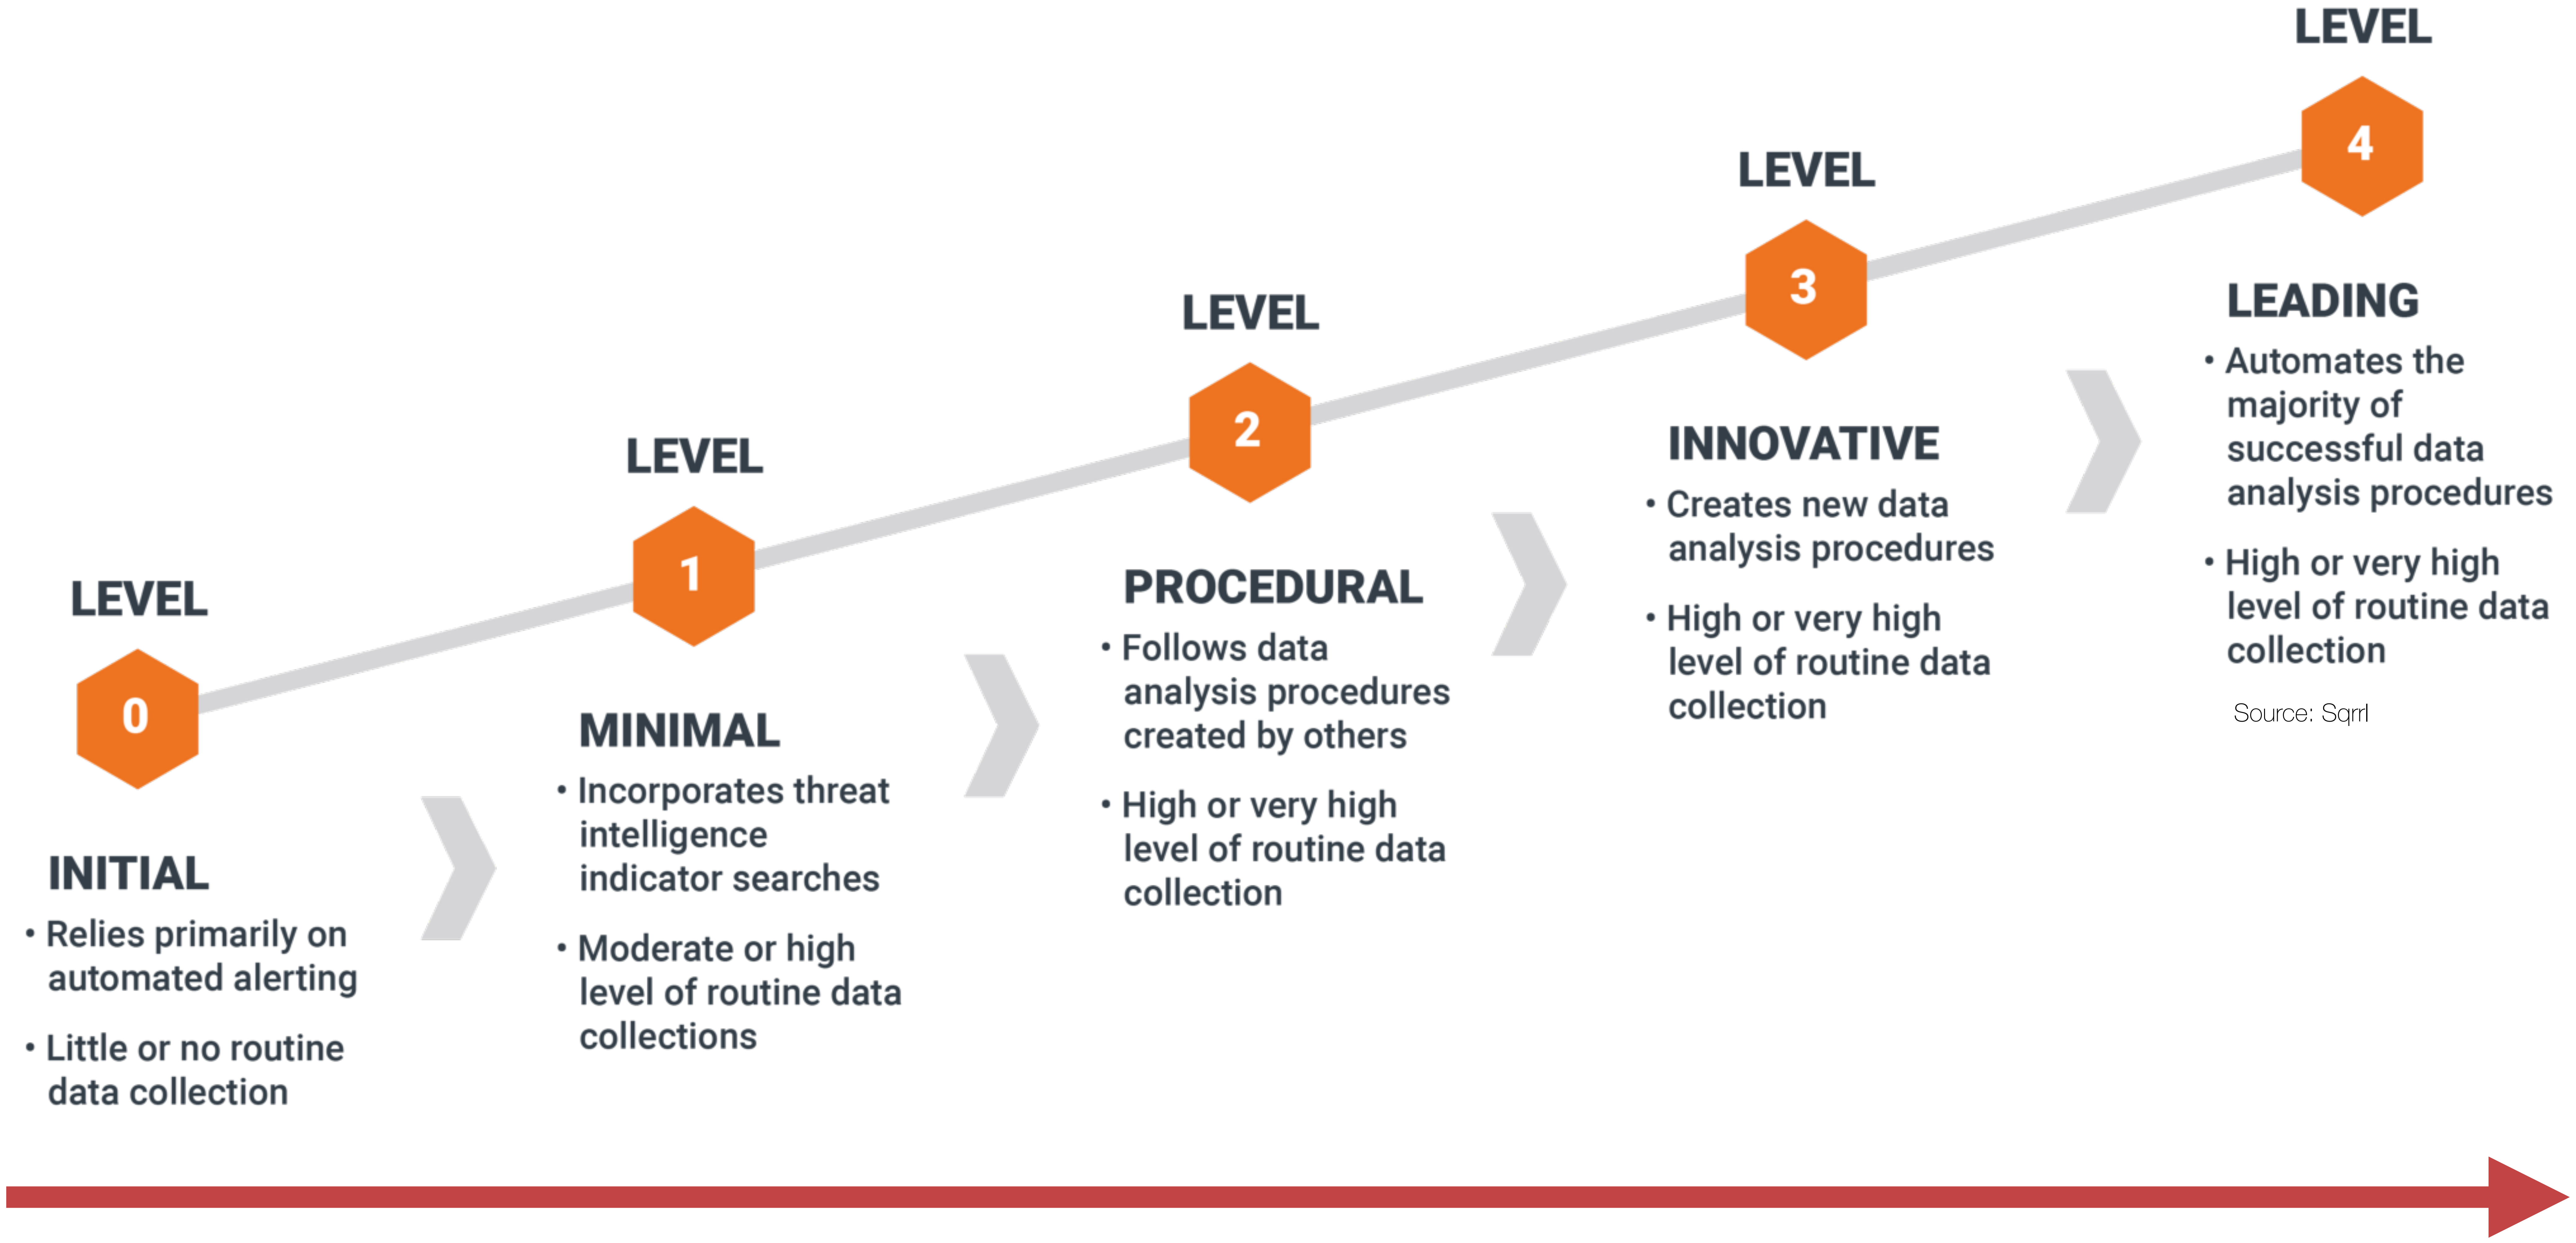

LEVEL 0

INITIAL

- Relies primarily on automated alerting
- Little or no routine data collection

LEVEL 1

MINIMAL

- Incorporates threat intelligence indicator searches
- Moderate or high level of routine data collections

LEVEL 2

PROCEDURAL

- Follows data analysis procedures created by others
- High or very high level of routine data collection

LEVEL 3

INNOVATIVE

- Creates new data analysis procedures
- High or very high level of routine data collection

LEVEL 4

LEADING

- Automates the majority of successful data analysis procedures
- High or very high level of routine data collection

Source: Sqrrl

Use AI to level up less experienced security teams

Conclusion

Support, not replace, the analyst

Automated intrusion detection

We remain limited to finding what we can describe

Paradigm Shift: Threat hunting

Assume you have been compromised already — find them

Analysts drive — tools support through visibility and automation

Add AI to the toolbox to support analysts

Provide triggers and insights; guide assessment and workflow

Use AI for what it’s good at: deriving patterns from existing data

Beyond Detection: AI's Potential For Supporting Threat Hunters

Robin Sommer

Corelight, Inc.

`robin@corelight.com`

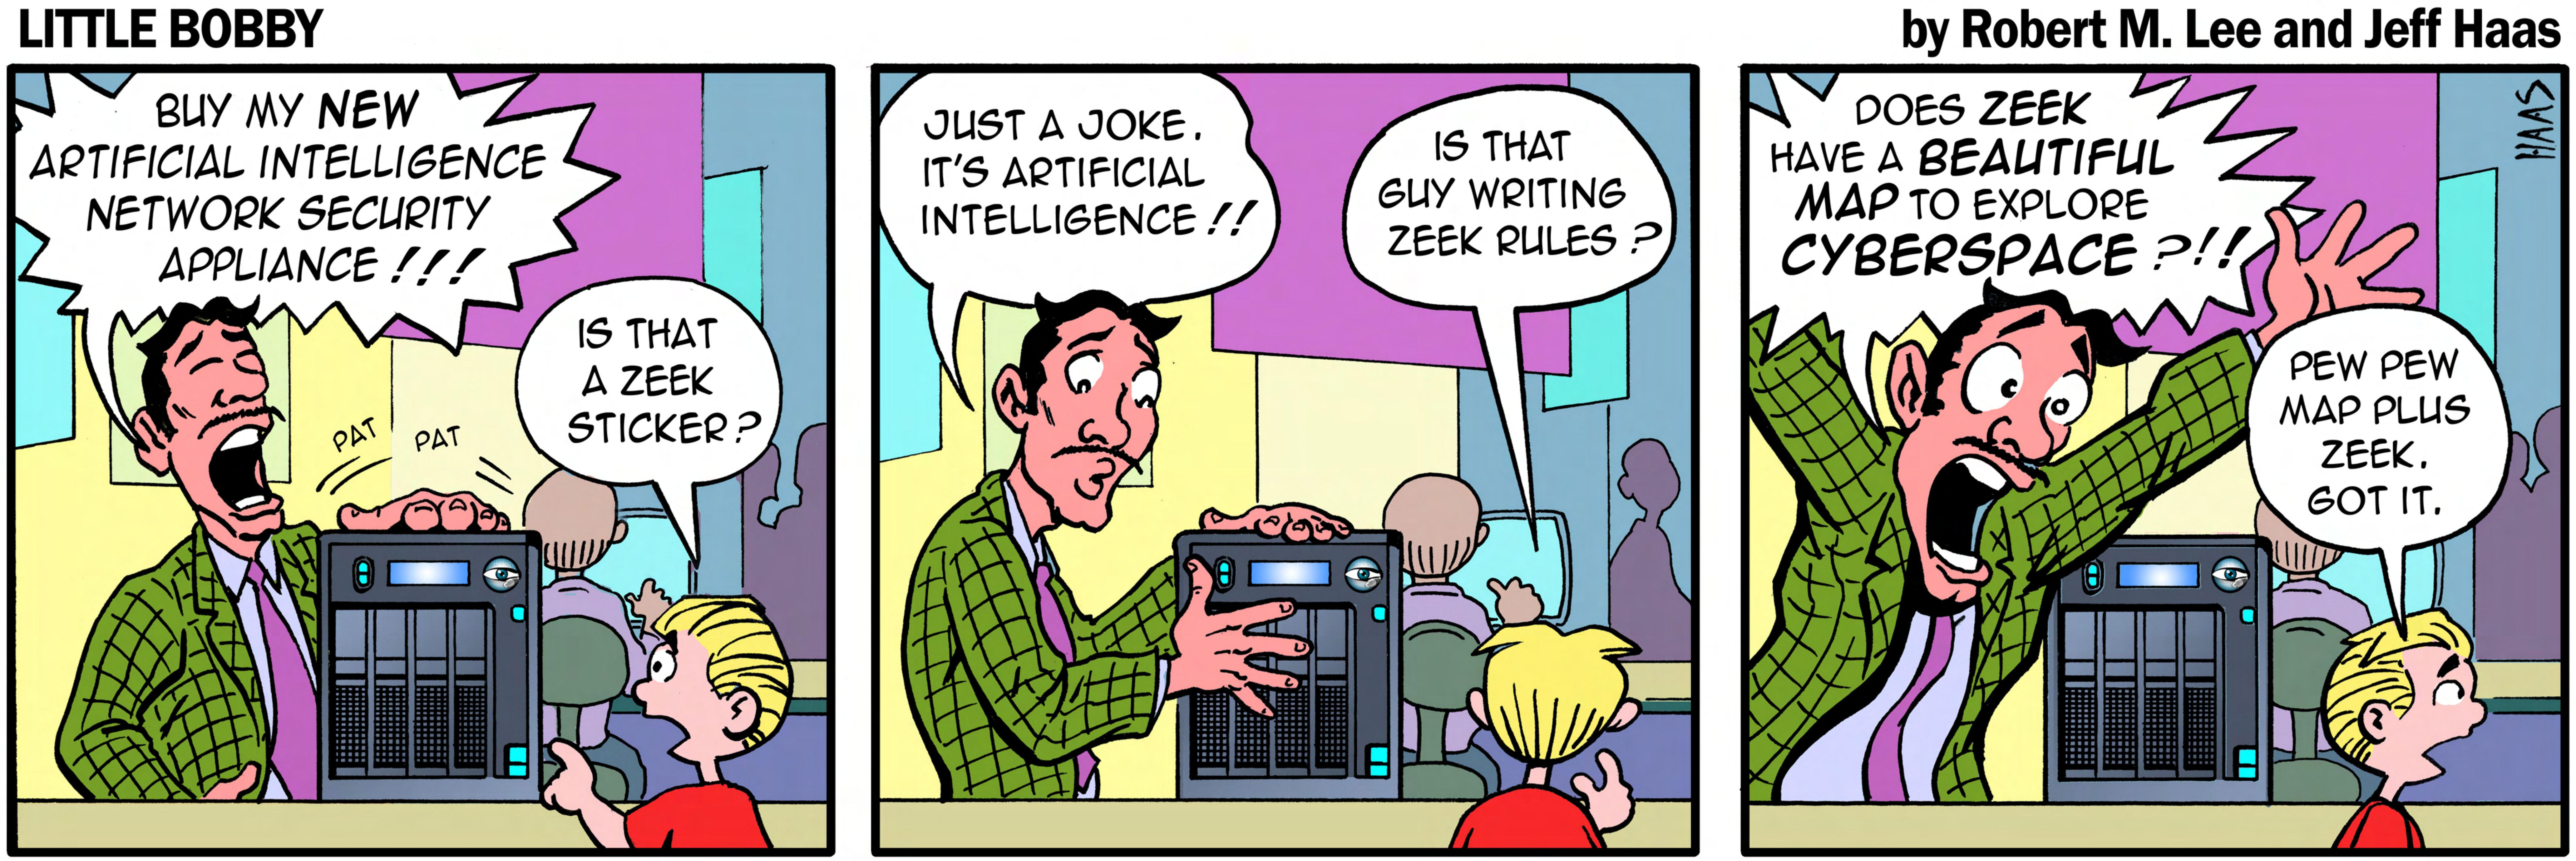
LITTLE BOBBY
by Robert M. Lee and Jeff Haas
BUY MY NEW ARTIFICIAL INTELLIGENCE NETWORK SECURITY APPLIANCE !!!
PAT PAT
IS THAT A ZEEK STICKER?
JUST A JOKE. IT'S ARTIFICIAL INTELLIGENCE !!
IS THAT GUY WRITING ZEEK RULES?
DOES ZEEK HAVE A BEAUTIFUL MAP TO EXPLORE CYBERSPACE ?!!
PEW PEW MAP PLUS ZEEK. GOT IT.
HAAS

www.littlebobbycomic.com/projects/week-239