SECURE DATA
TRUSTED CYBERSECURITY EXPERTS

Charl Van Der Walt @charlvdwalt
Sid Pillarisetty @4n0m4l1

black hat EUROPE 2018

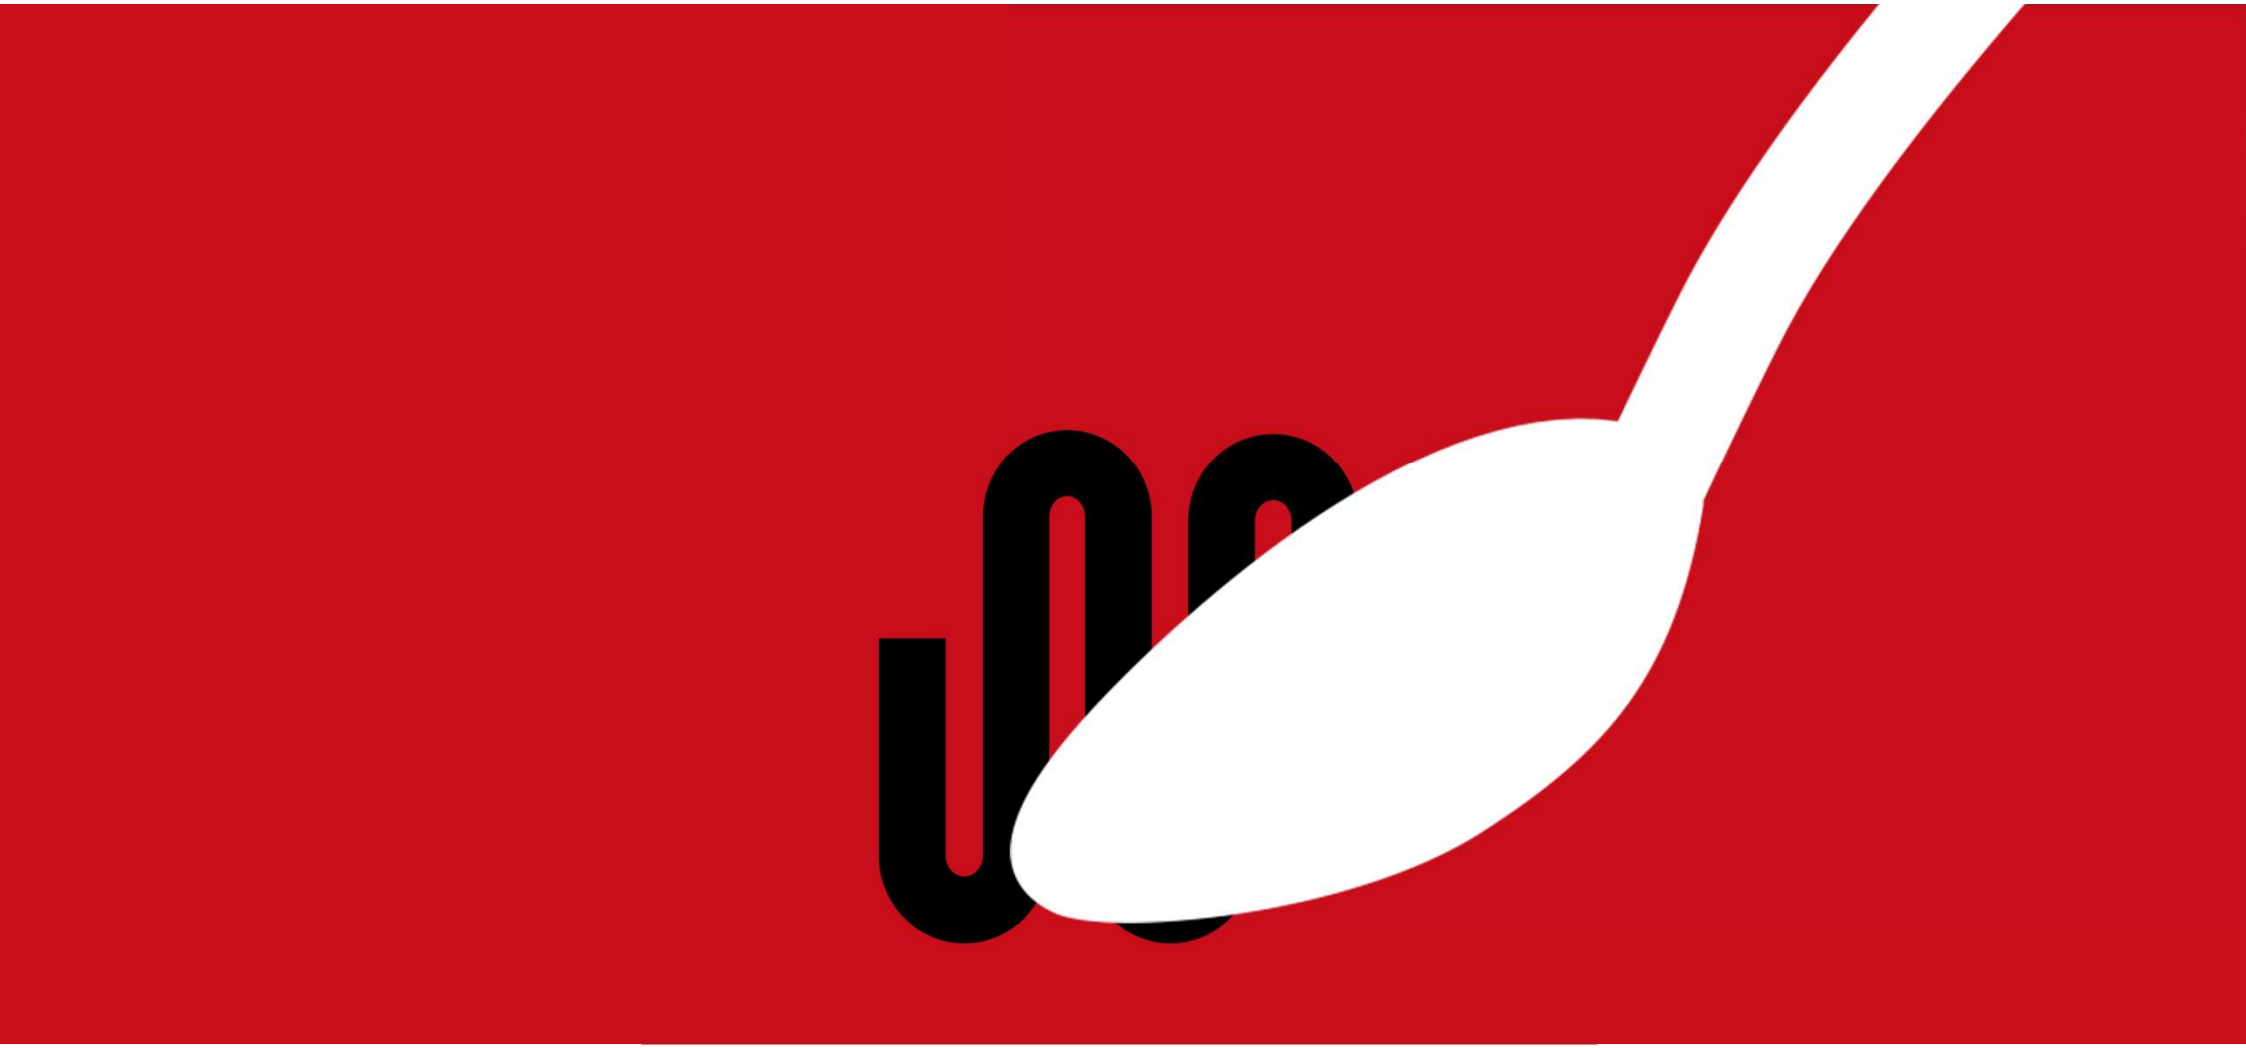

Why this research


Not secure | www.intelligenceledsecurity.com
Create your website today.
Start Now
Home
Subscribe
ILS
Intelligence Led Security LLC
OUR NEW SITE IS
COMING SOON
STAY TUNED!

Intelligence led security is the collection, aggregation, correlation and analysis of both internal and external data to understand risks, identify threat actors, discover and minimize attacks or losses already underway, and understand and predict the methods and actions of likely adversaries.

http://www.centurylink.com/
business/enterprise/blog/thinkgig/3-major-benefits-of-intelligence-led-security/

TREND: COMPLEX INDICATORS ARE MORE LIKELY TO DETECT UNKNOWN APT-RELATED ACTIVITY

Detecting the APT is incredibly difficult and many organizations are not prepared to effectively identify that they have been compromised. In most cases, initial notification of an APT intrusion originated from a third-party, primarily law enforcement. The primary reason organizations fail to identify the APT is that most of their security devices examine inbound traffic at the perimeter. Most organizations rely solely on anti-virus solutions to provide host-based monitoring. In addition, implementing the ability to monitor internal to internal communications on a network is costly and challenging. In both instances, being able to respond quickly and to deploy APT indicators is difficult, as organizations' security arsenals are not configured to monitor using this methodology.

Host- and network-based signatures used to detect malicious activity have previously consisted of data like MD5, file size, file name, and service name, etc. Although useful, the lifespan of these type of signatures is often short because attackers can routinely modify their malware to avoid detection. Although those signatures will periodically work to identify attacker activity, MANDIANT has found greater success in adapting specific signatures into what are known as Indicators of Compromise ("IOC" or "indicators").

These indicators not only look for specific file and system information, but also use logical statements that characterize malicious activity in greater detail.

MANDIANT has determined that the majority of APT custom-developed tools typically contain code segments from other, similarly developed malware. The code segments could also be upgrades to previously identified malware. Indicators derived from this information remain fairly consistent between the various malware and their subsequent upgrades. Victims are more likely to detect APT-related activity using code segments when it is possible new APT malware might be used. In many cases, previously unidentified malware and backdoors were identified through the use of these indicators in both network traffic and host-based information.

The combination of both host- and network-based indicators continues to be the most reliable way to identify APT-related malware on a network. In two separate investigations, network-based information from a generic packed file transfer revealed suspected malicious activity. Upon further research, the file transfer was identified as malicious activity that was then immediately validated through the use of host-based indicators and forensic analysis.

The first documented appearance of the term indicators of compromise, or IOCs, in the modern context is from the first Mandiant M-Trends report, published on **25 Jan 2010**

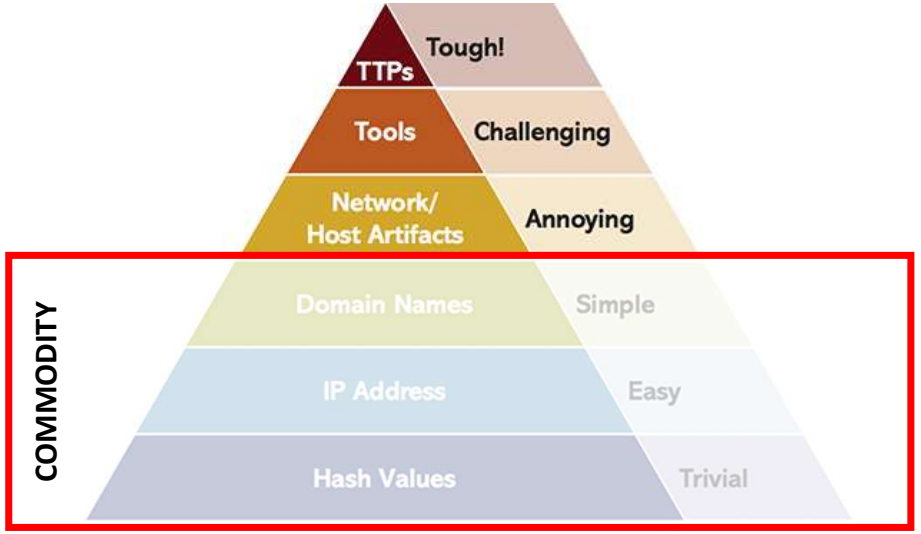

Source: David J. Bianco, personal blog

DATA MONITORED MONTHLY

12 MILLION
Unique URLs

640 MILLION
Unique Users

1.2 BILLION
Unique Devices

2 Million Threat Events Every Hour
8 Million Unique Compromised Devices Daily

This IP list is a composition of other IP lists.

The objective is to create a blacklist that can be safe enough to be used on all systems, with a firewall, to block access entirely, from and to its listed IPs.

The key prerequisite for this cause, is to have no false positives. All IPs listed should be bad and should be blocked, without exceptions.

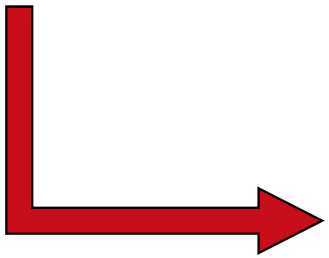

ipset entries	6,801	min: 6,706 max: 6,848
unique IPs	632,286,314	min: 632,286,314 max: 632,811,060
source	(not a url)	
local copy	download local copy	
changesets	github commit log	
check frequency	1 minute	
average update frequency	48 minutes	

```
6500  207.10.232.16
6501  207.10.232.21
6502  207.22.192.0/18
6503  207.32.128.0/19
6504  207.32.208.0/20
6505  207.45.224.0/20
6506  207.47.71.46
6507  207.58.163.118
6508  207.58.168.91
6509  207.[illegible]5
6510  207.107.101.210
6511  207.110.64.0/18
6512  207.110.128.0/18
6513  207.134.189.64
6514  207.140.14.141
6515  207.177.101.10
6516  207.183.192.0/19
```

0.2%?

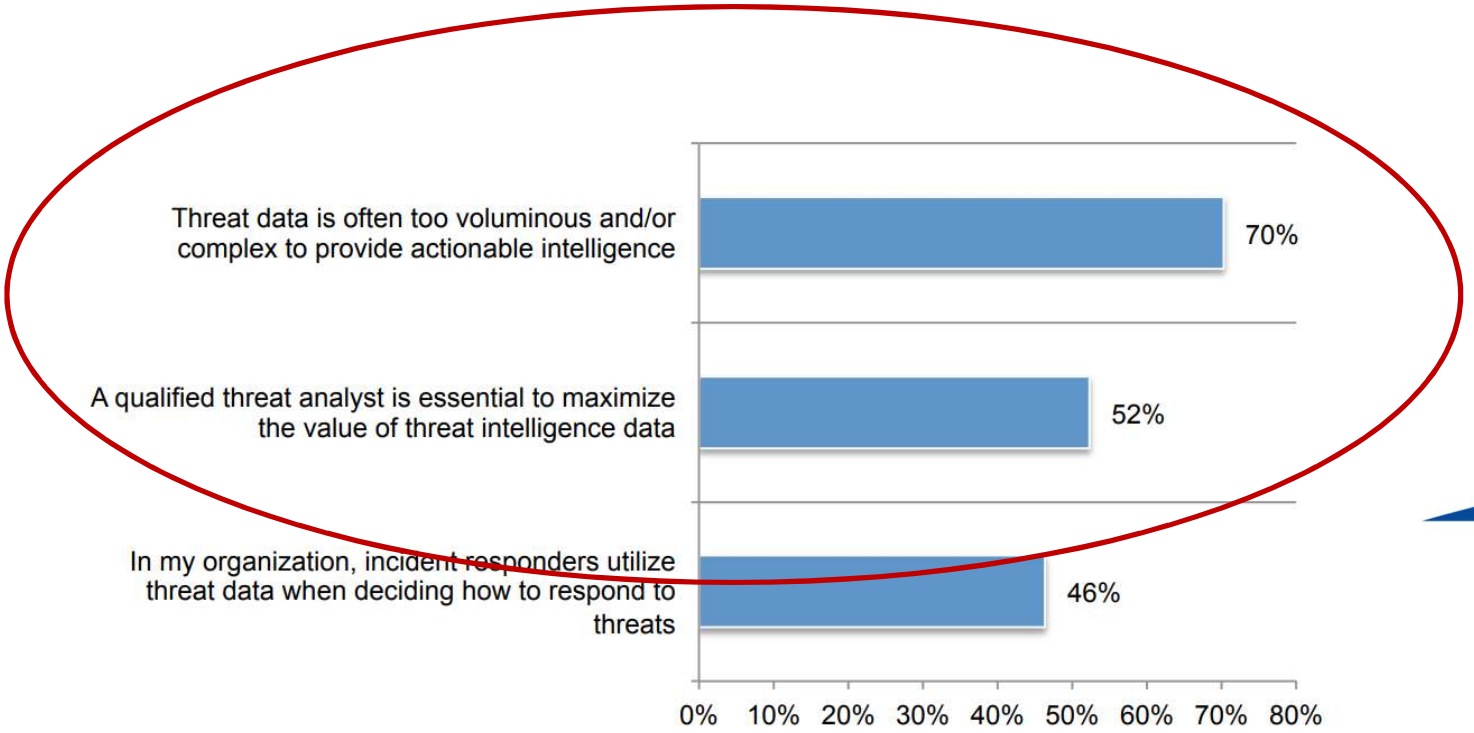

Ponemon INSTITUTE

Dealing with the volumes.

Gartner

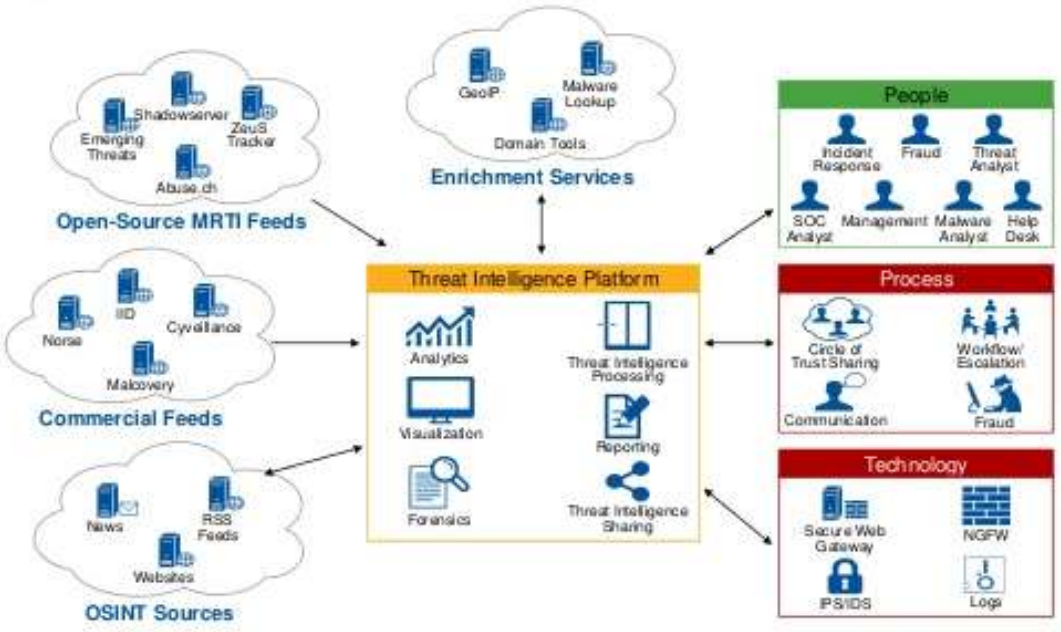

Intelligence Analysts

Intelligence Analysts face a tremendous workload in combating cyber threats. To improve the odds, they need tools that quickly sort through structured and unstructured information for relevancy; that enable collaboration through a single, centralized workspace; and that eliminate manual and repetitive work.

EclecticIQ Platform

EclecticIQ Platform empowers analysts to optimize their workflow using with automation tools based on analytics. Instead of manually crunching through data, analysts can better spend their time on collaboration with peers, working to enrich, qualify, analyze and share threat information to stakeholders.

- Automation based on analytics
- Analyze and share threat information to stakeholders

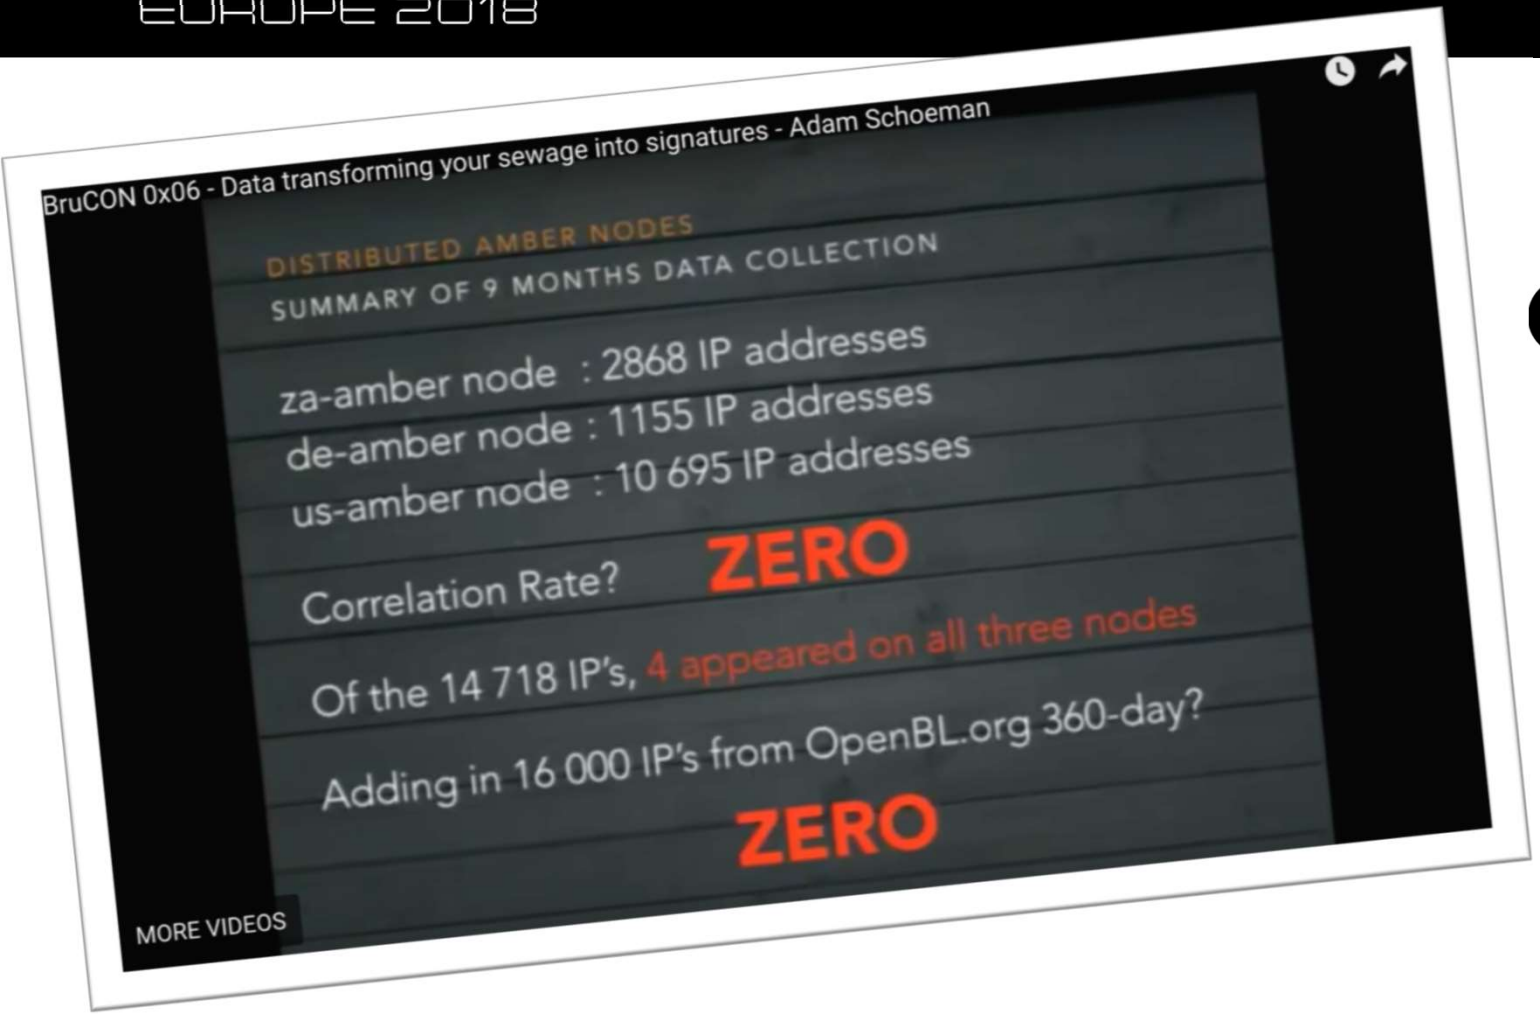

https://www.youtube.com/watch?v=M_BppG-wXC8

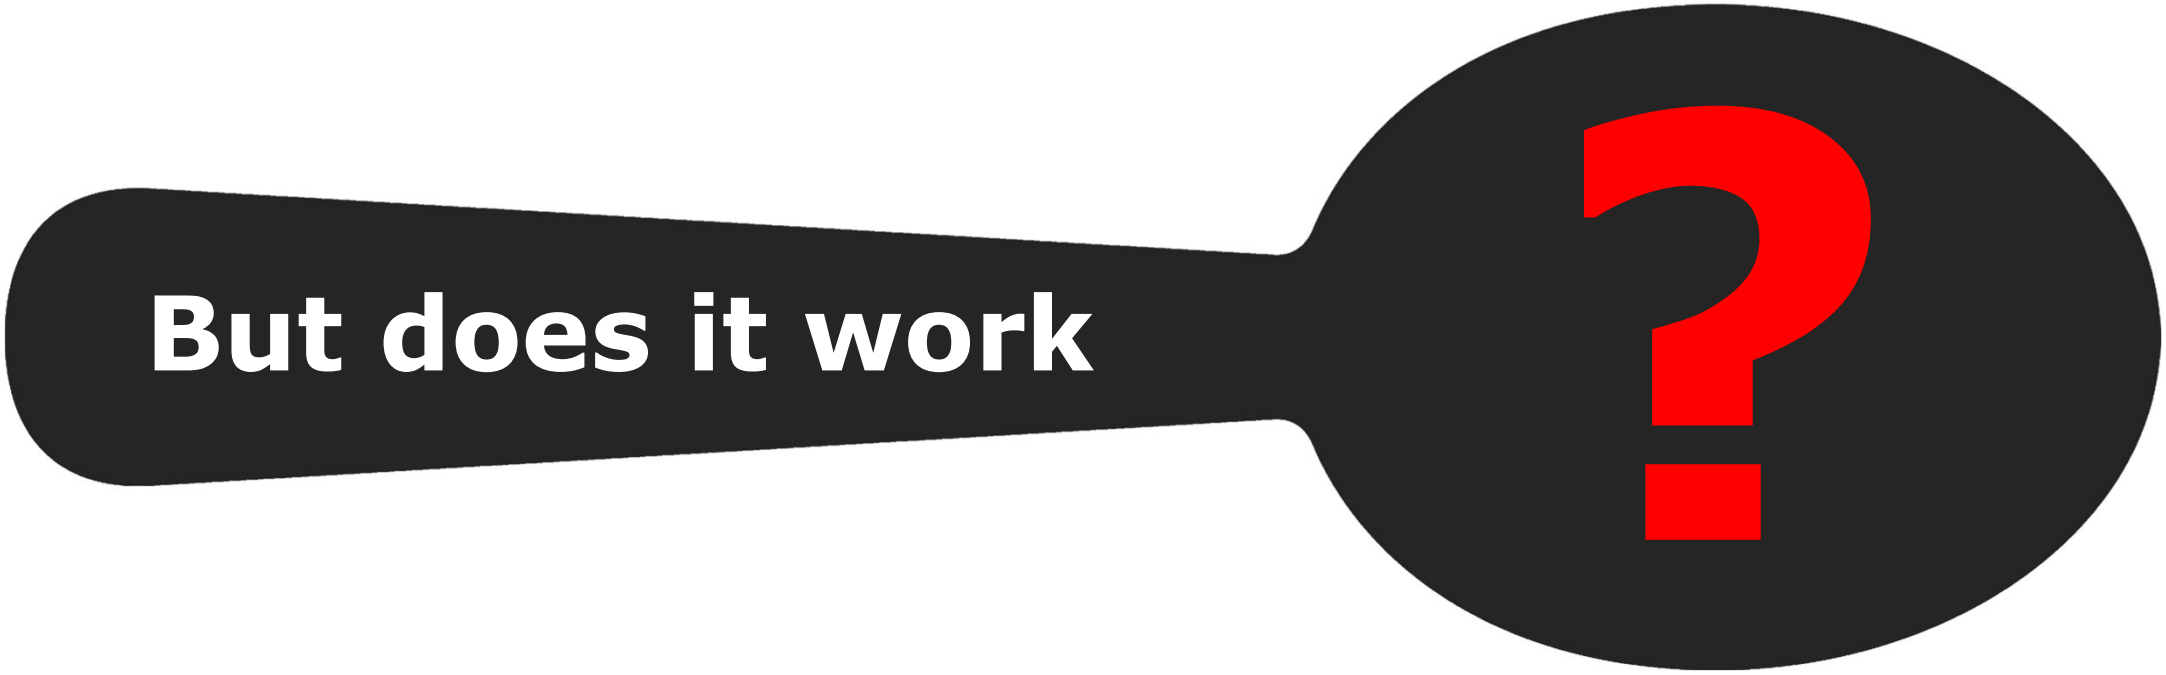
But does it work
?

Prior Work.

Paper	Authors	Date	Reference
Everything You Wanted to Know About Blacklists But Were Afraid to Ask	Leigh Metcalf Jonathan M. Spring CERT Network Situational Awareness Group	September 2013	https://christian-rossow.de/publications/blacklists-raid2014.pdf
On Comparing Threat Intelligence Feeds	Anton Chuvakin	January 2014	https://blogs.gartner.com/anton-chuvakin/2014/01/07/on-comparing-threat-intelligence-feeds/
Measuring the IQ of your Threat Intelligence Feeds (#tiqtest)	Alex Pinto Kyle Maxwell	August 2014	https://www.slideshare.net/AlexandrePinto10/defcon-22-measuring-the
Evaluating Threat Intelligence Feeds	Paweł Pawlinski Andrew Kompanek	February 2016	https://www.first.org/resources/papers/munich2016/kompanek-pawlinski-evaluating-threat-ntelligence-feeds.pdf

Prior Work.

NOVELTY: How frequently are lists updated?

OVERLAP: How unique are the lists?

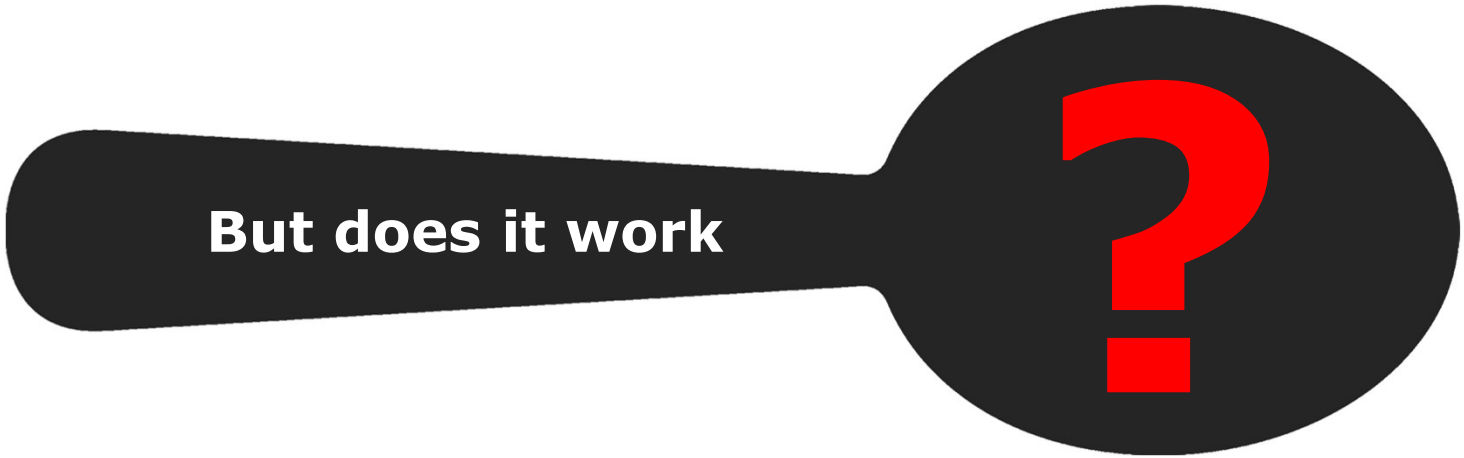

How efficient is Threat Intelligence about the behaviour of an IP in predicting future behaviour by that same IP

- A controlled experiment
- To answer a very specific question
- About *Internet* Threat Intelligence feeds
- Using a transparent methodology
- On a (limited) proprietary dataset
- Share findings, observations and emerging new questions

As luck would have it, we may be able to confuse this issue with some *facts*.

https://github.com/SecureDataLabs/BlackHat-EU-2018

Perimeter Security Logs

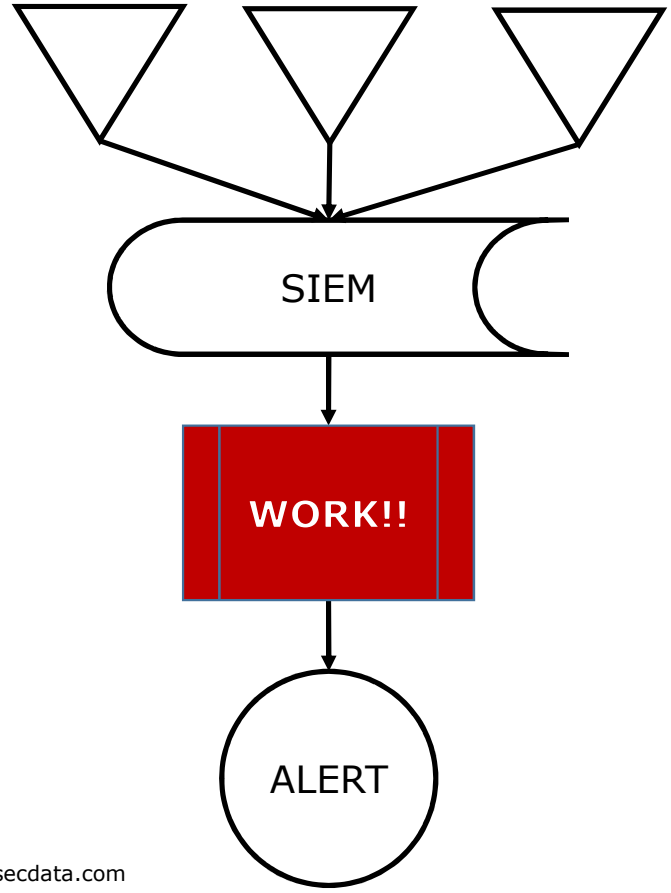

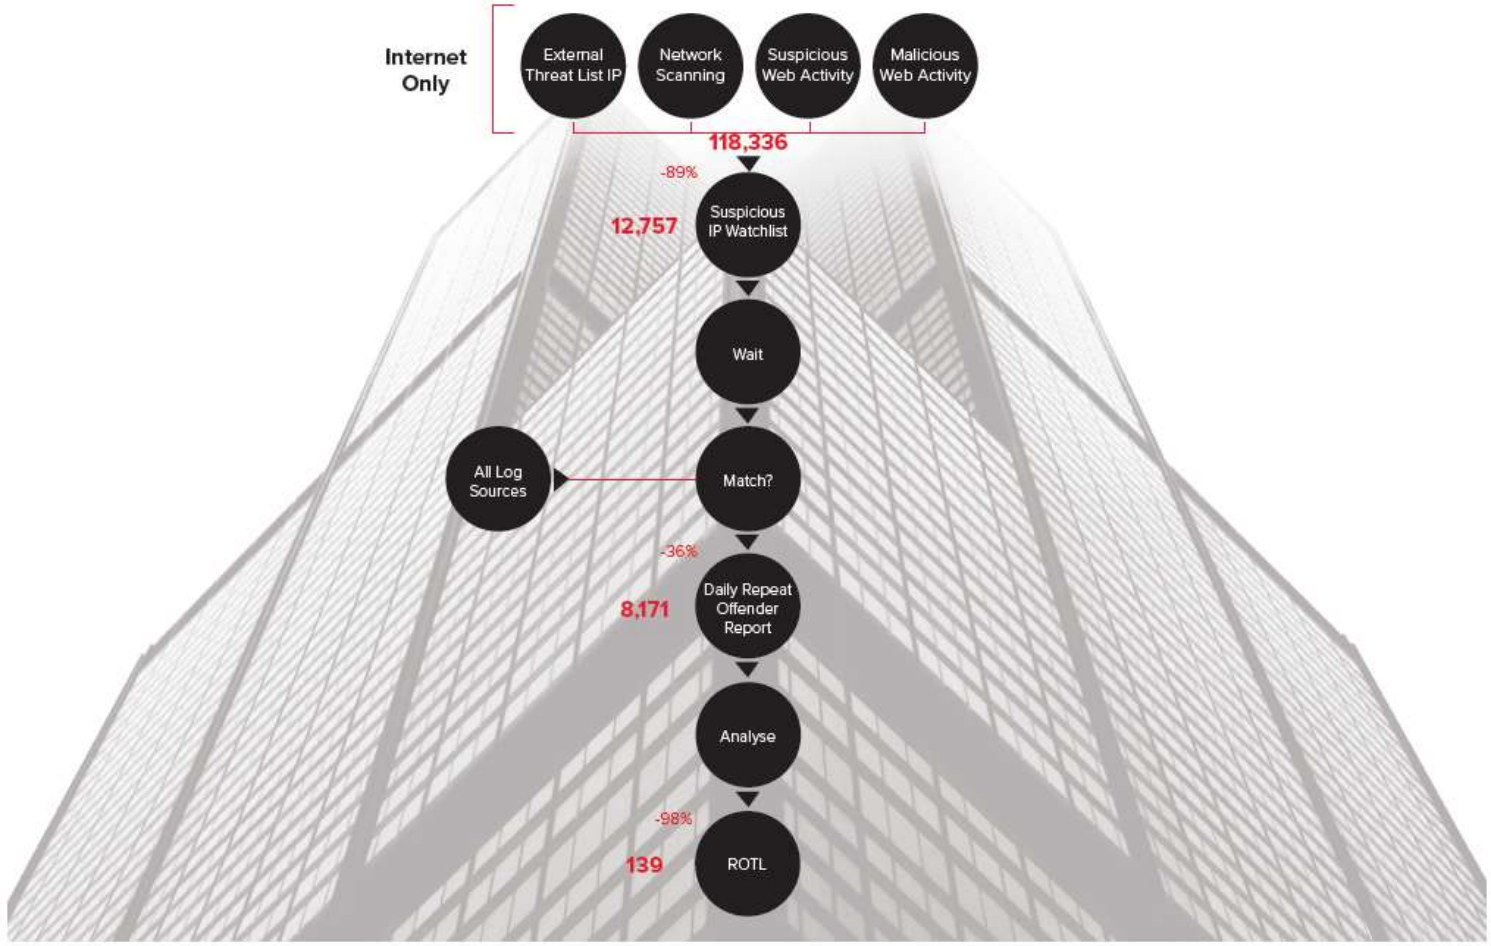

9 'Sensors'
SIEM Alarms on Internet-facing log sources

41 'Entities'
Separate customers or customer locations

	Finance	General	Web Services	Insurance	Media	Retail	Solicitors	Technology	HoneyNet
Suspicious and persistent				Day 2					
External Threat Intelligence									
Suspicious Web Activity			Day 2		Day 2				
Malicious Web Activity									
Suspicious Internet Activity	Day 1		Day 3						Day 7
Malicious Internet Activity									

Prediction

Prediction

Prediction

Day 1	Finance	Suspicious Internet Activity
Day 2	Insurance	Suspicious and persistent
Day 7	HoneyNet	Suspicious Internet Activity

Threat Data

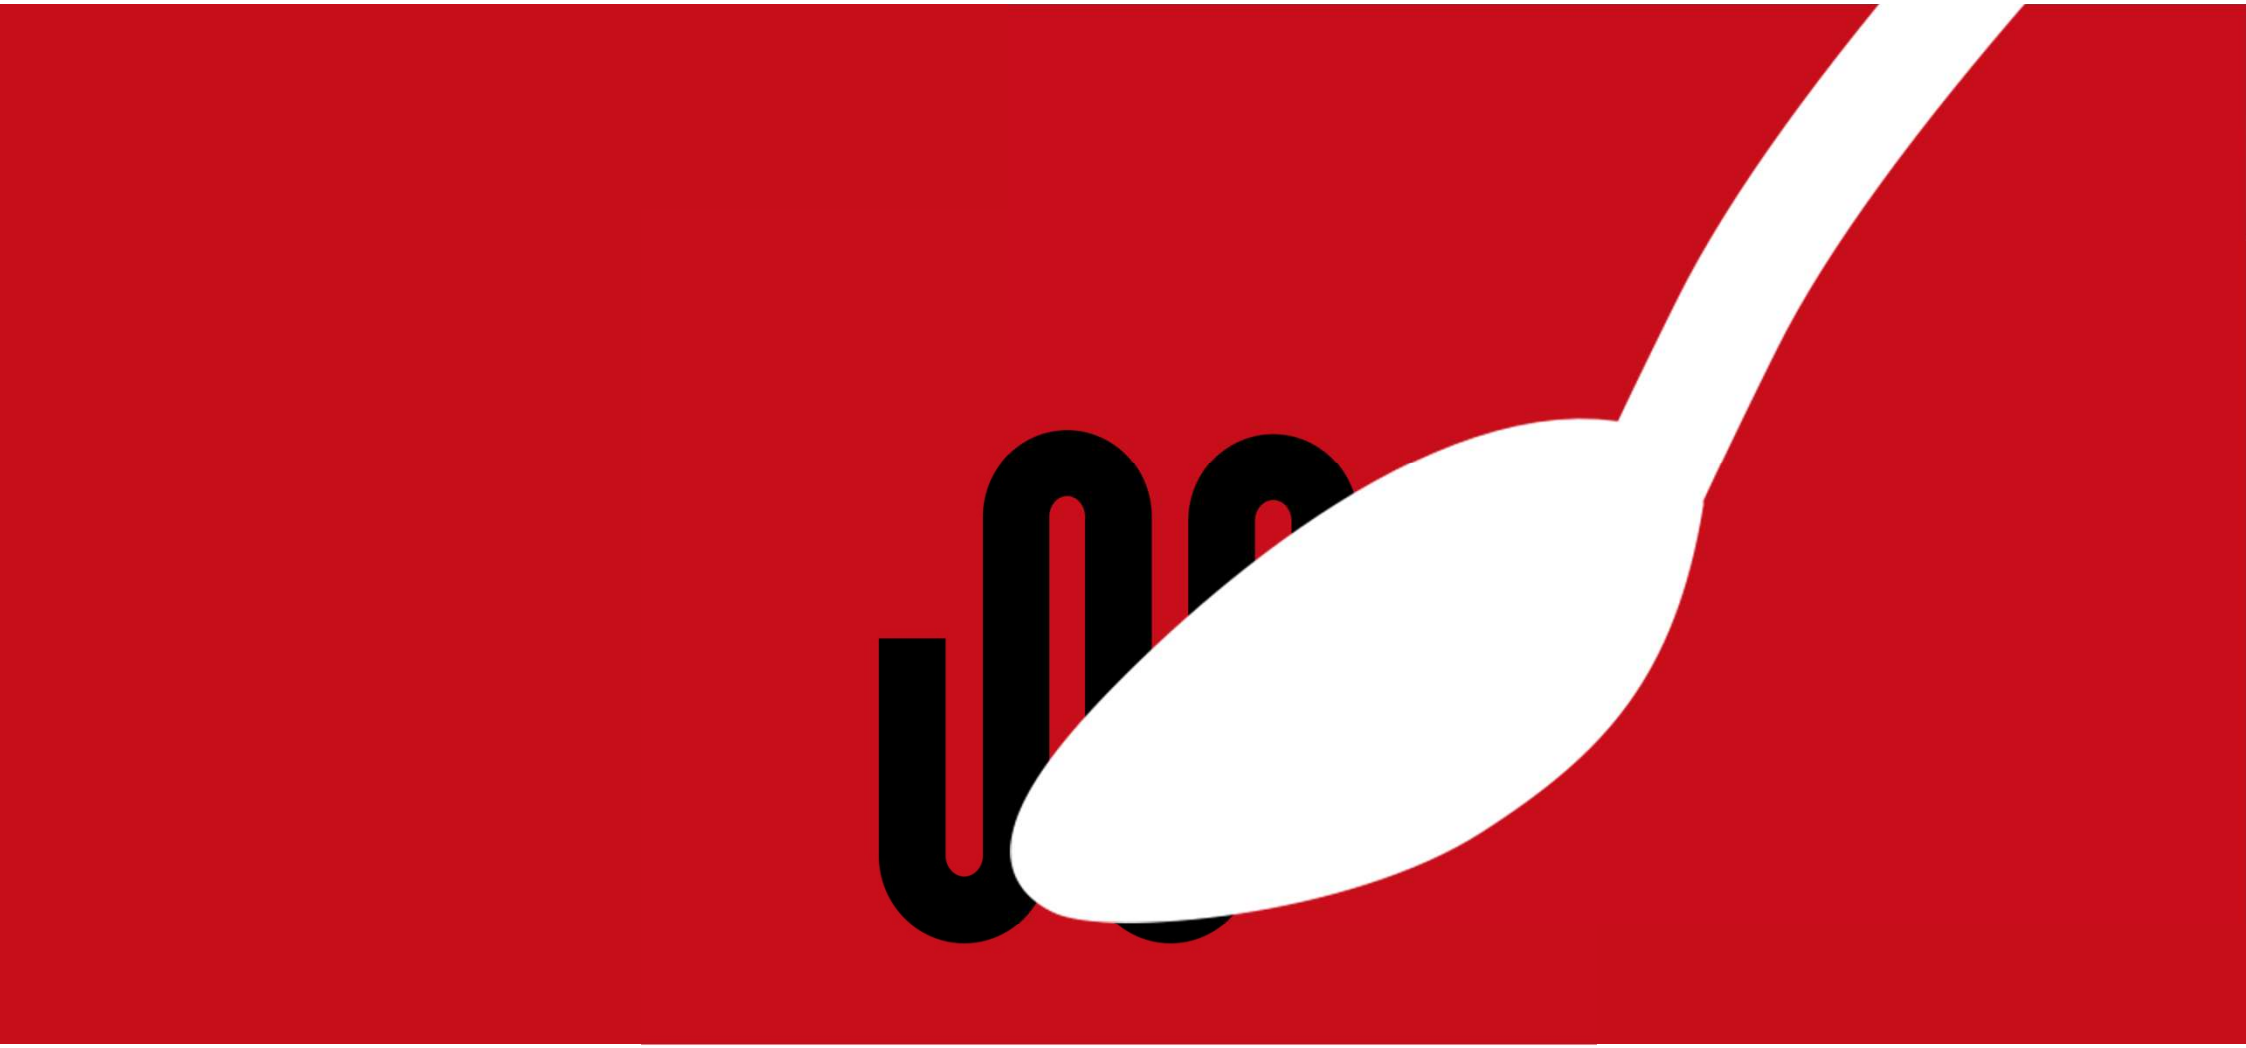

Introducing the data

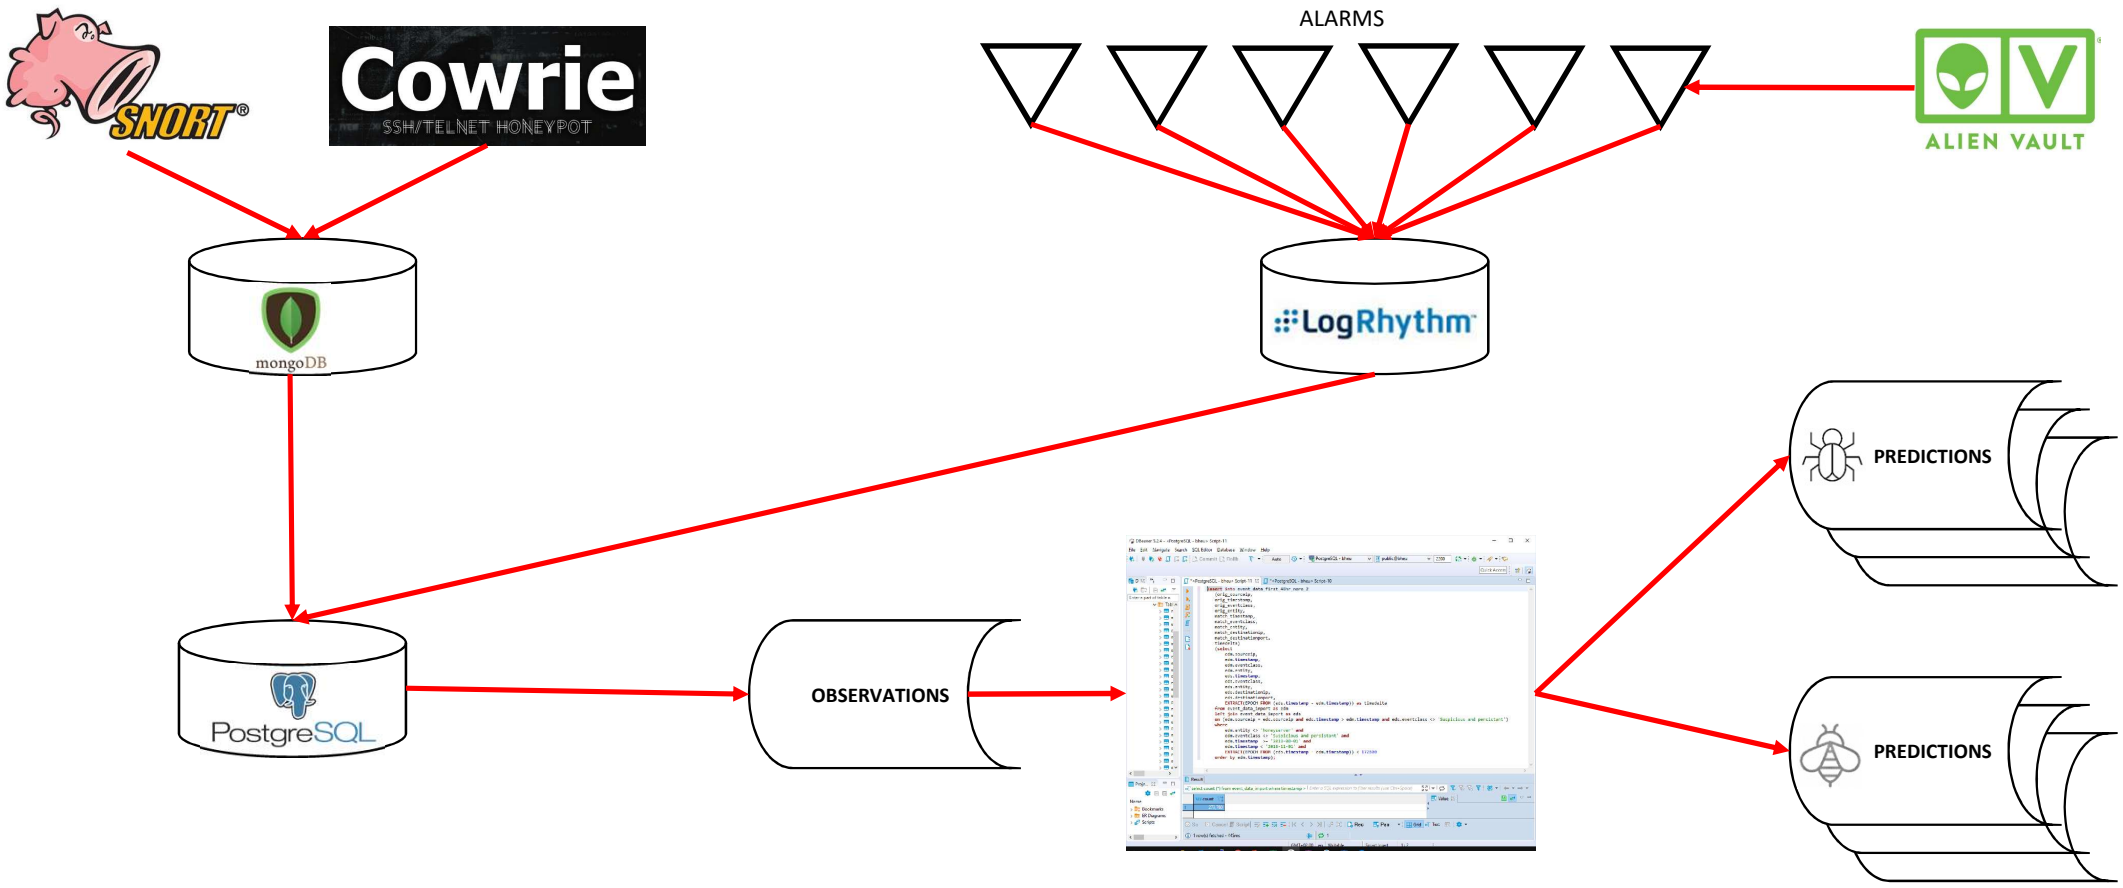
SNORT
Cowrie
SSH/TELNET HONEYPOT
ALARMS
ALIEN VAULT
mongoDB
LogRhythm
PostgreSQL
OBSERVATIONS
PREDICTIONS
PREDICTIONS

https://threatstream.github.io/mhn/

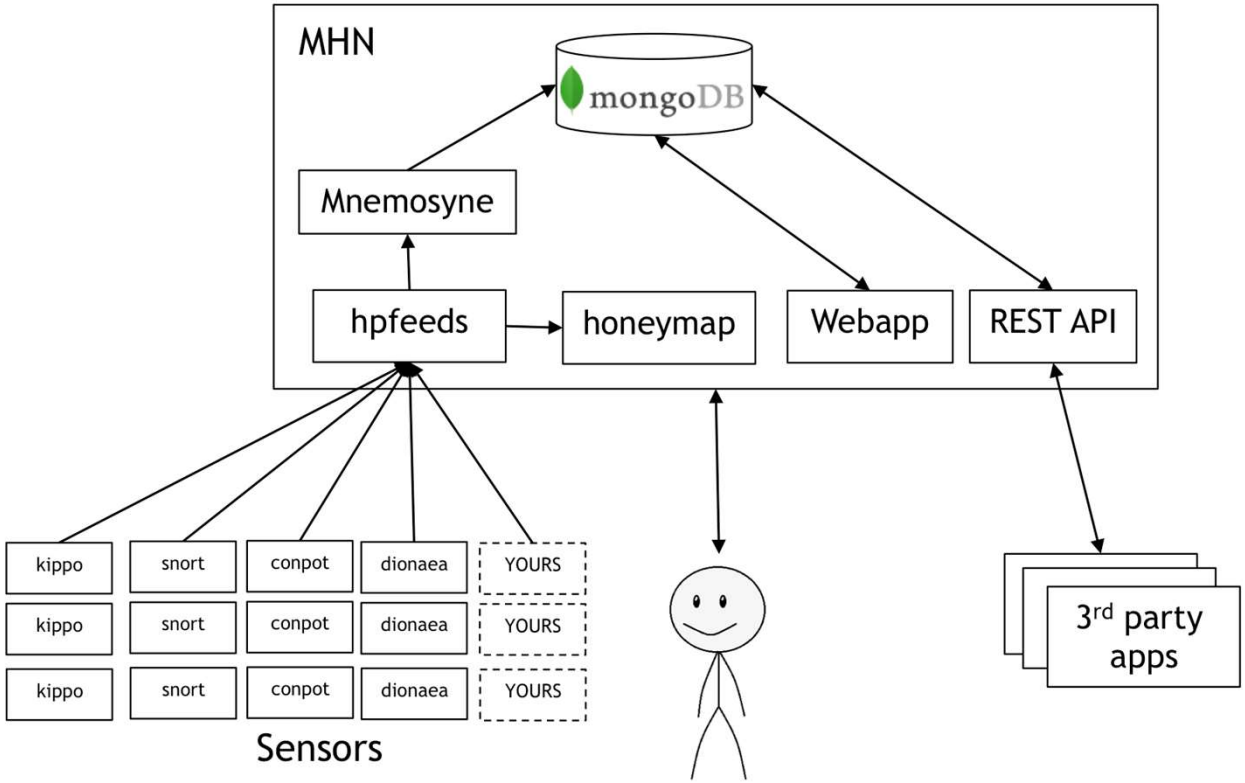

Modern Honey Network

- **3 honeypots**
  - Australia
  - Great Britain
  - USA
- **Snort**
  Open Source Emerging Threats
- **Cowrie**
  SSH

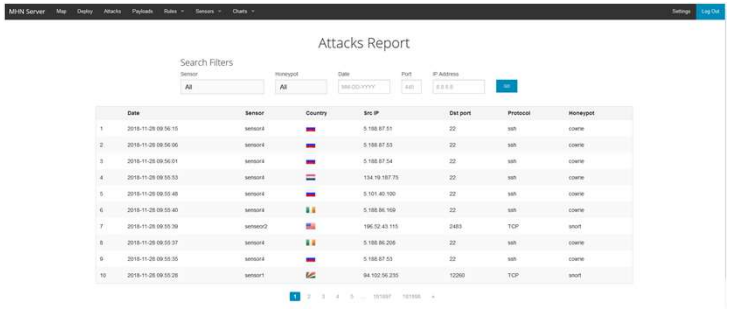

Observations Collected over Time

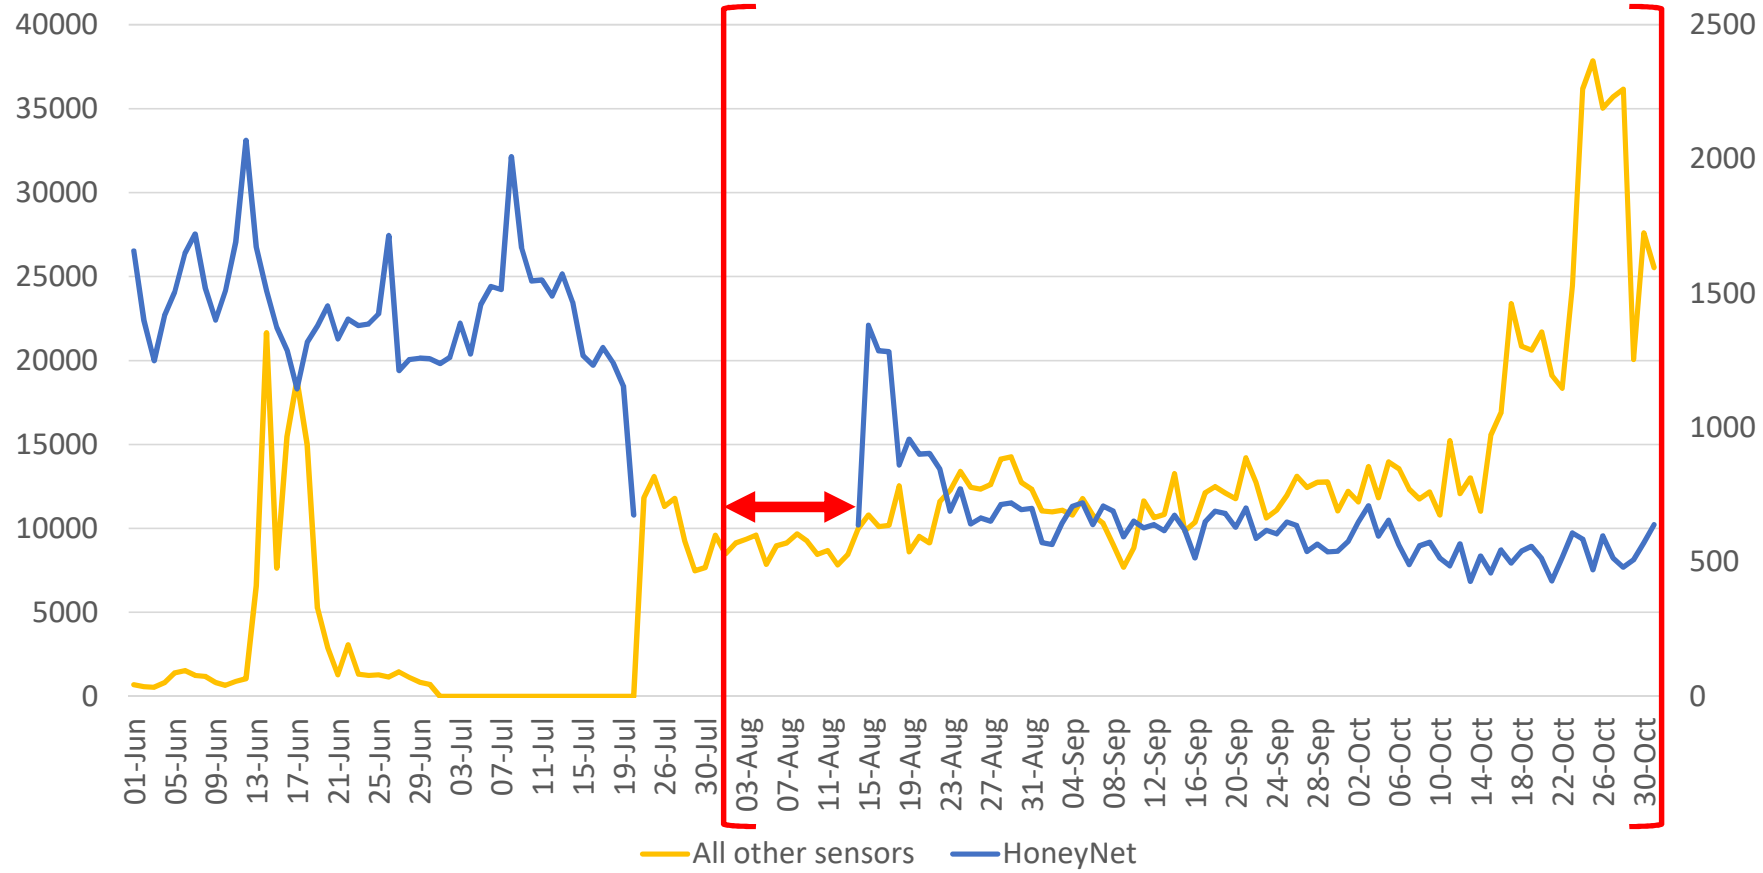

1,475,055
Observations by our own Sensors

124,241
Observations across three honeypots

Observations by the Number.

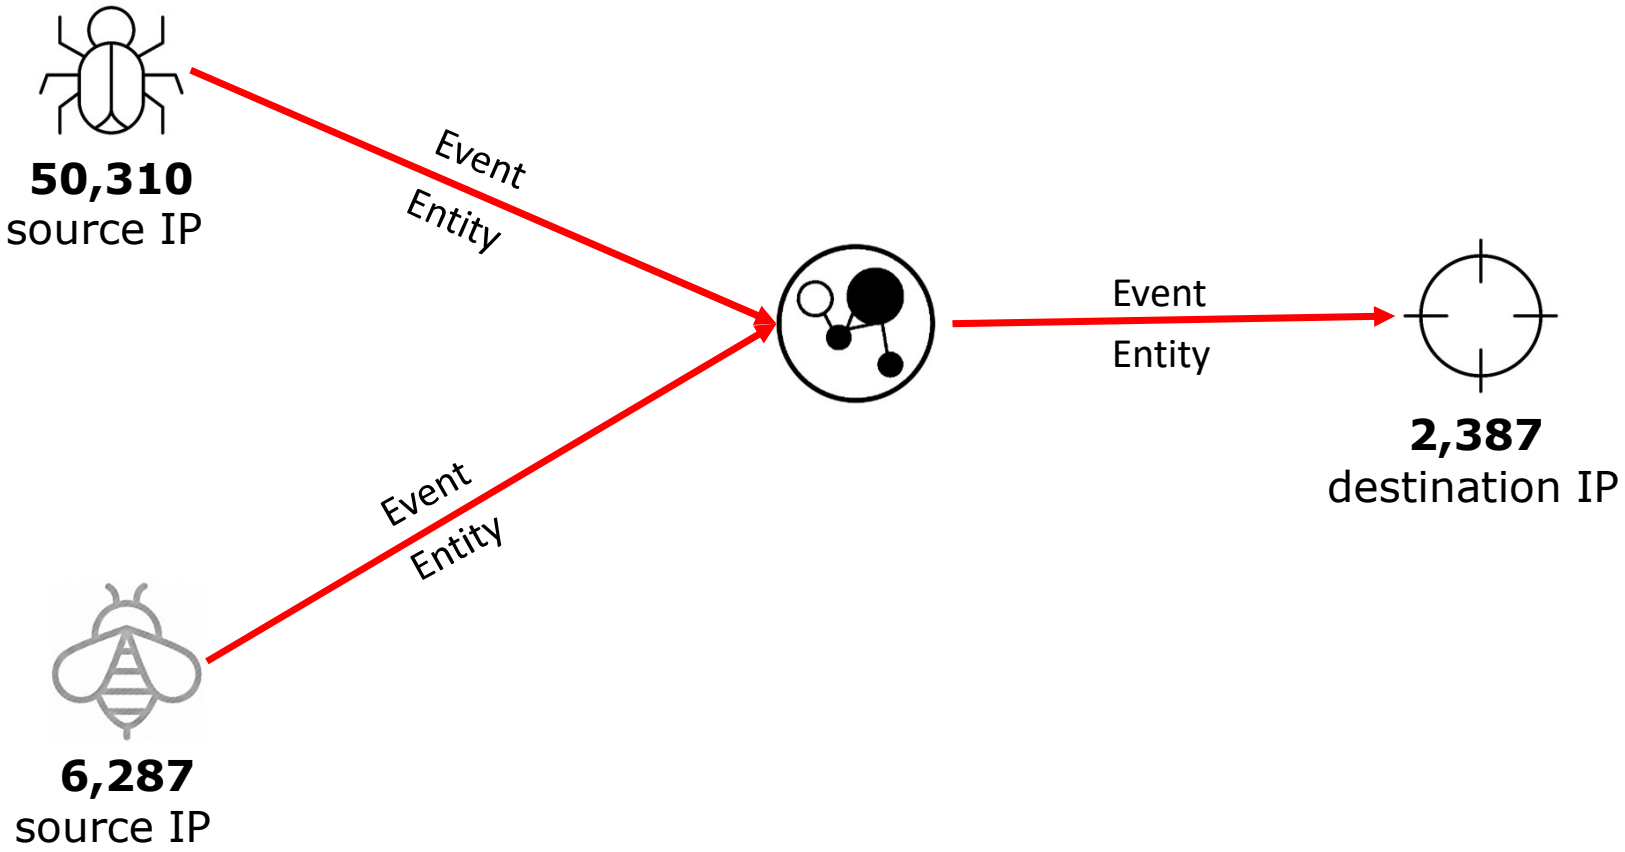

222,437
observations by our own Sensors

51,065
observations by honeypots

Unique IP per Sector

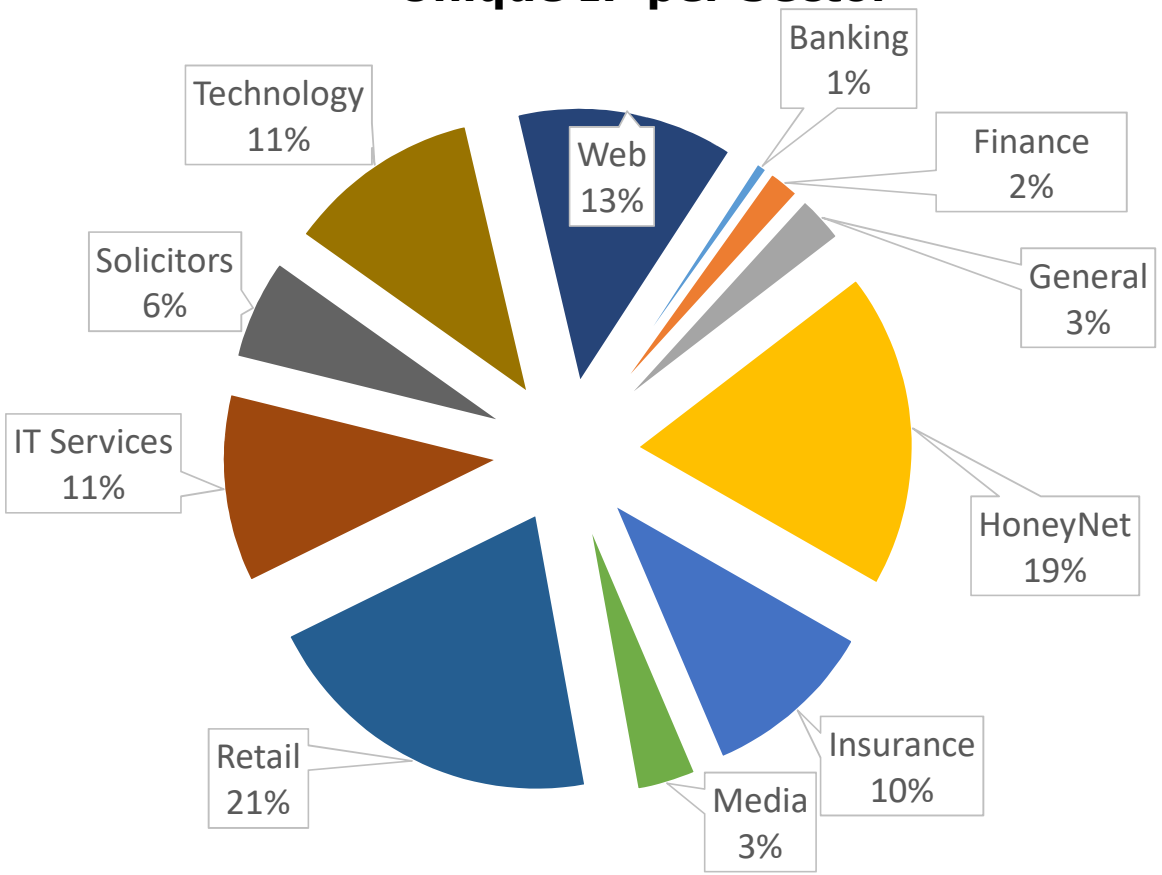

50,310
Unique IP observed by our own Sensors

6,287
Unique IP observed by honeypots

Rule Name	Category
Repeat Offender	Suspicious and persistent
Network Anomaly: Ext : Threat List IP - Allow	External Threat Intelligence
Arbor Blocked IP Then seen on ASM	Malicious Web Activity
F5 WAF Alarm Triggered	Malicious Web Activity
External IPS high severity Alert	Malicious Internet Activity
Recon - Port Scan	Suspicious Internet Activity
Suspect - URL Request Rate	Suspicious Web Activity
Suspicious Web Activity	Suspicious Web Activity
Suspicious - HTTP Error Code Rate	Suspicious Web Activity
Sucuri WAF Alerts	Malicious Web Activity

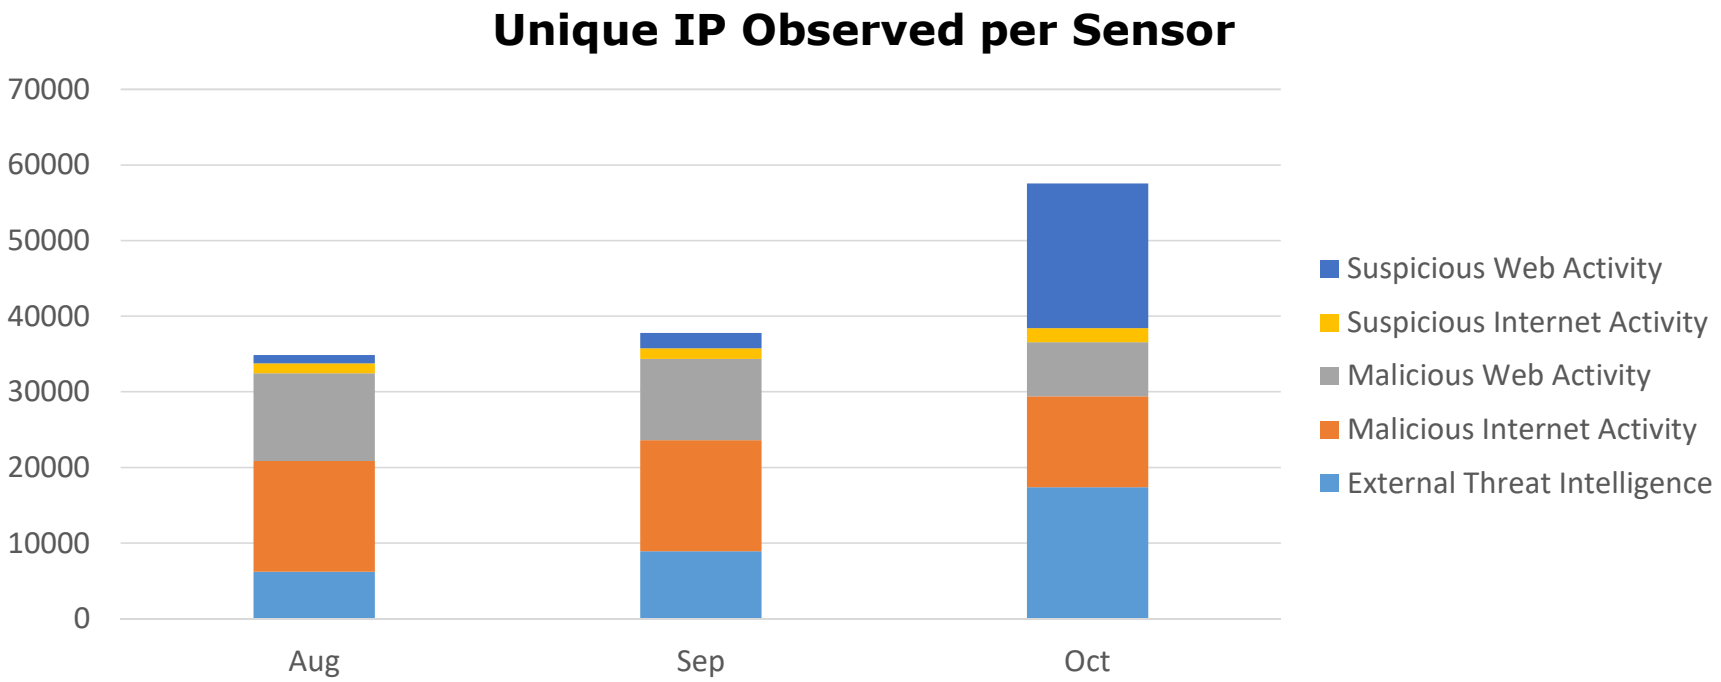

50,310
Unique IP observed by our own Sensors

6,287
Unique IP observed by honeypots

Proportion of Unique IP Observations per Honeypot

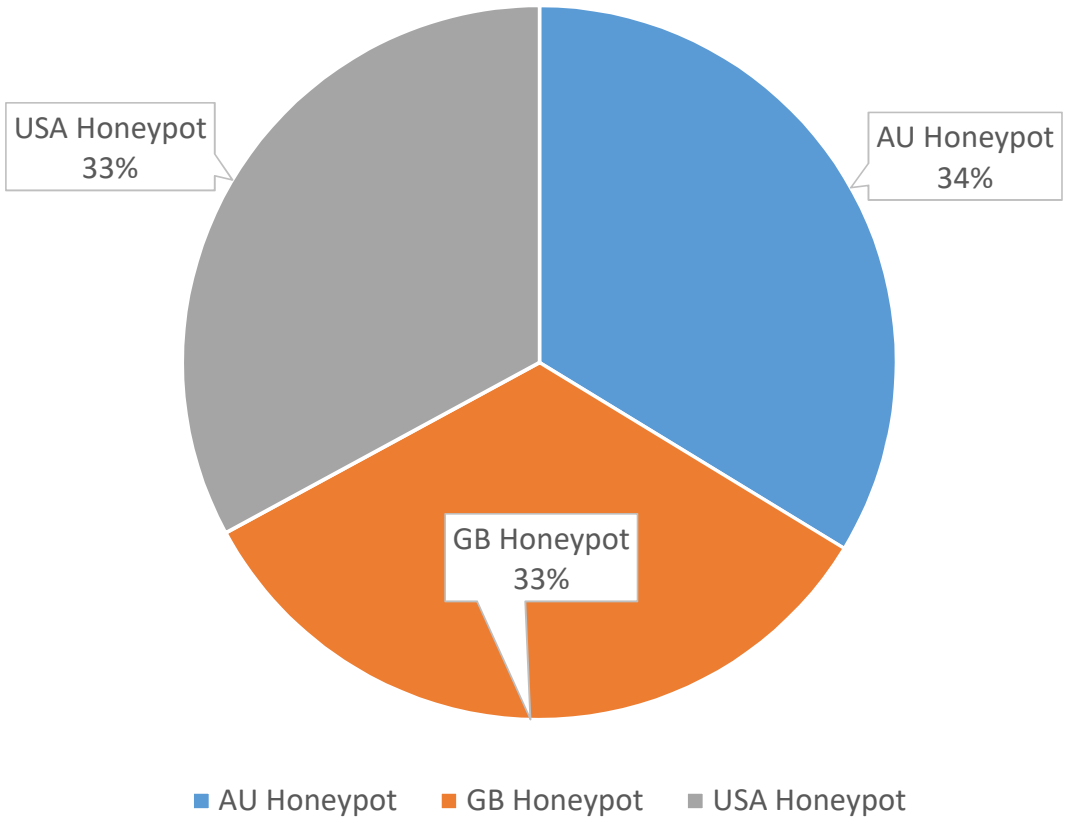

6,287
Unique IP observed by honeypots

- **Mondern Honey Network**
- 3 honeypots
- Snort
- Cowry

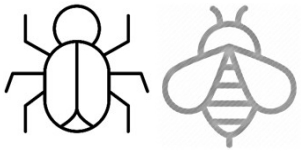

Example *Observation*.

ID	Timestamp	Entity	Event	oIP	dIP
1723823	01/06/2018 11:07	General G 1	Suspicious Web Activity	159.xxx.yyy.70	
1723825	01/06/2018 11:07	Web service A 1	Malicious Web Activity	77.xxx.yyy.108	
1723830	01/06/2018 11:18	Media A 1	External Threat Intelligence	209.xxx.yyy.4	195.xxx.yyy.196

oIP is detected by ***Sensor***[x] at an ***Entity***[x] at ***Time***[x]

Example Prediction.

oIP	oTimeStamp	oEventClass	oEntity	pTimeStamp	pEventClass	pEntity	deltaT
159.xxx.yyy.70	01/08/2018 11:07	Suspicious Web Activity	General G 1	01/09/2018 11:06	Suspicious Web Activity	General G 1	2678341
159.xxx.yyy.70	02/08/2018 11:44	Suspicious and persistent	General G 1	12/10/2018 06:53	Suspicious and persistent	Banking A 1	6116949

oIP is observed by Sensor[***x***] at an Entity[***x***] at Time[***x***] before being observed by another Sensor[***y***] at Entity[***y***] at Time[***y***] within ***Delta[t]***

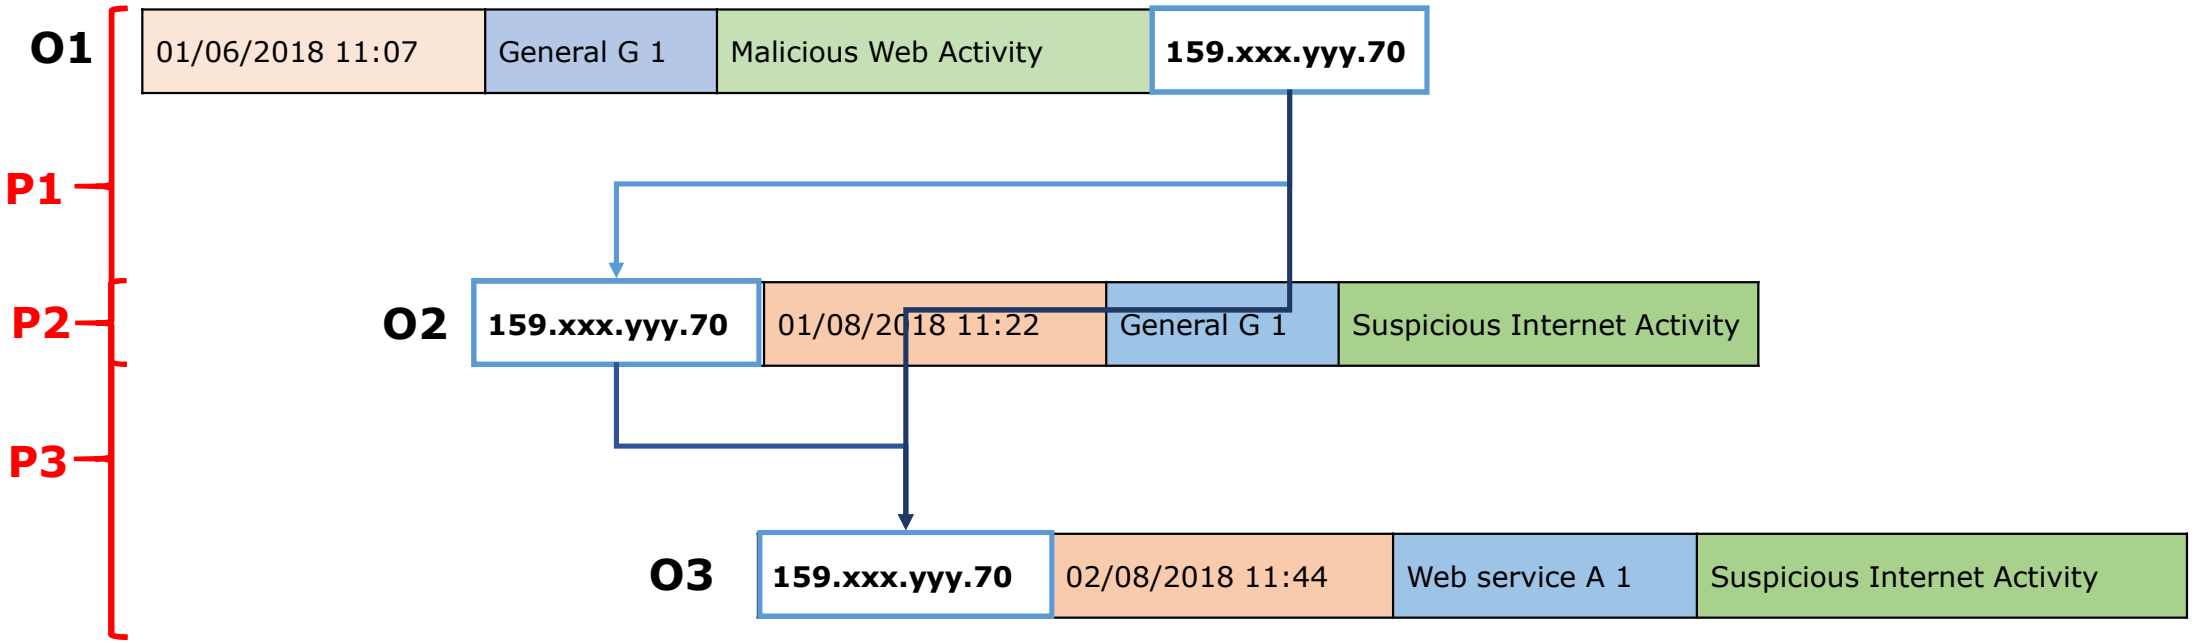
O1
01/06/2018 11:07
General G 1
Malicious Web Activity
159.xxx.yyy.70
P1
P2
O2
159.xxx.yyy.70
01/08/2018 11:22
General G 1
Suspicious Internet Activity
P3
O3
159.xxx.yyy.70
02/08/2018 11:44
Web service A 1
Suspicious Internet Activity

Number of Predictions per Prediction Time Window

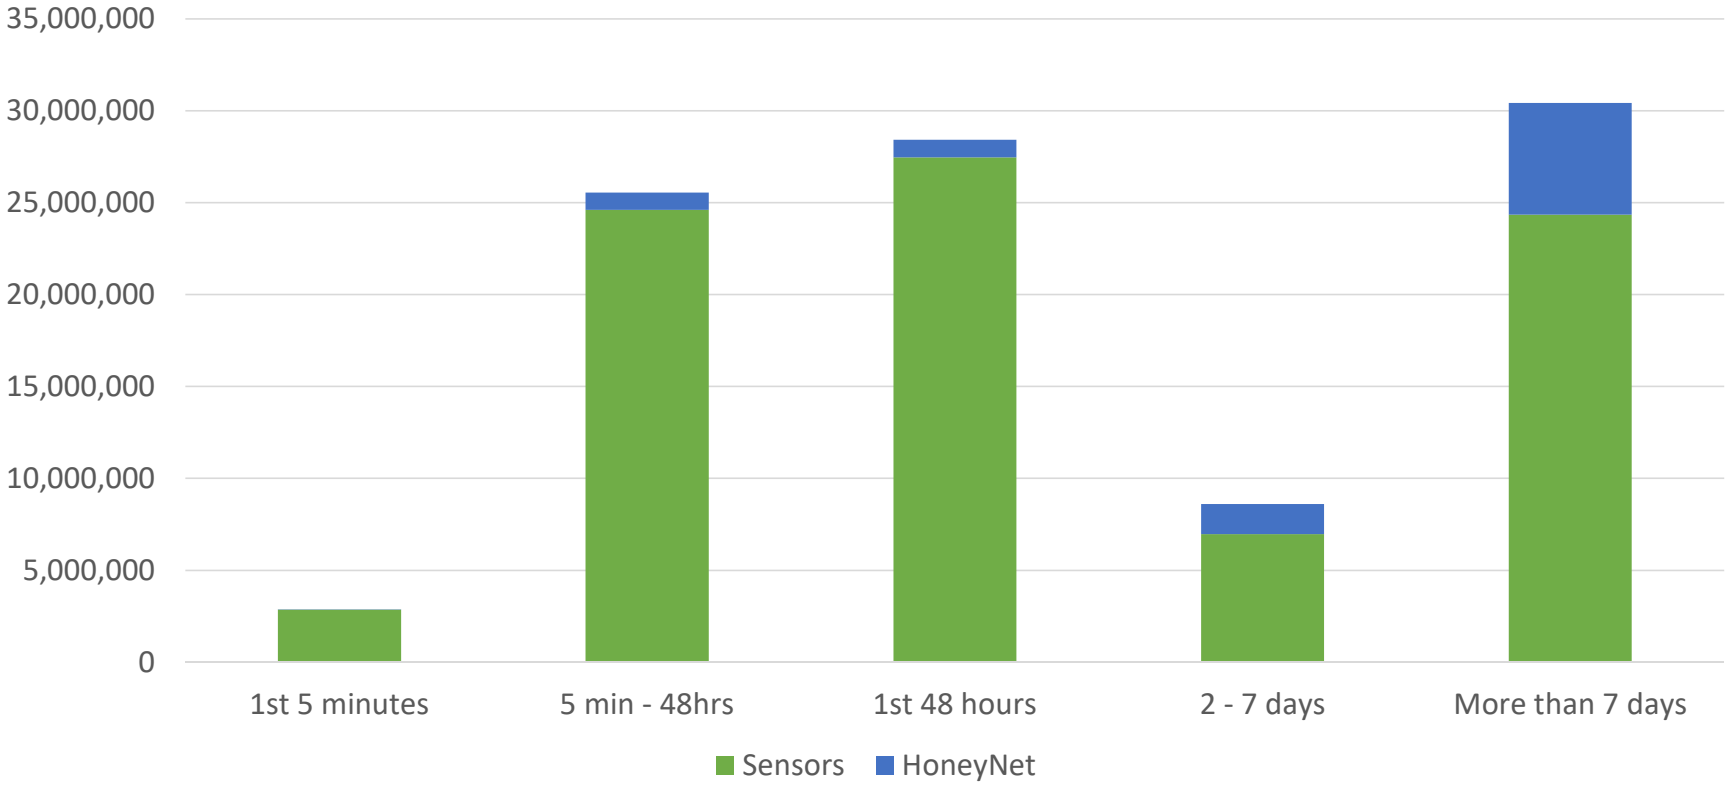

1,3 billion
Predictions from 1,599,296 Observations in our raw data set.

95,911,086
Predictions from 1,599,296 Observations in our cleaned, working set

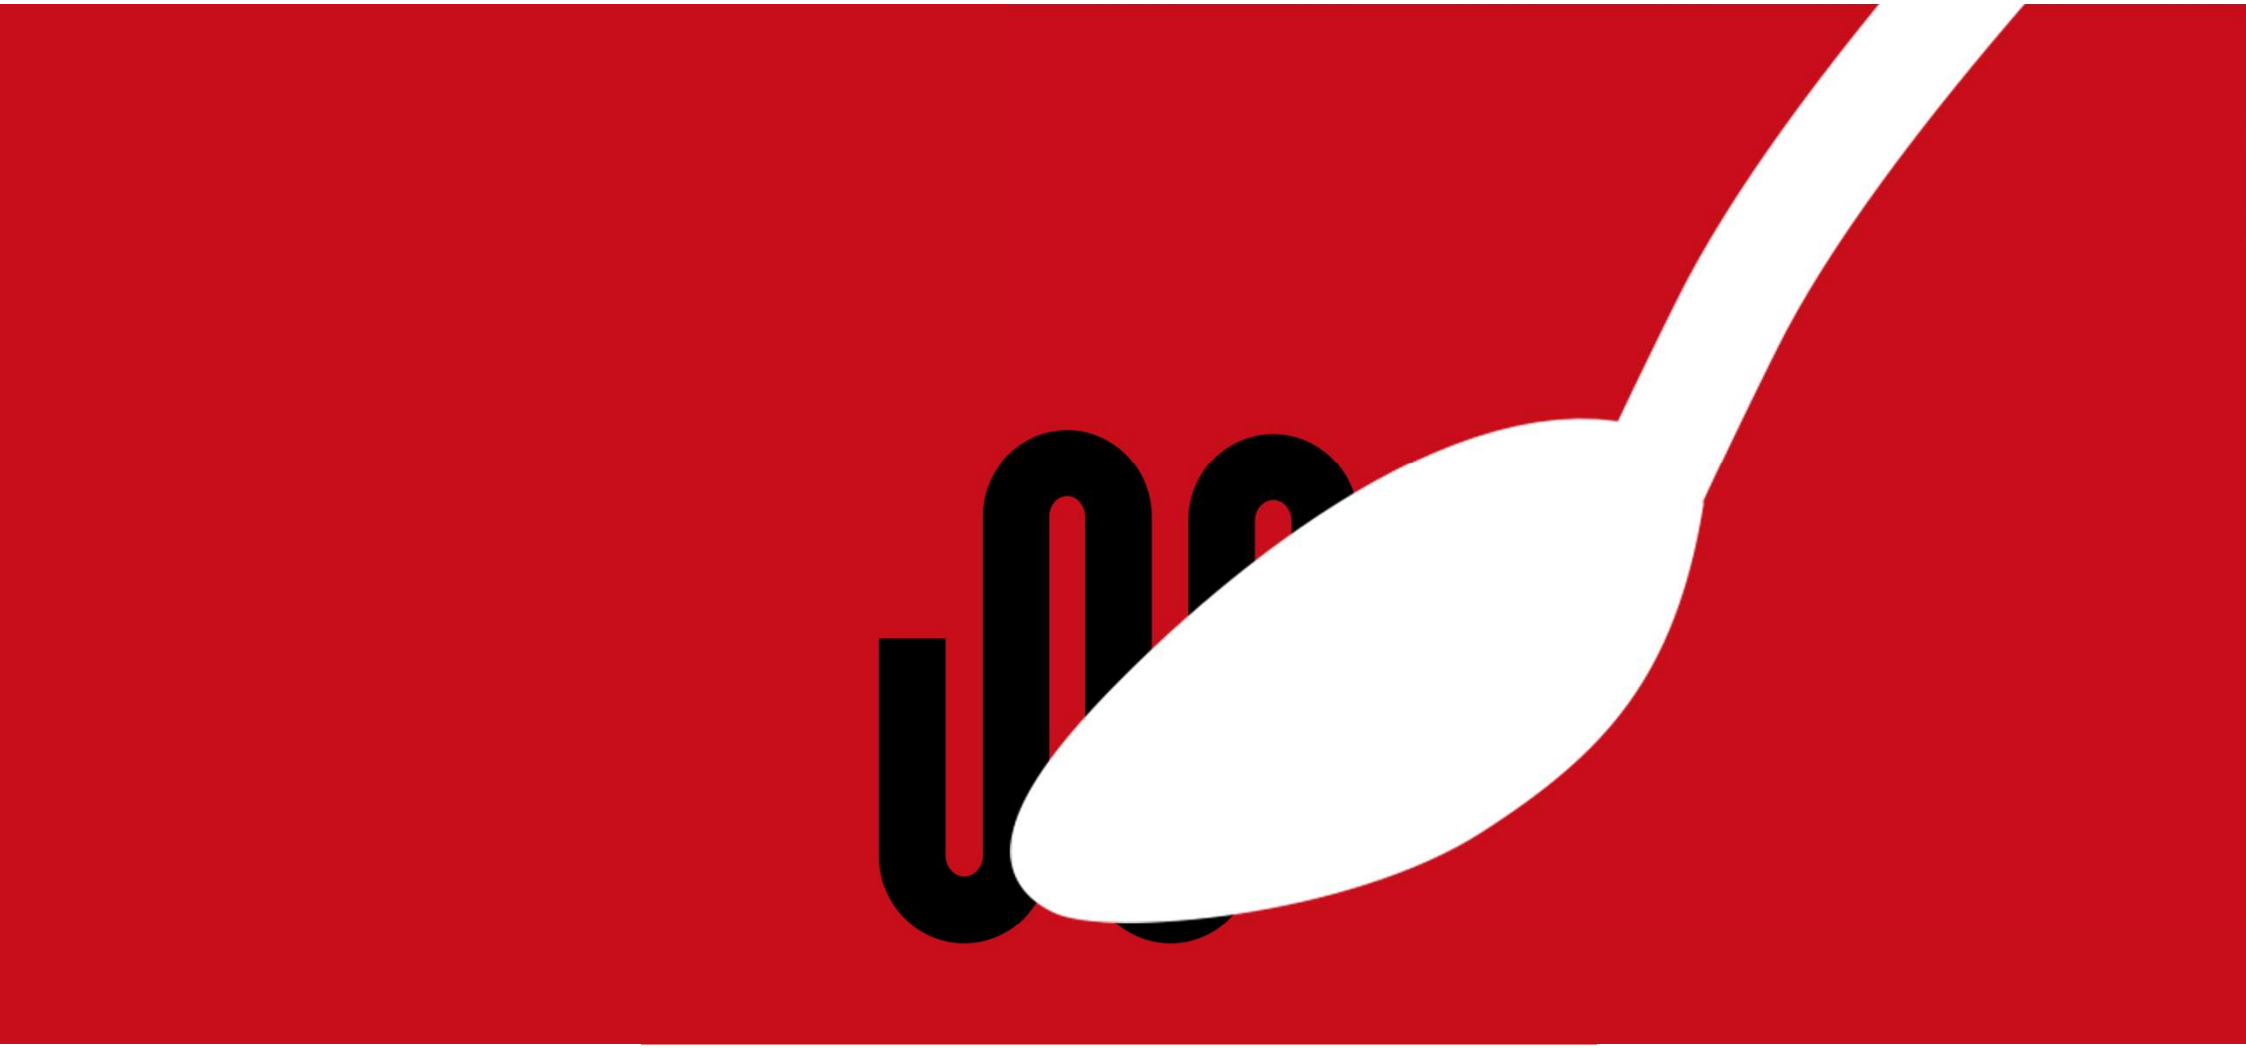

Key findings

Prediction Timeframes Distribution in Days

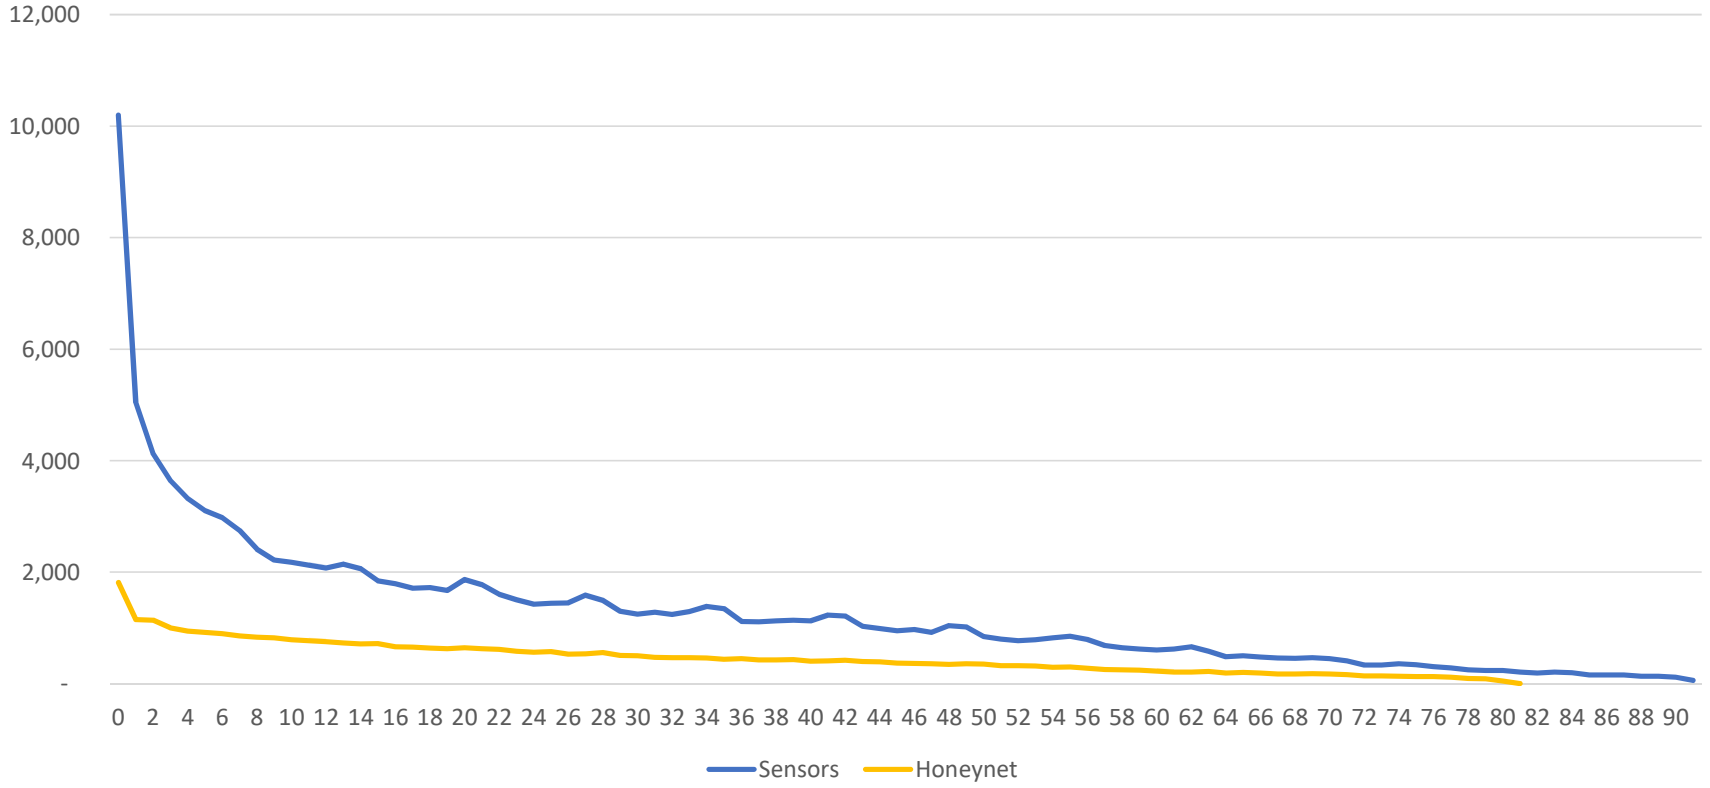

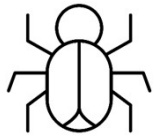

68%
of all Unique Predictions occurred within the ***1st 48hrs***

55%
of all Unique Predictions occurred within the ***1st 48hrs***

	External Threat Intelligence	Malicious Internet Activity	Malicious Web Activity	Suspicious Internet Activity	Suspicious Web Activity
External Threat Intelligence	**80.76%**	17.97%	0.67%	0.28%	0.32%
Malicious Internet Activity	29.85%	**68.16%**	0.34%	1.50%	0.15%
Malicious Web Activity	0.97%	0.97%	**97.81%**	0.00%	0.25%
Suspicious Internet Activity	2.07%	18.22%	0.00%	**78.29%**	1.42%
Suspicious Web Activity	0.40%	0.30%	0.09%	0.09%	**99.12%**

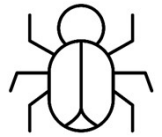

On average
87%
of all Predictions predicted a similar event

Summary of Diversity Events Predicted per IP

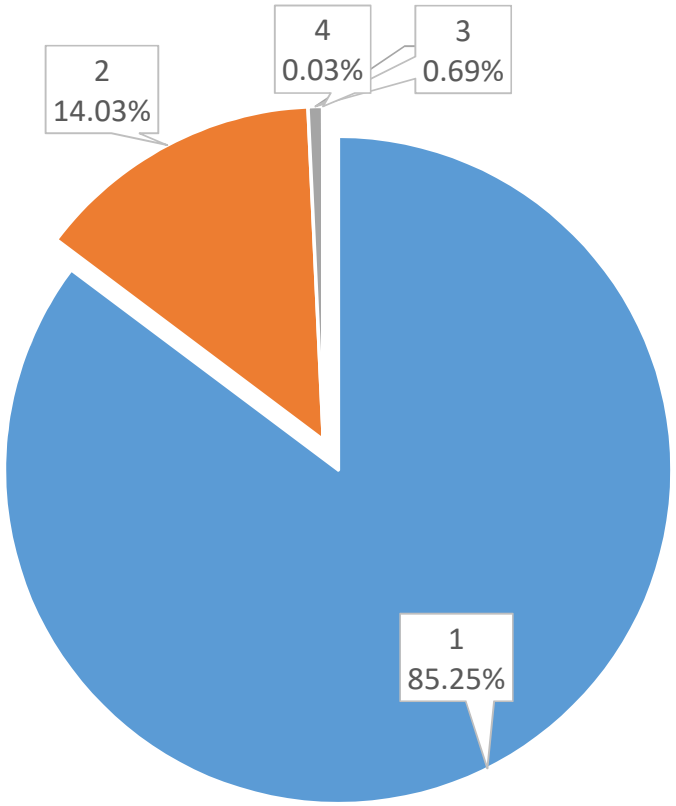

In
85%
of cases an IP that was observed acting suspiciously more than once, was still observed doing the same kind of thing.

37%

of IP addresses detected on our honeypots were detected again later by our sensors

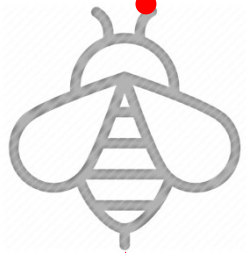

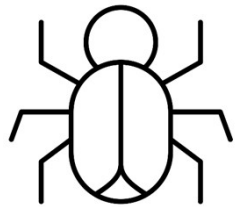

3%

of IP addresses detected on our sensors were detected again later by our honeynet

Observation

A suspicious security event detected and reported by a sensor
oIP is detected by ***Sensor***[x] at an ***Entity***[x] at ***Time***[x]

Prediction

A suspicious security event by an IP that serves as an early warning of another event by the same IP
oIP is observed by Sensor[**x**] at an Entity[**x**] at Time[x] before being observed by another Sensor[**y**] at Entity[**y**] at Time[**y**] within **Delta[t]**

Precision

Given that an IP is observed behaving suspiciously, with what *Precision* does it predict future suspicious behavior by the same IP
Pv = Meaningful Predictions / Observations

	PREDICTED = 1	PREDICTED = 0
SUSPICIOUS = 1	TRUE POSITIVE	FALSE POSITIVE
SUSPICIOUS = 0	FALSE NEGATIVE	~~TRUE NEGATIVE~~

TRUE POSITIVE

Joint probability, given Observations

$$pV = \frac{Unique\ Predictions}{Unique\ Observations}$$

Using maximum likelihood

FALSE POSITIVE

Joint probability, given Observations

$$C = \frac{Observations - Predictions}{Observations}$$

Precision.

$$P(correctly\ predicted = 1 \mid observed = 1)$$

Given that a specific IP is given to be acting suspiciously by a Threat Intelligence source, what is the ***probability*** that the IP will be observed acting suspiciously again later?

3.59%

Threat Intelligence Lab
Our T.I. petri dish environment

9.23%

Honeynet Lab
Our honeynet petri dish environment

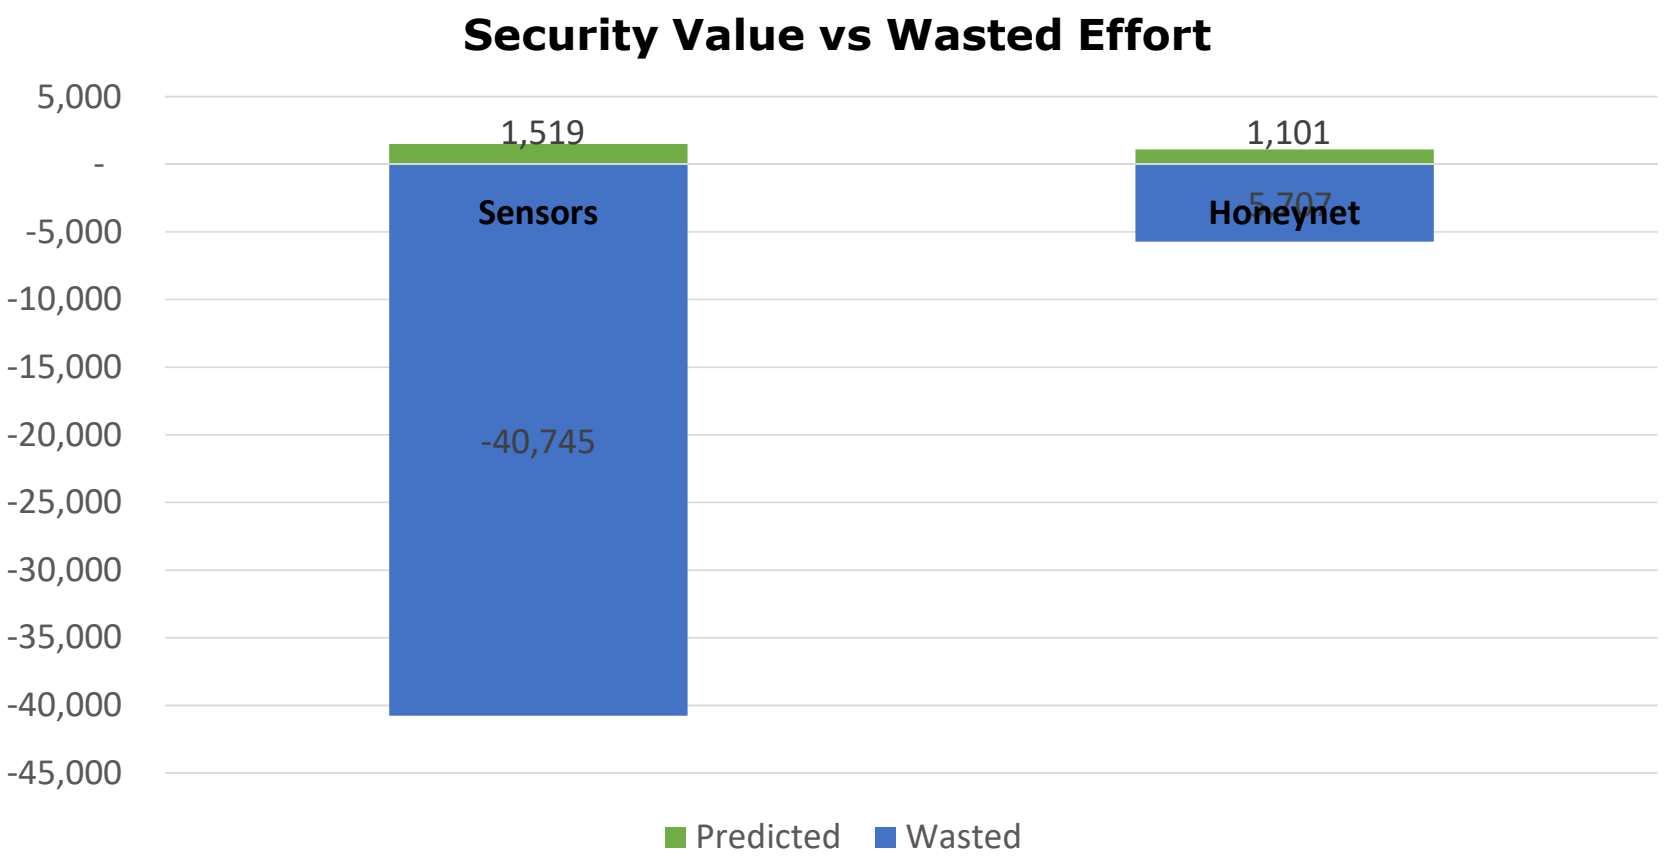

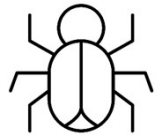

3.59%

precision, with normalized wastage of **0.81**.

9.23%

precision, with normalized wastage of **0.11**.

Normalised Overhead.

Given that an IoA False Positive represents wasted work, no matter how small, what is the relative cost of Threat Intelligence, normalized for comparison.

0.81

Threat Intelligence Lab
Our T.I. petri dish environment

0.11

Honeynet Lab
Our honeynet petri dish environment

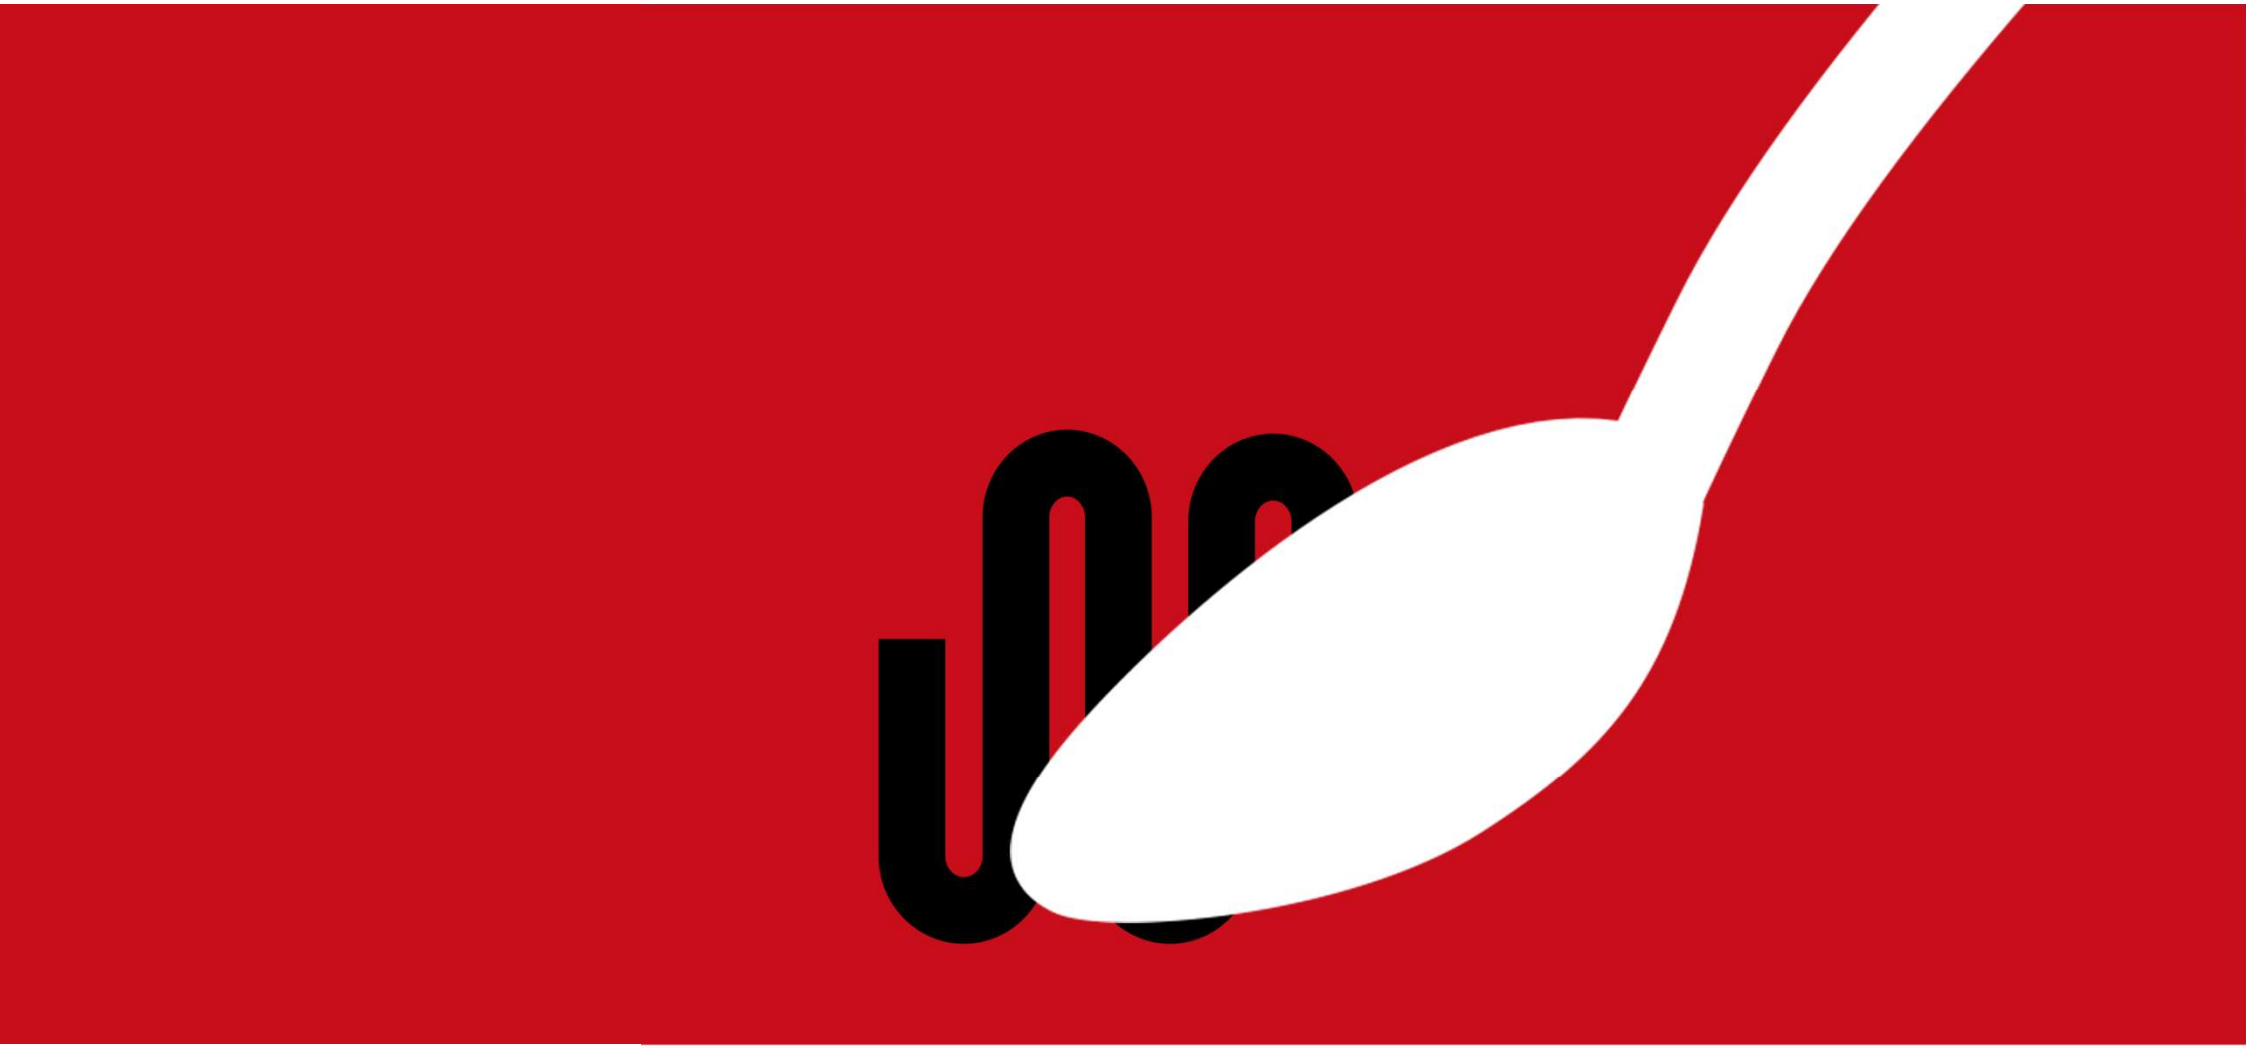

Additional Observations

The estimated amount of time, in man-days, over the 90-day experiment period, that would be required to deal with all the False Positives generated by our sensor feed.	**48.6 DAYS** **Threat Intelligence Lab** Our T.I. petri dish environment
The estimated amount of time, in man-days, over the 90-day experiment period, that would be required to deal with all the False Positives generated by our honeynet feed.	**8.26 DAYS** **Honeynet Lab** Our honeynet petri dish environment

black hat EUROPE 2018

SECUREDATA
TRUSTED CYBERSECURITY EXPERTS

Don't eat Spaghetti with a spoon

Honeynet Effectiveness

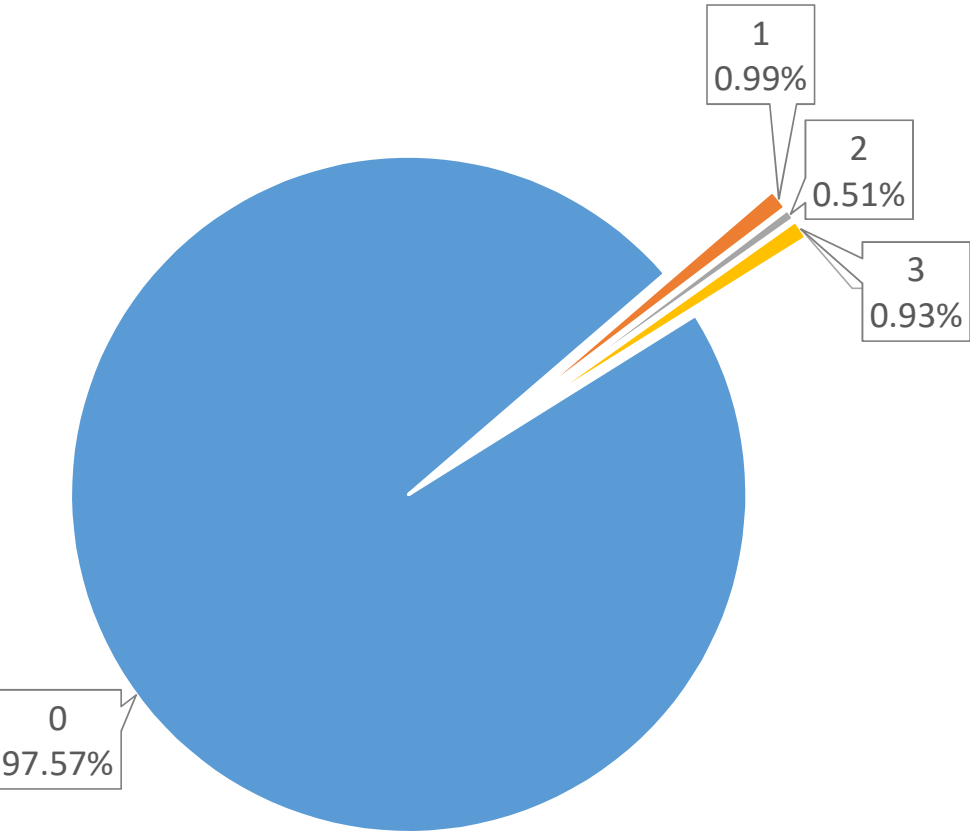

Honeypot Correlation Summary

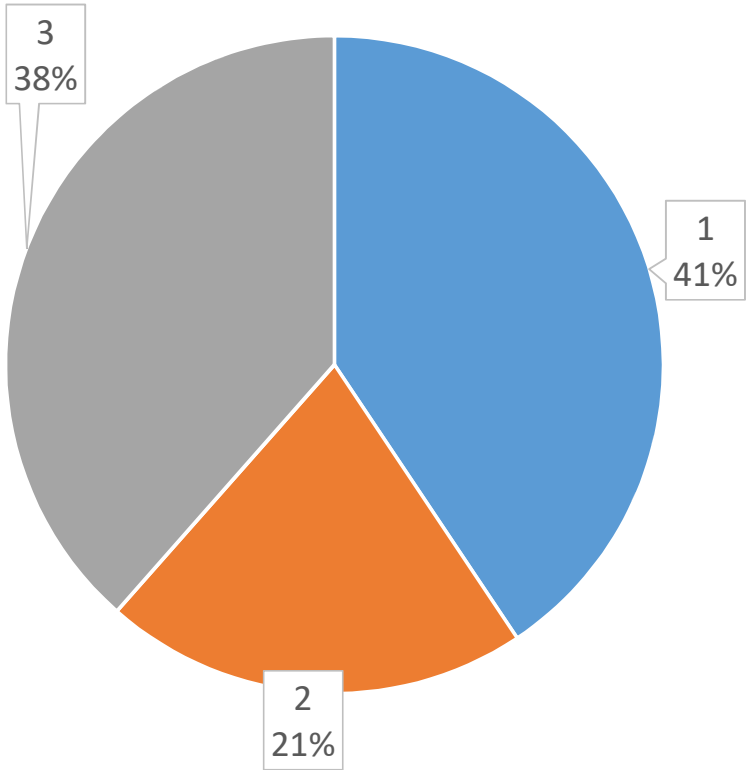

Only **2.5%** of IPs observed in this experiment were observed by our honeynet.

Of those, **41%** were only observed by only one honeypot.

Honeynet Scaling Behaviour

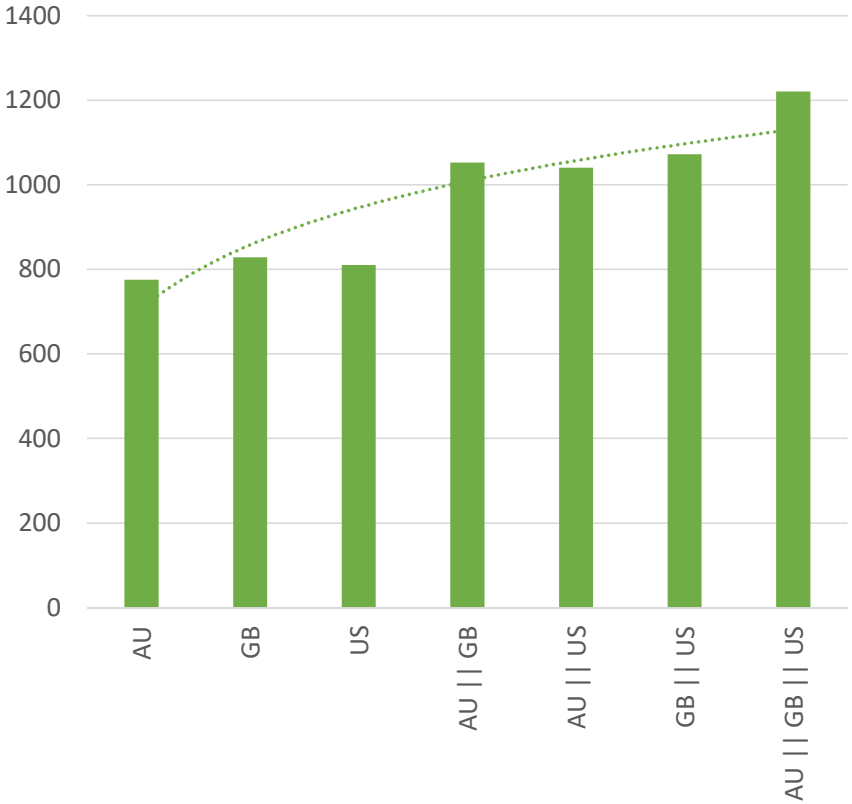

Improvement with Scale?

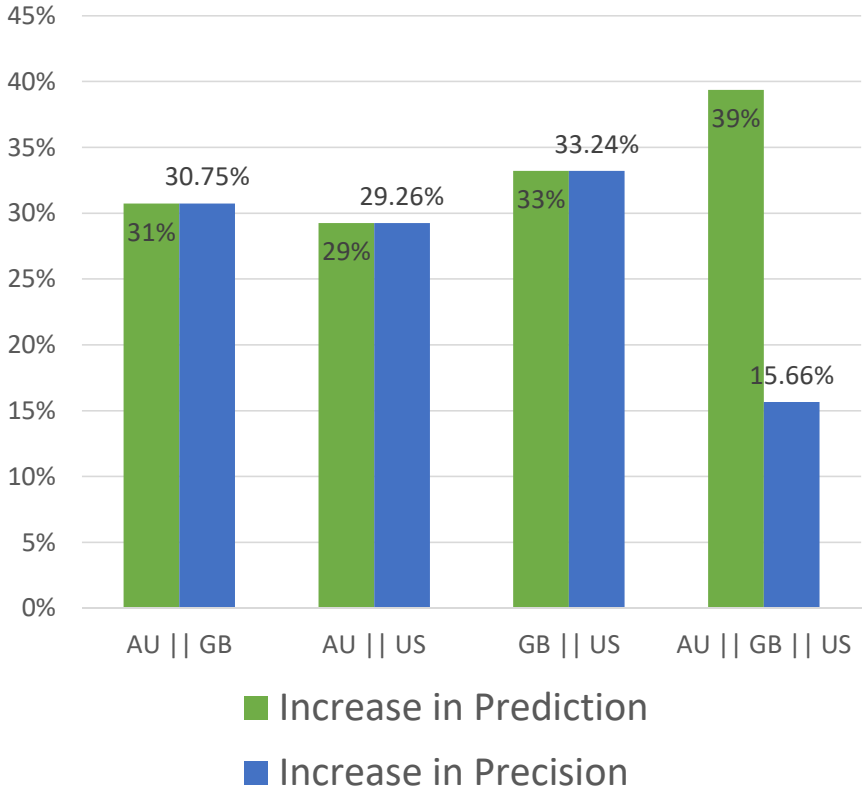

Effectiveness grows with additional honeypots

Increase in Prediction initially at **~30%** on average

But the Increase in Precision appears to drop of quite quickly

black hat EUROPE 2018

SECURE DATA
TRUSTED CYBERSECURITY EXPERTS

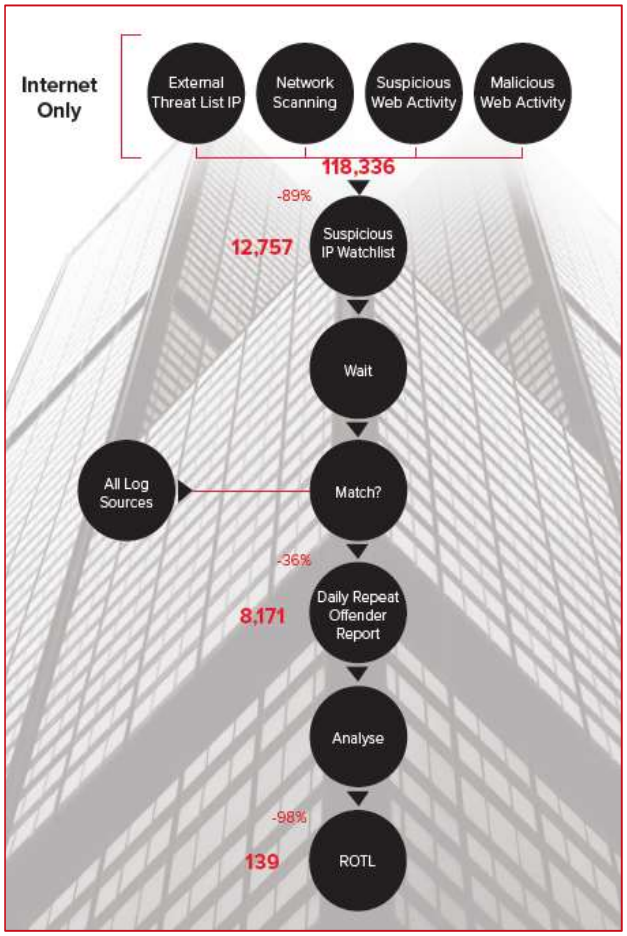

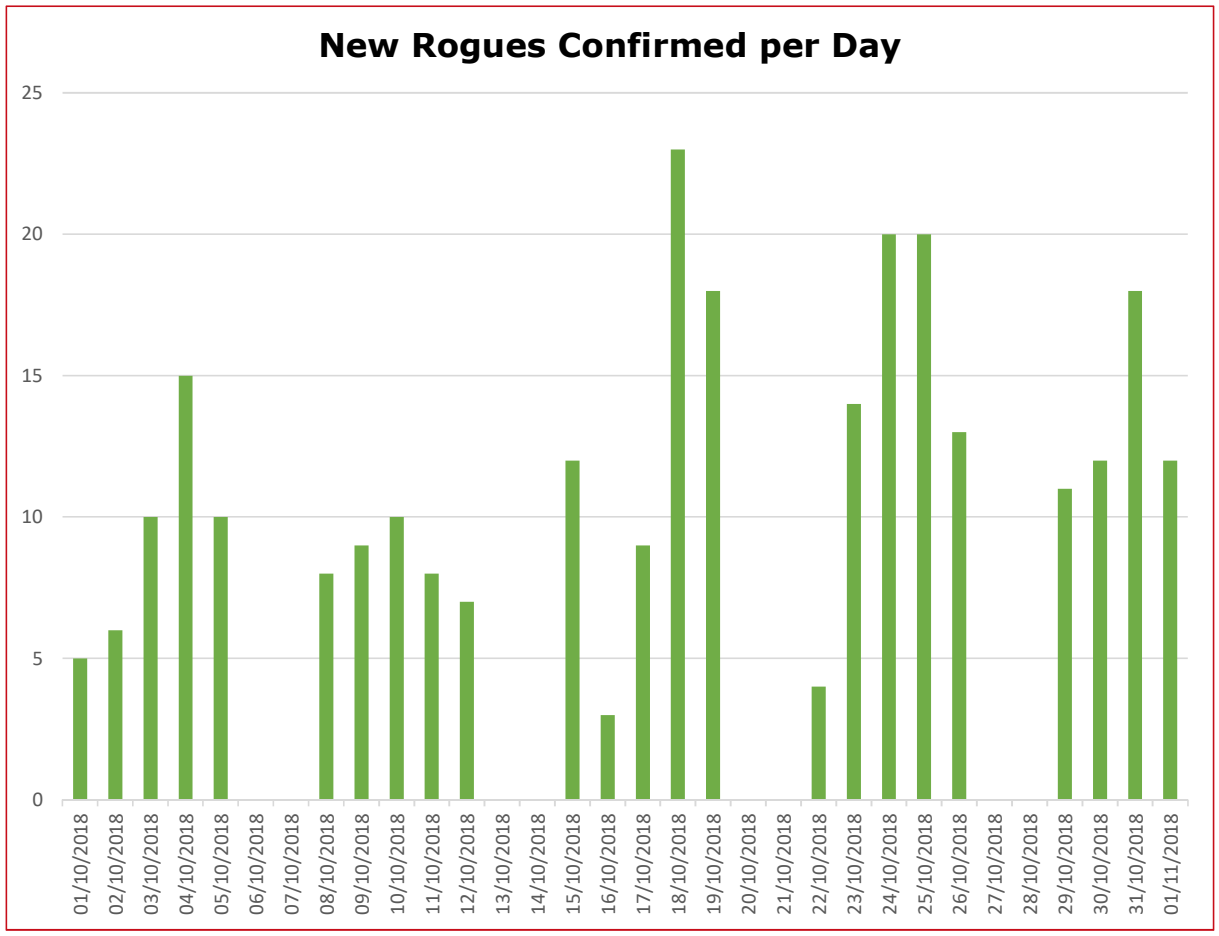

171
rogues confirmed by our analysts as *suspicious* and *persistent* during the period.

Precision on Rogue List.

$$P(correctly\ predicted = 1 \mid observed = 1)$$

Given that a specific IP is given to be acting suspiciously by a Threat Intelligence source, what is the probability that the IP will finally be confirmed by our analysts as a **rogue**

0.25%

Threat Intelligence Lab
Our T.I. petri dish environment

0.84%

Honeynet Lab
Our honeynet petri dish environment

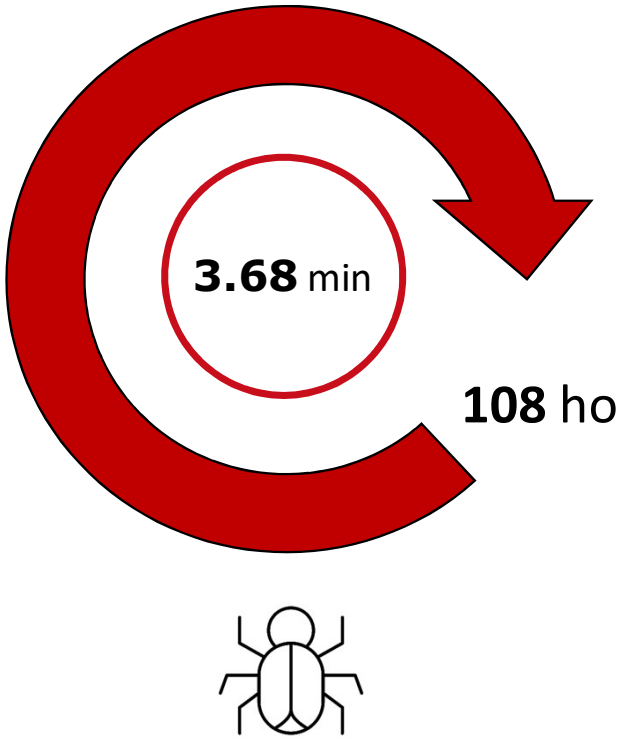

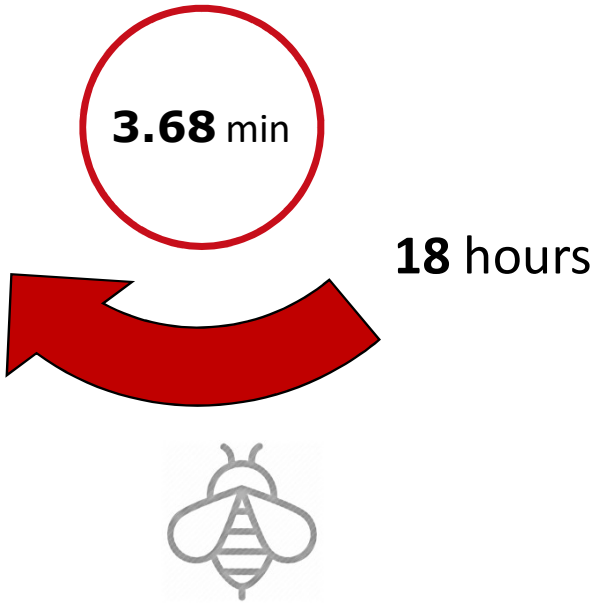

Based on an internal analysis of our own processes we estimate it takes an analyst **3.68** minutes to process a suspect IP.

Applying this to the number of False Positives involved we can estimate that a manual process of confirming the false positives from our Sensors amounted to **108 hours** of wasted effort

Performance against Rogue List

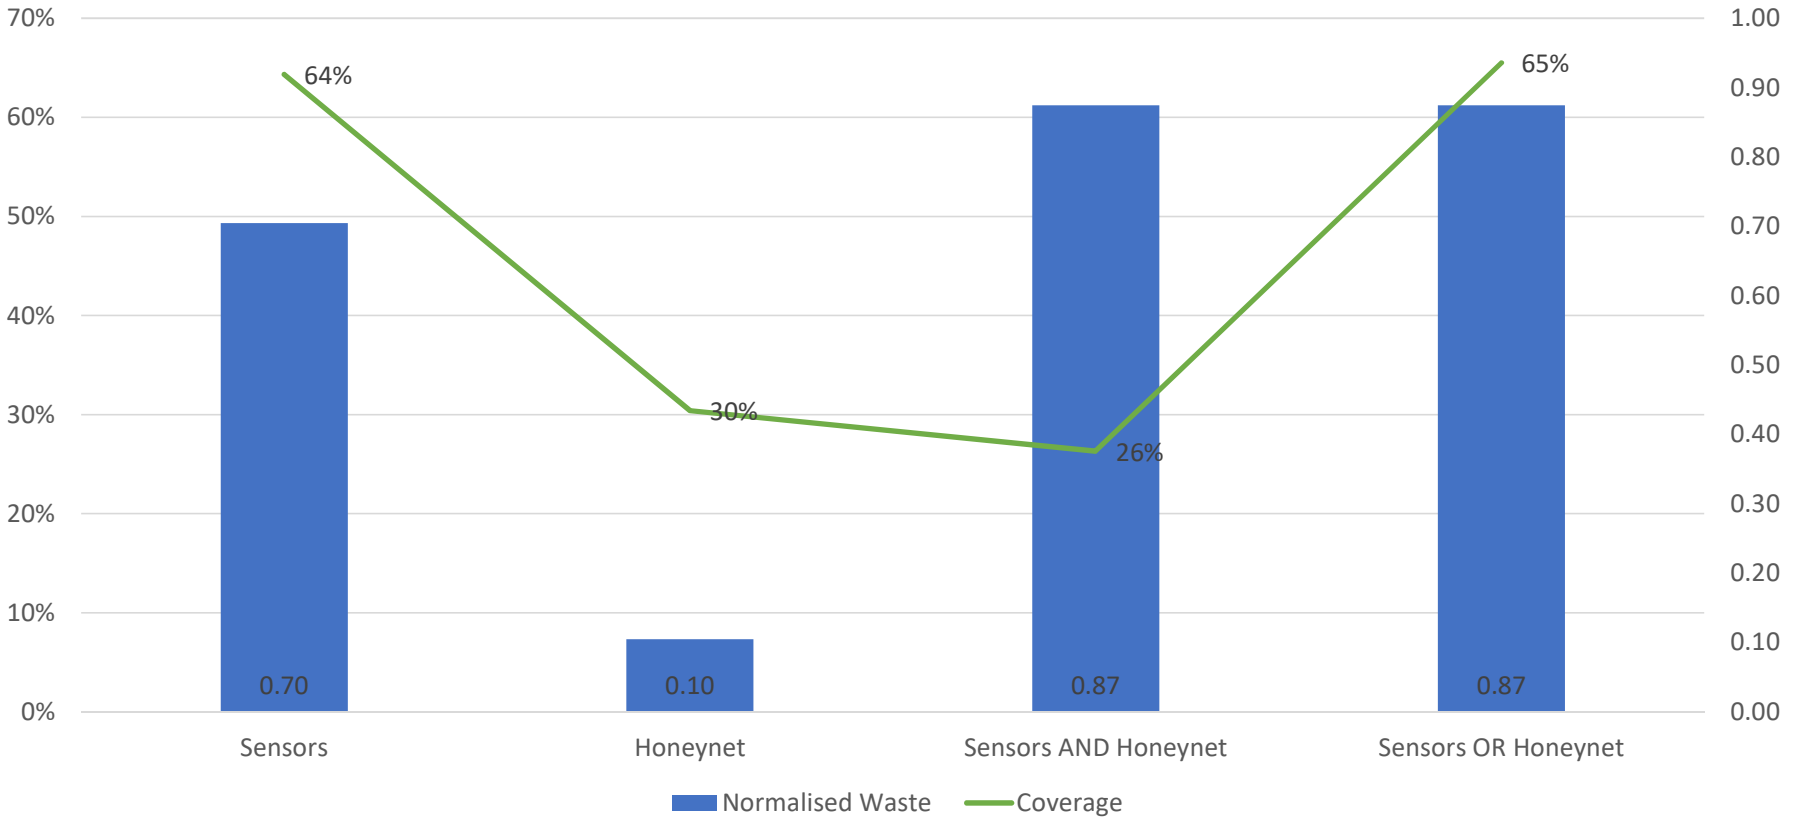

Sensors are **2x** as effective as the *Honeynet* at predicting rogues, but at **7x** the cost in wasted effort.

Combining the Sensors and Honeynet improves Coverage, but also the effort.

black hat EUROPE 2018

SECUREDATA
TRUSTED CYBERSECURITY EXPERTS

Don't eat Spaghetti with a spoon

Security Value vs Wasted Effort

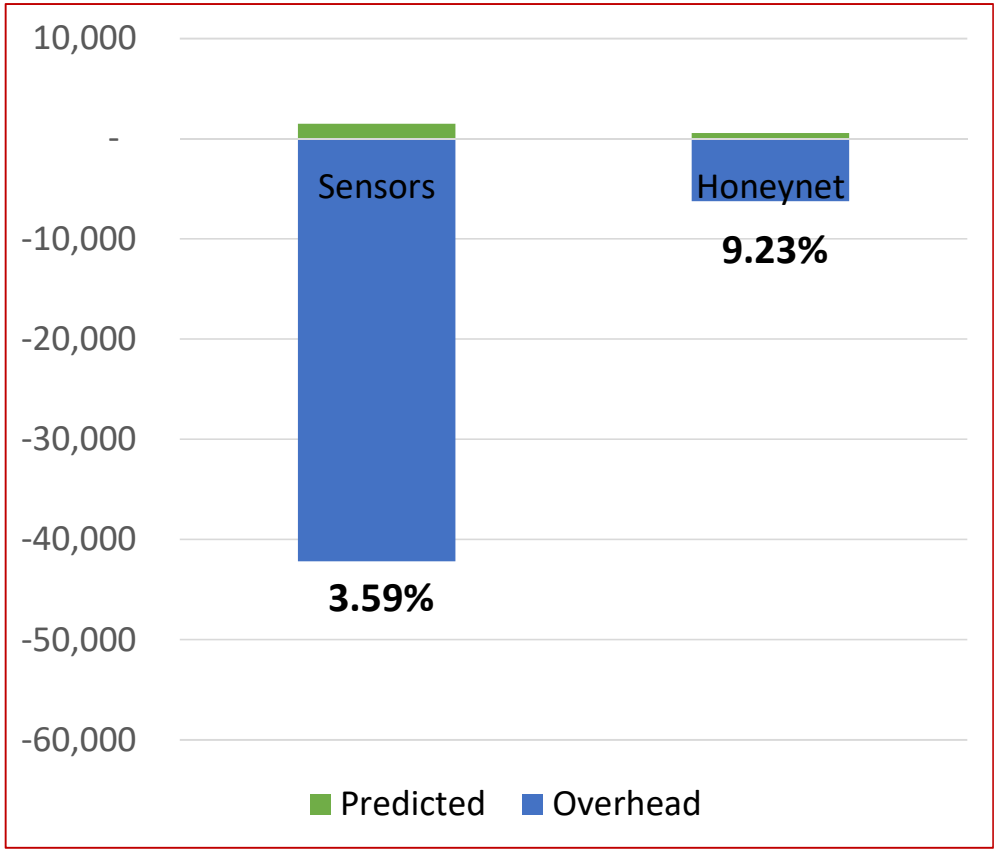

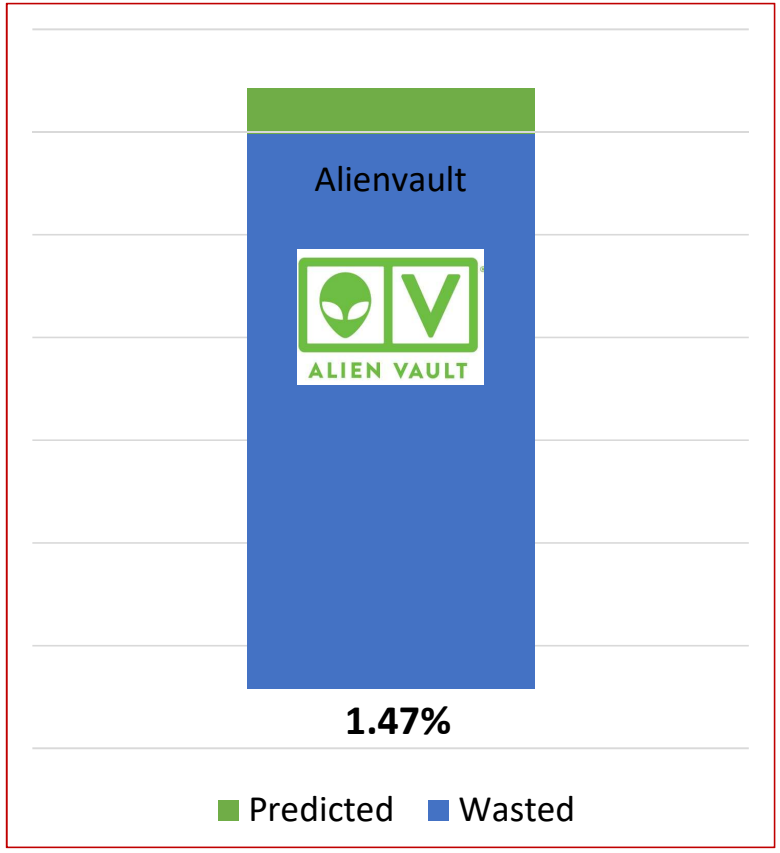

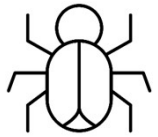

Threat List predicted **three times** as much as our sensors, but at of **39%** more wasted effort.

Threat List predicted **seven times** as much as our honeynet, but at of **9x** more wasted effort.

Commercial Threat List Sample.

471
Source IP overlap with our Sensors

4
Source IP overlap with our Rogues List

17
Source IP overlap with our Honeynet

110 & **6**
Source IP overlap with our Sensor & Honeynet Predictions

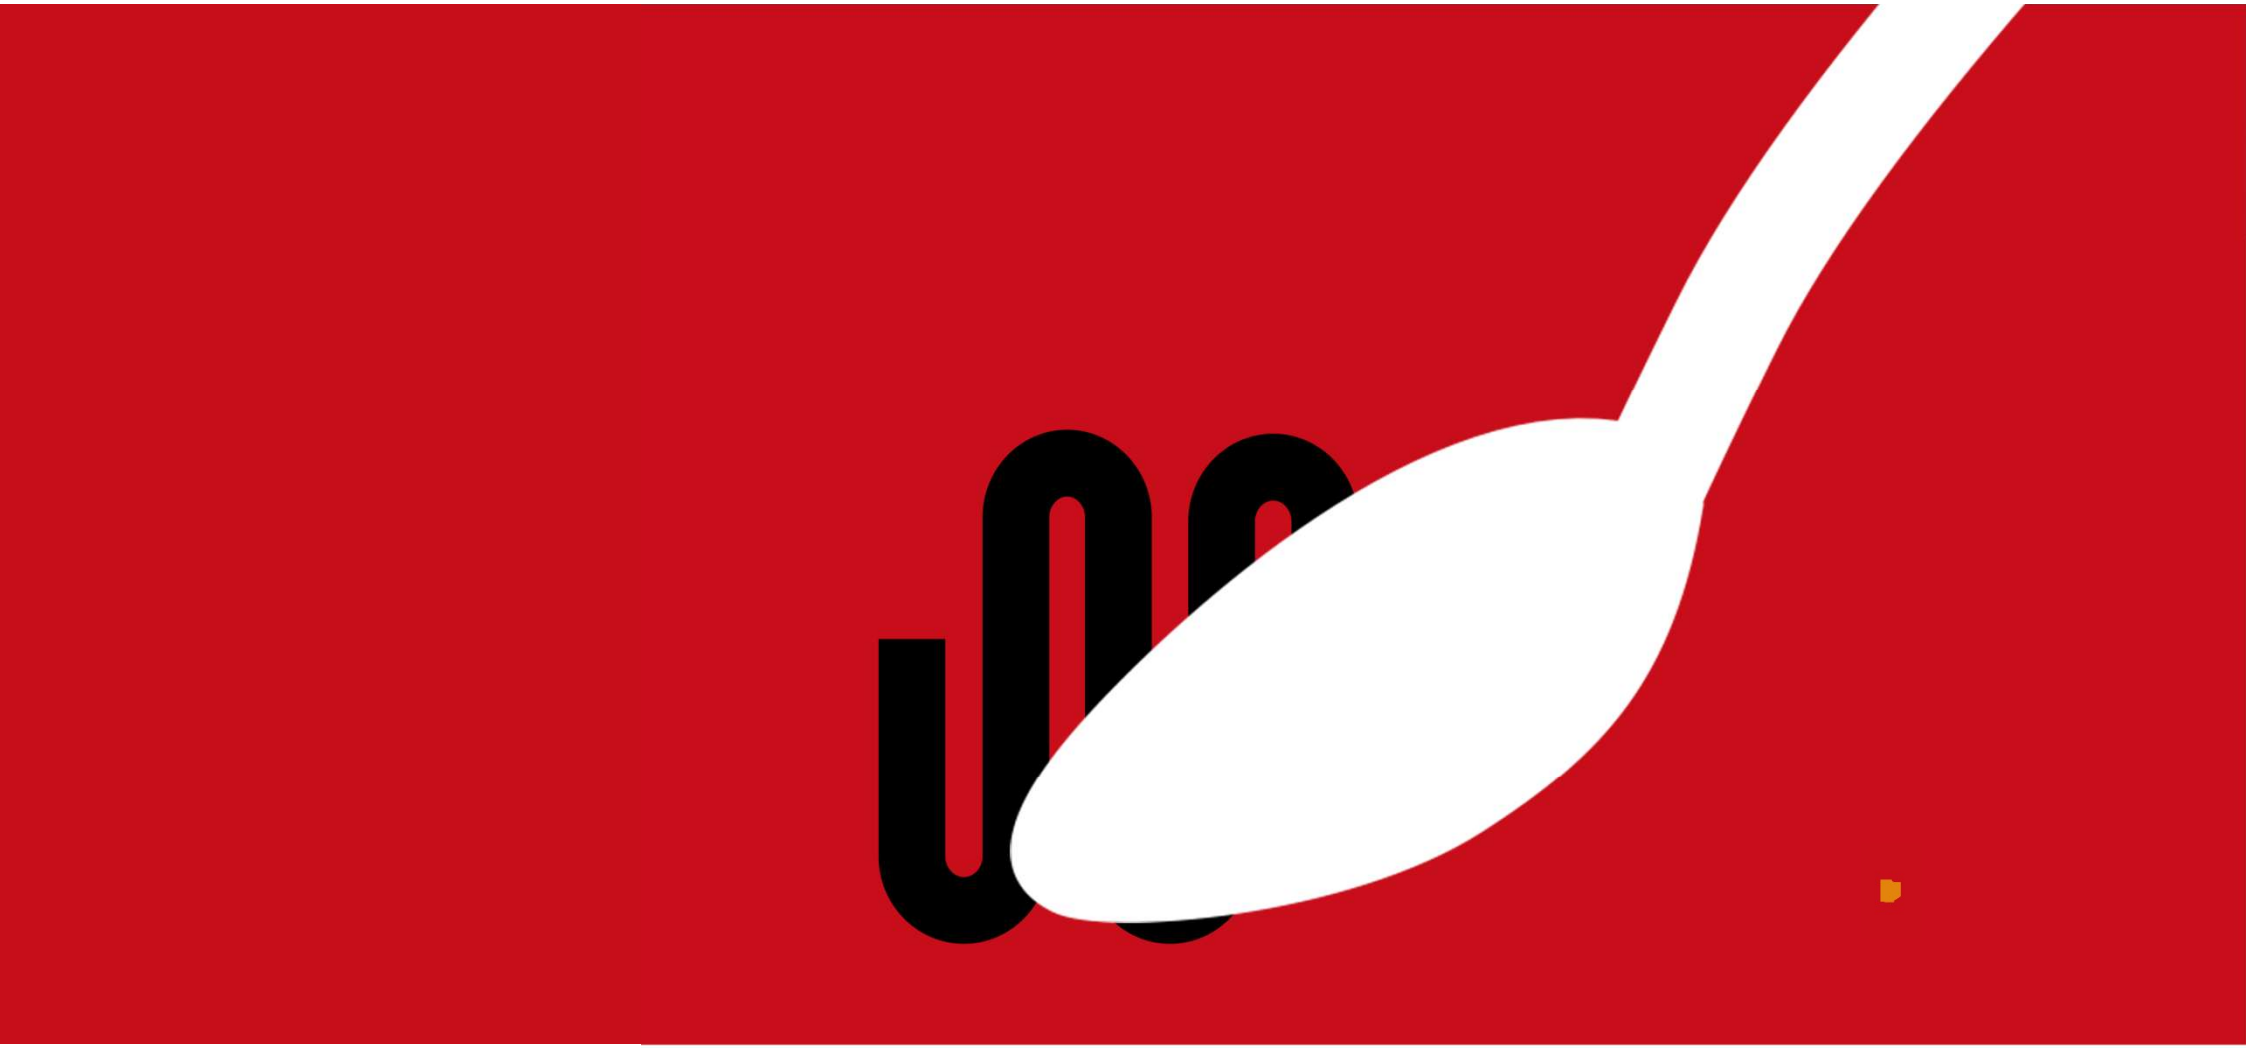

The digestive

A question of philosophy.

All forms of intelligence-led security suffer from the same tension between three factors – False Positives, Limited Resources & Unknown Unknowns.

At what levels do these come into balance and, given that we will never know the Unknown Unknowns, is there any real logic in pursuing it?

Would our limited resources not be better spent in proactively engineering robust systems?

This dilemma holds not only for Threat Intelligence, but also for Threat Detection, Bug Hunting, Vulnerability Scanning and other domains.

T: +44 (0)1622 723400 | **E:** info@secdata.com | **W:** www.secdata.com

Parting thoughts.

So what to make of all of this…?

Honeypot appear much more effective
Our simple Honeynet faired twice as well as our Threat Intelligence petri dish, and at a quarter the 'effort'

But all the list tested basically suck
Less than 10% of all the IPs we produced as 'intelligence' were involved in other suspicious behavior. For actual Threat Lists and for all practical purposes, the performance was much worse than that.

This was just an experiment
These are the results of a staged and limited experiment, not an evaluation of any commercial project

More work is needed to test these results with actual Threat Lists
This work arguably offers more questions than answers.

SECUREDATA
TRUSTED CYBERSECURITY EXPERTS

https://github.com/SecureDataLabs/BlackHat-EU-2018