Battle between hackers and machine learning

Current status and trends

Mikhail Kader
Distinguished System Engineer

Alexey Lukatsky
Business Development Manager

July, 5 2018

Our agenda

AI for cyber security

Hackers / threats trends

Expectations

HACKER / THREATS

Current status

Process of Attacks

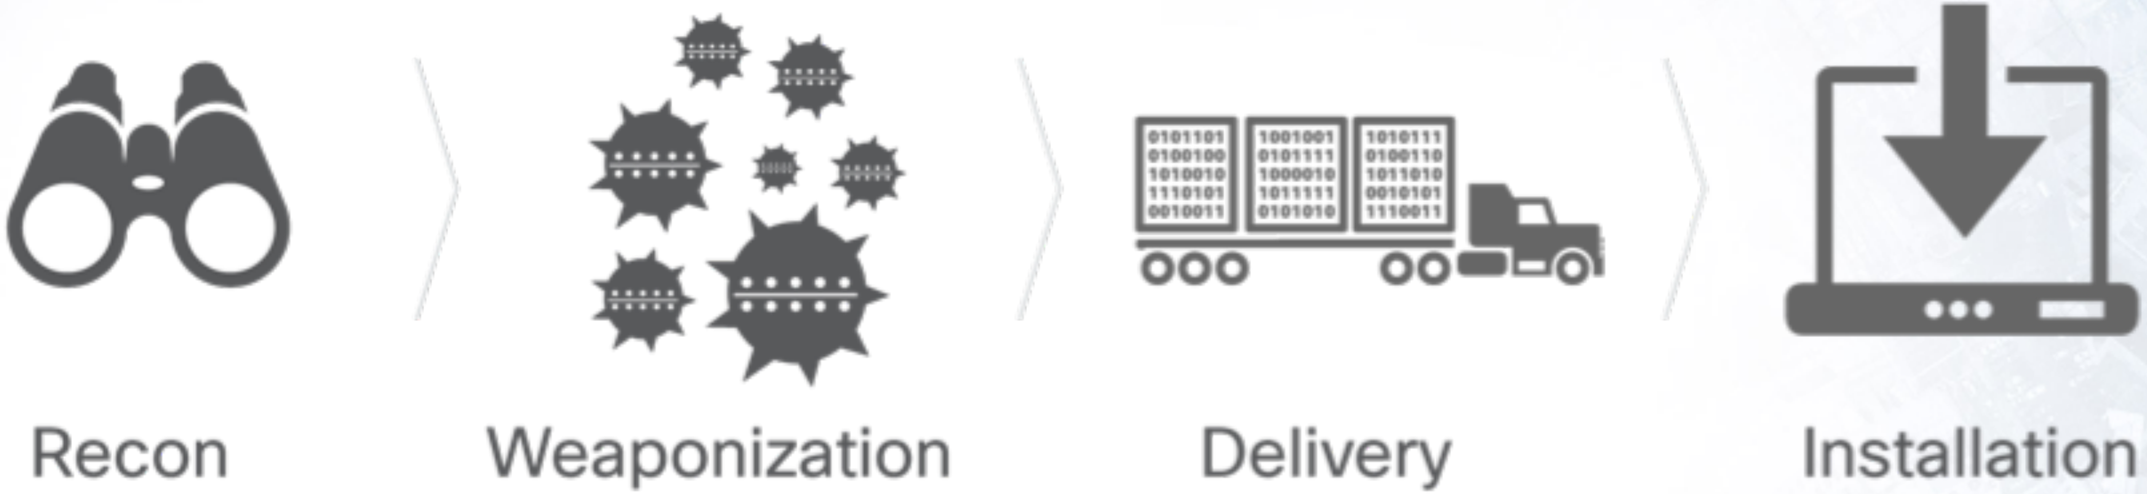

Research, identify and select targets

Pair remote access malware with exploits

Deliver cyberweapons by email, website and attachments

Install payloads to gain persistent access

High Severity Vulnerabilities and Patch Management

High severity is driven by headlines

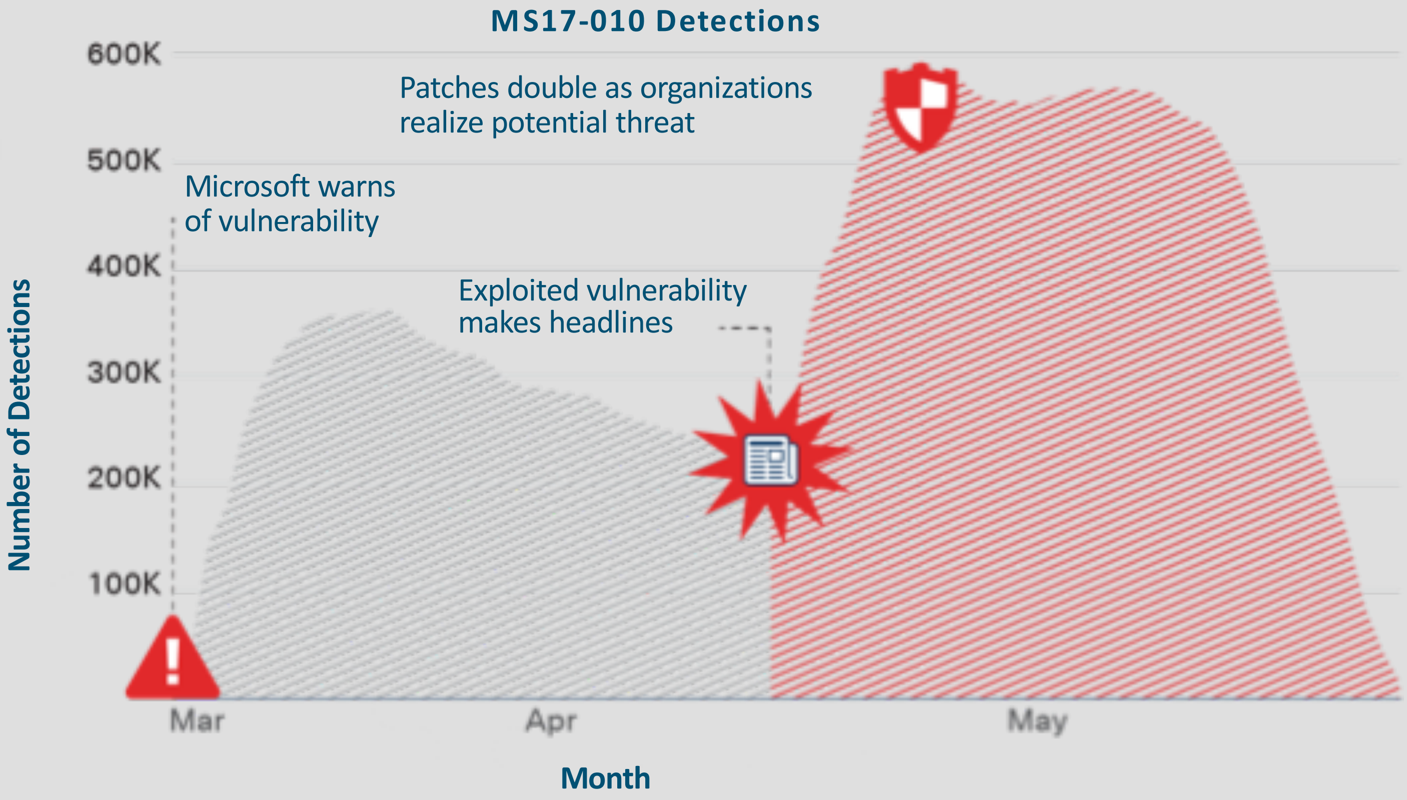

Source: Qualys

We need a better way to improve patch management processes. Can you patch all of your systems, for example, ICS?

Network-based Ransomware

WannaCry and Nyetya: rapid-moving, self-propagating network-based attacks

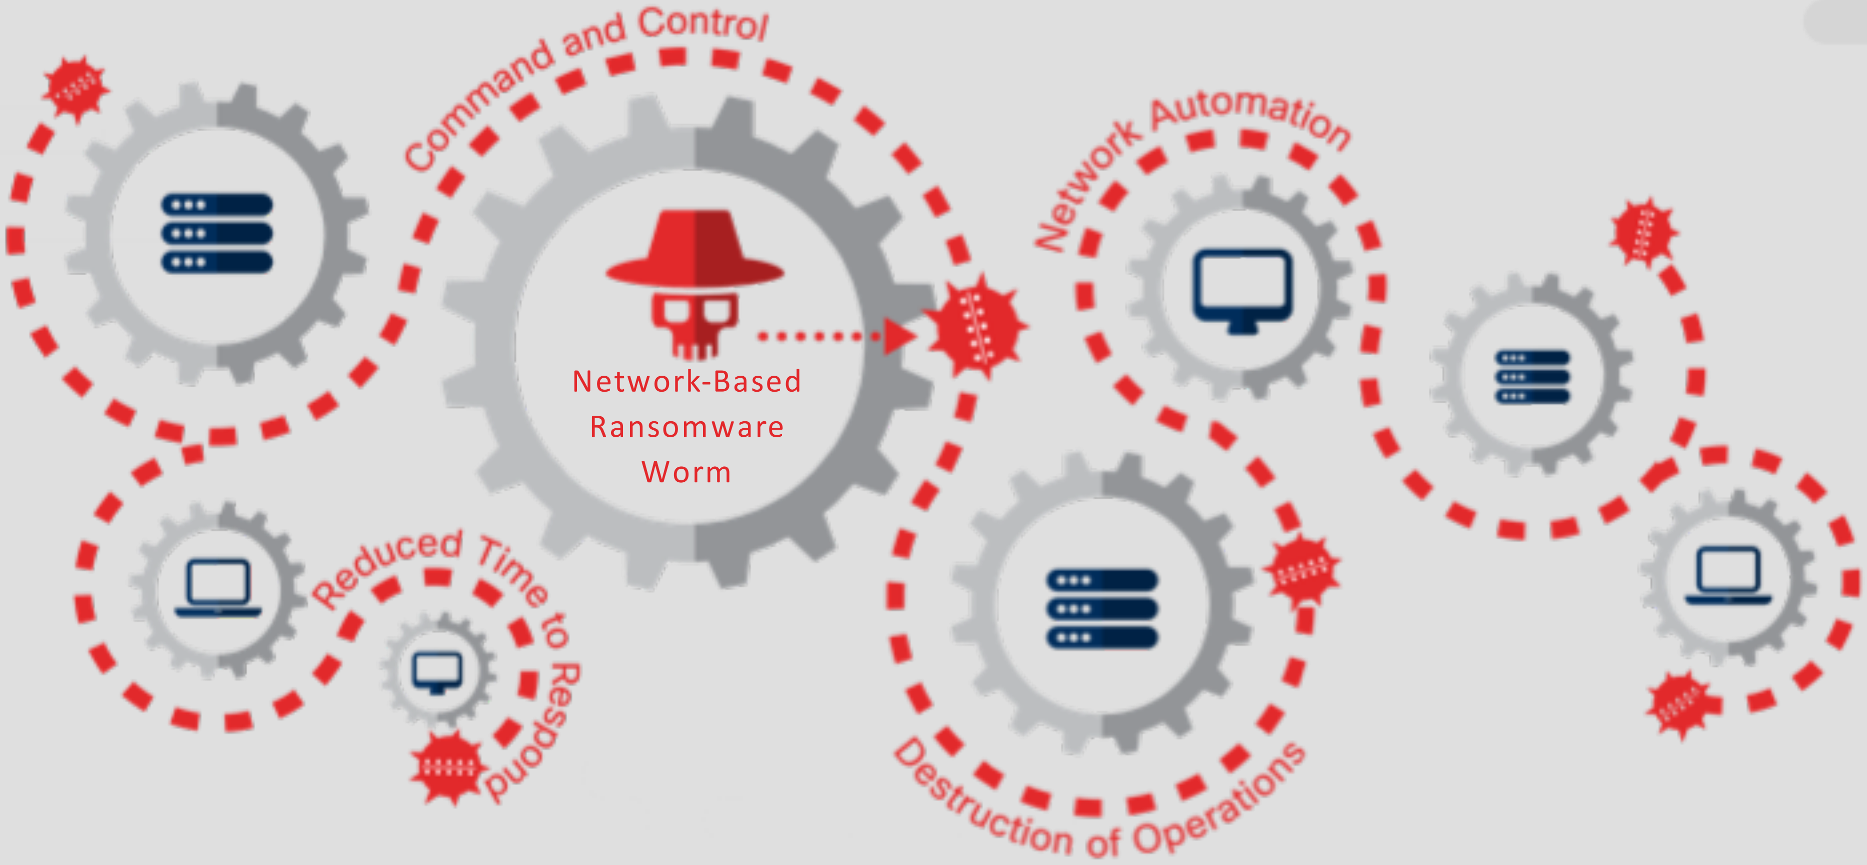

With active, unpatched machines, these automated worms will attack again. Have you secured your network?

Adware and Malvertising Shift Into High Gear

Malvertising

Using brokers (gates) to increase speed and agility

Switching quickly between servers without changing redirection

ShadowGate: a cost-effective campaign

Adware

75%

of organizations investigated had adware infections

Spam Attacks: Snowshoe and Hailstorm

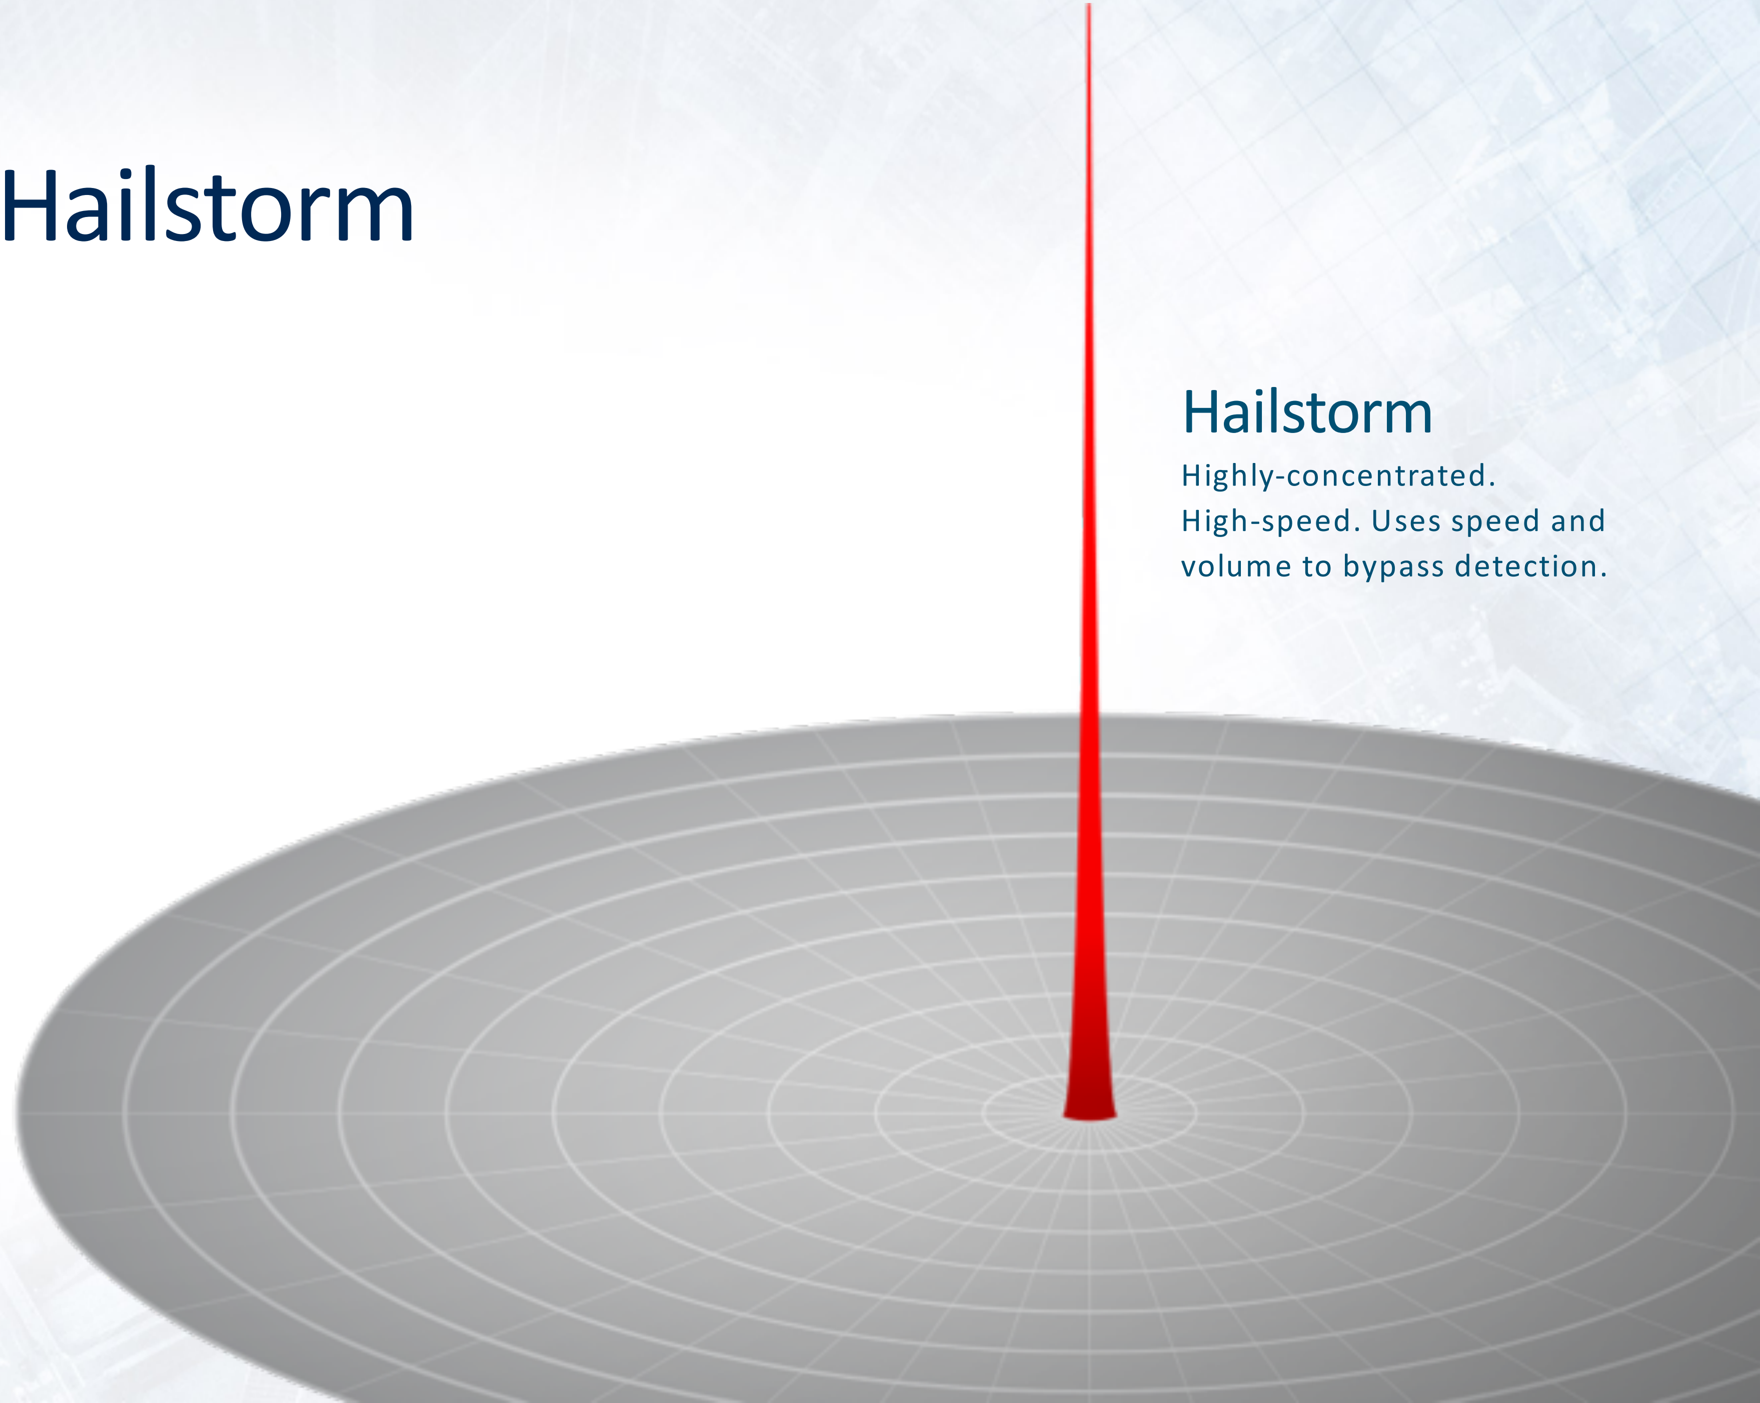

Hailstorm

Highly-concentrated. High-speed. Uses speed and volume to bypass detection.

Snowshoe

Uses various IP address. Hides from detection with low volume.

TTE: Time To Evolve

Malware Families Behaving Badly; Closing Window of Opportunity

File Types

Attackers cycle through various file types such as .zip,
.exe, .js, .docm, .wsf

Delivery Mechanisms

Attackers deploy through both web and emails

Speed of Evolution

Attackers quickly evolve and generate new files as the old ones become less effective

TTD

Defenders need to reduce TTD to force attackers' hands

TTE: File Delivery Mechanisms (Locky)

Adversaries Shift Vectors Often and Quickly to Evade Detection

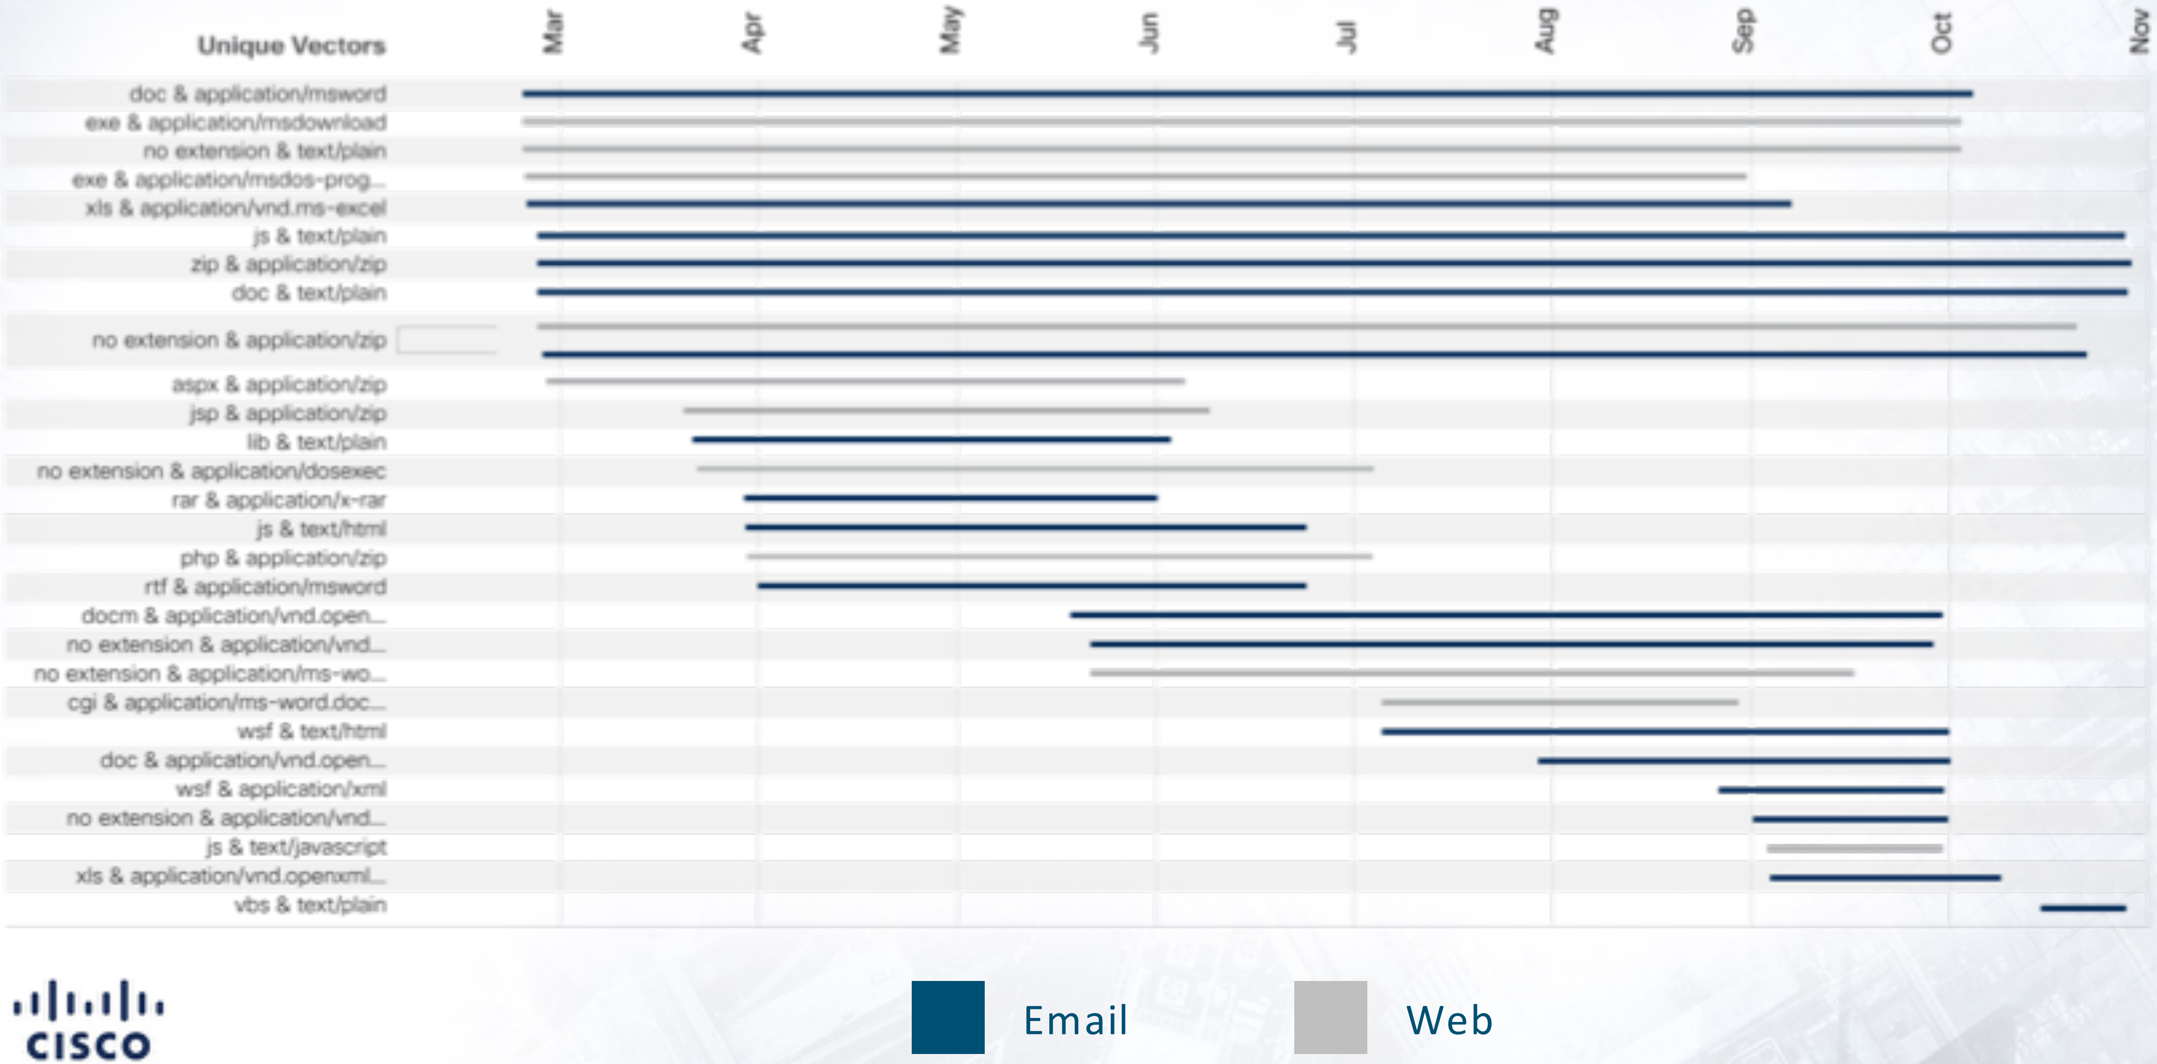

TTE: Unique Hashes vs. TTD (Locky)

Reducing TTD Forces Adversaries to Speed Up Their Effort Just to Keep Up

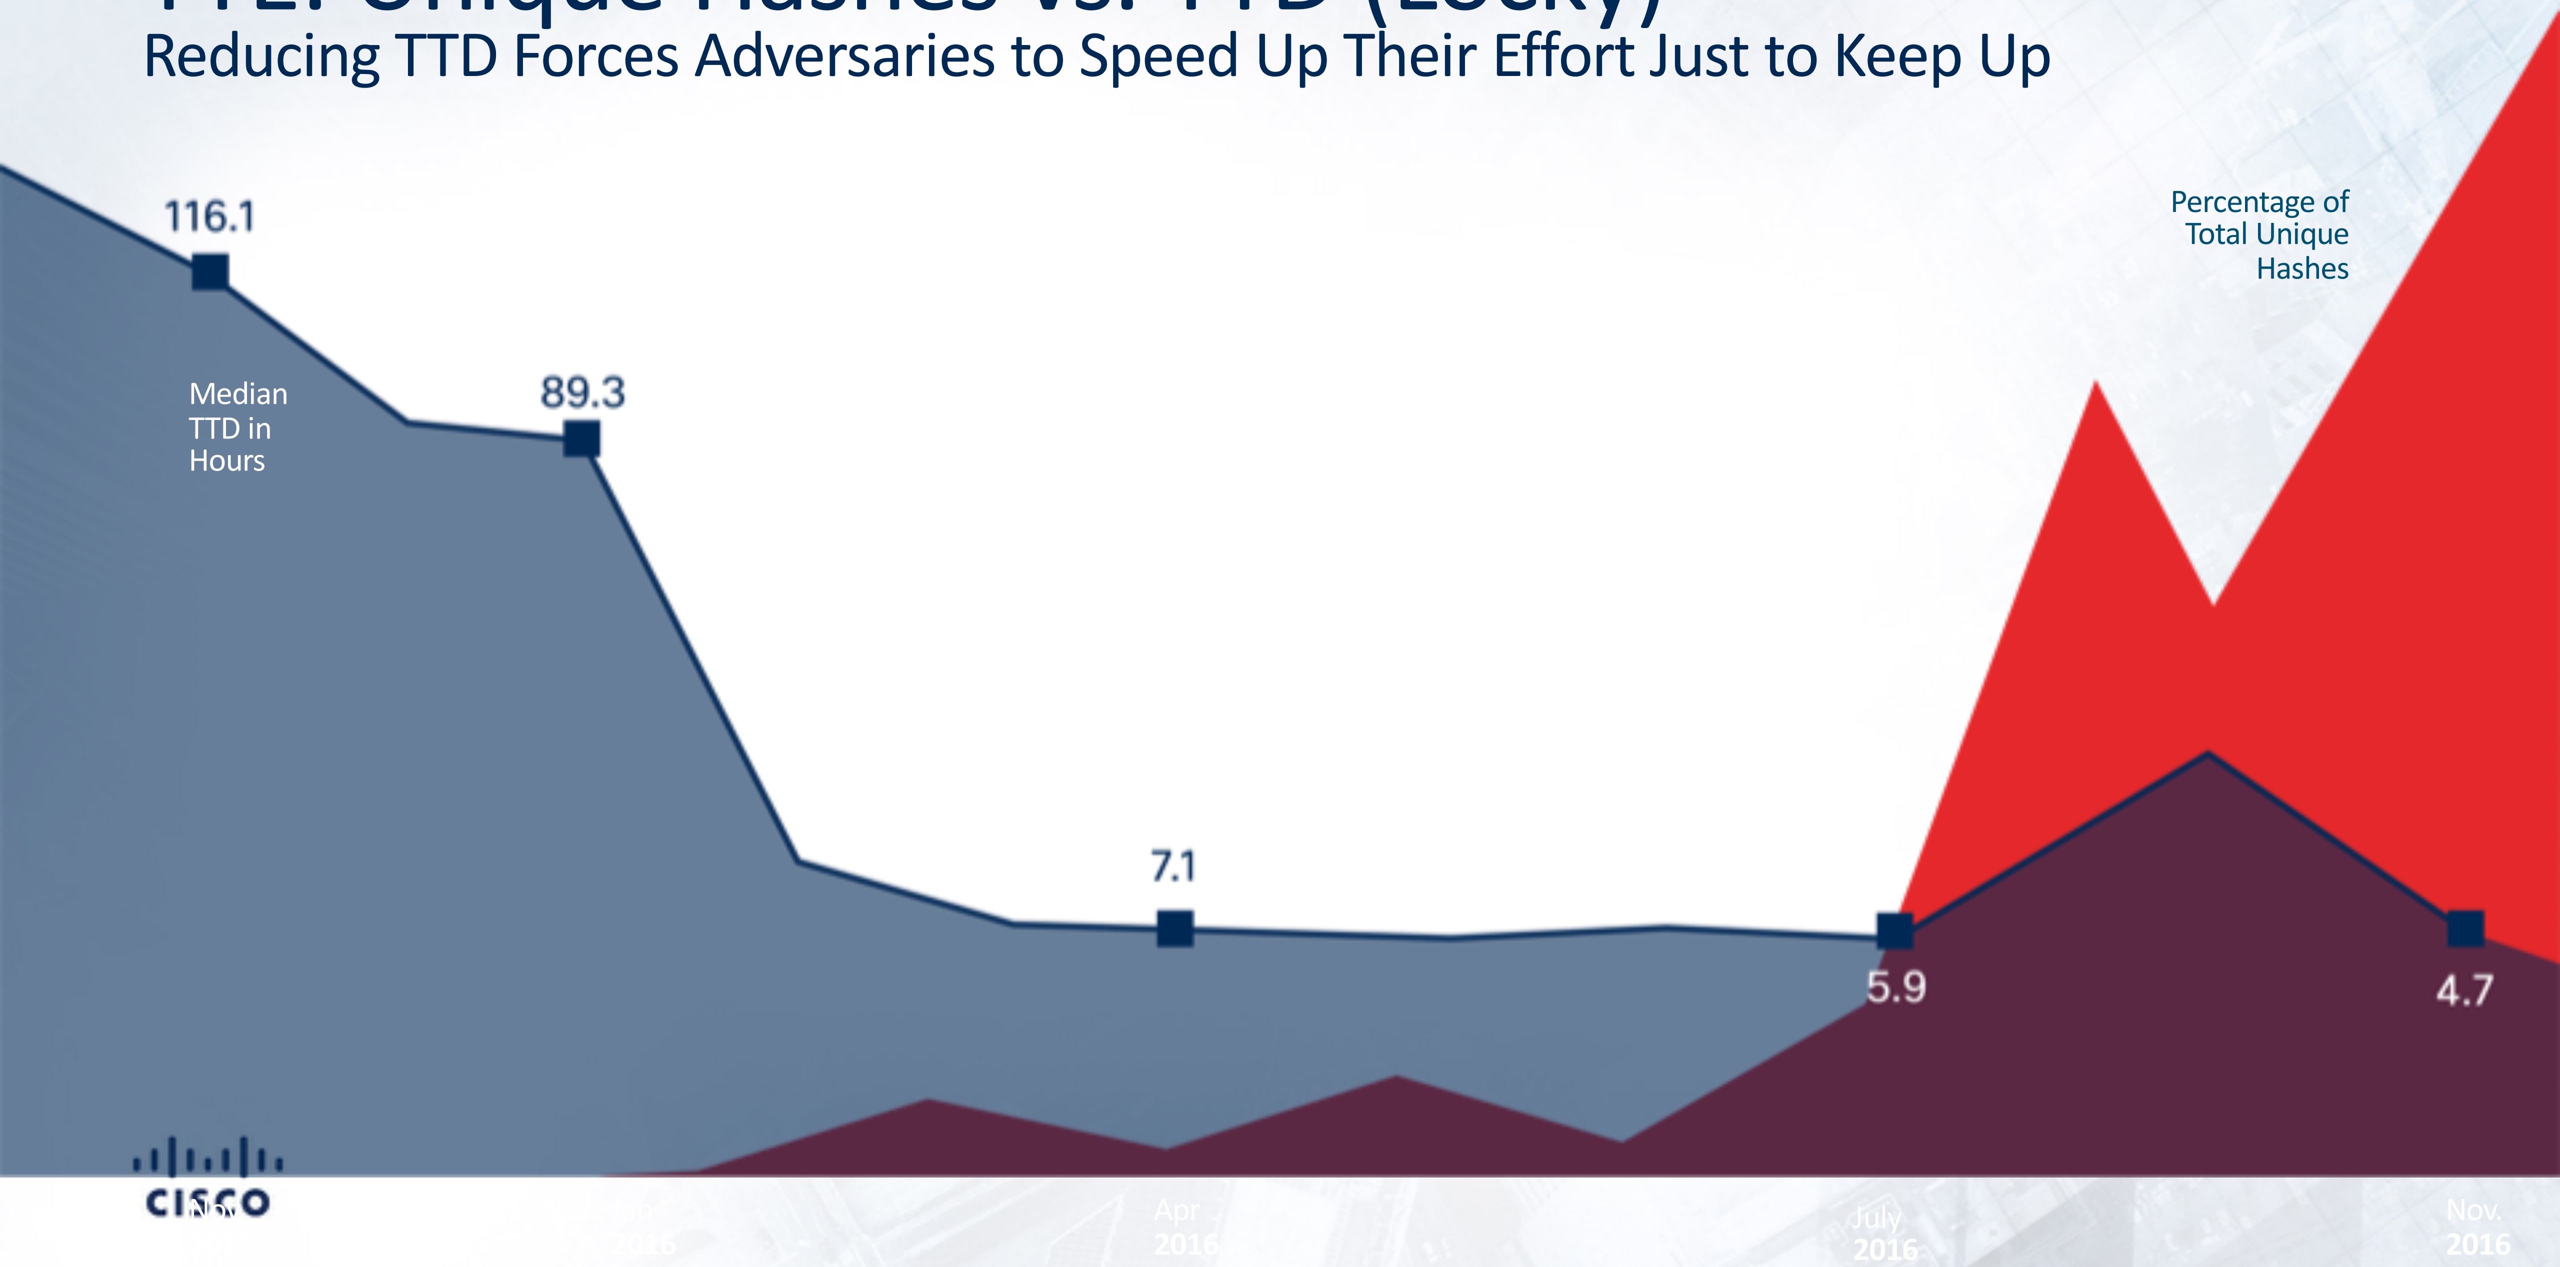

IoT and DDos

Application-layer attacks are rising, network-layer attacks are declining

Burst attacks are increasing

- Complexity
- Frequency
- Duration

Amplification attacks

2/5

of businesses experienced a reflection amplification attack in 2017

2/3 of those organizations mitigated the attacks

Source: Radware

IT/OT Attack Sentiment

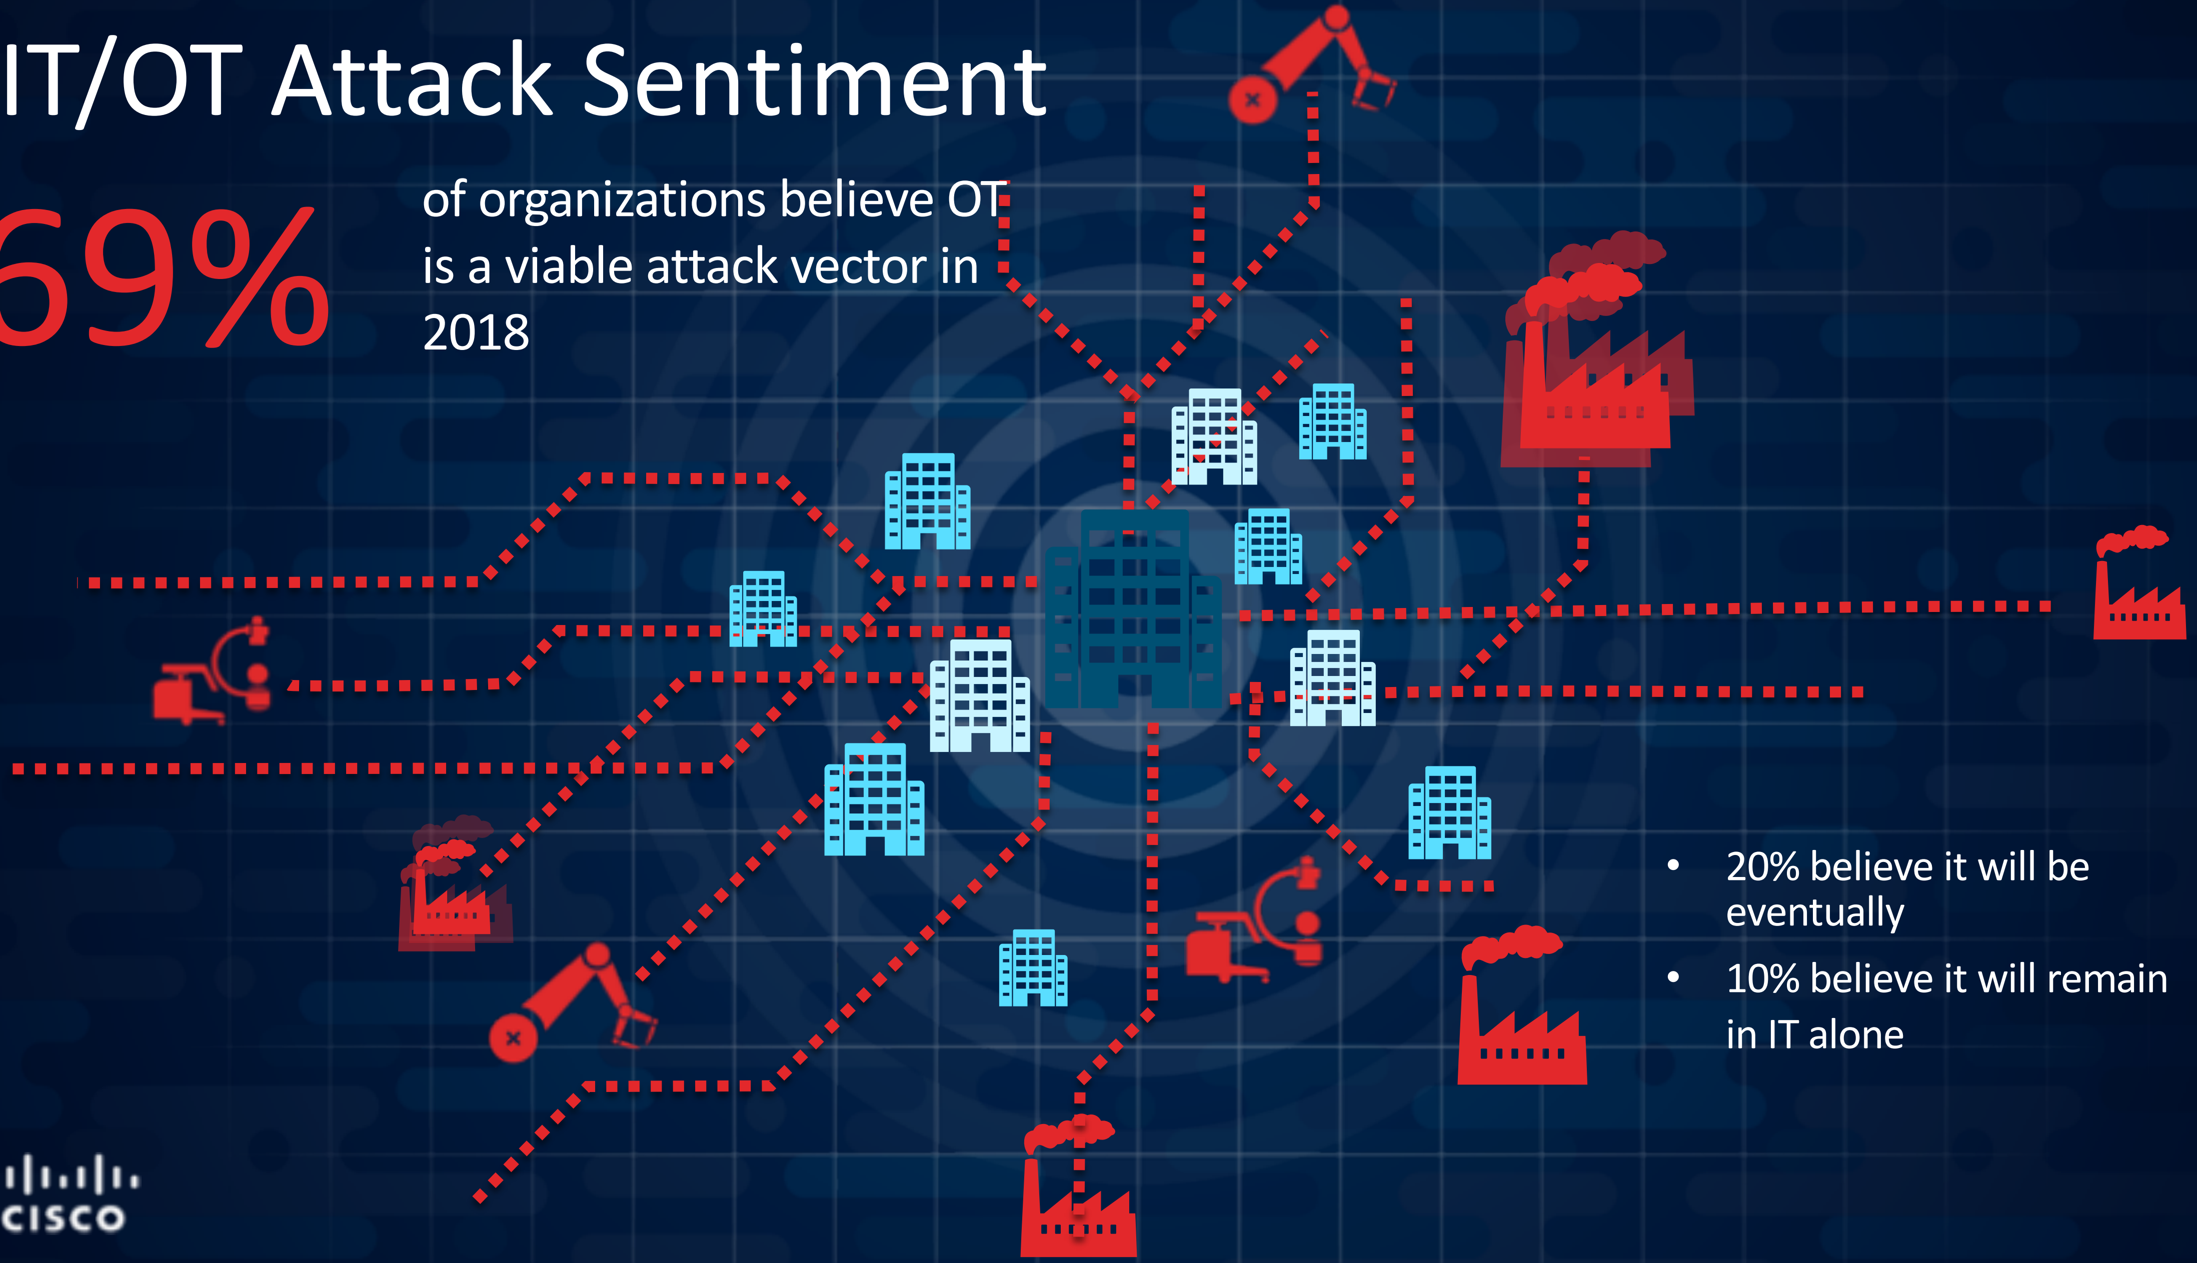

69% of organizations believe OT is a viable attack vector in 2018

- 20% believe it will be eventually
- 10% believe it will remain in IT alone

What is result?

Uninvestigated alerts still create huge business risk

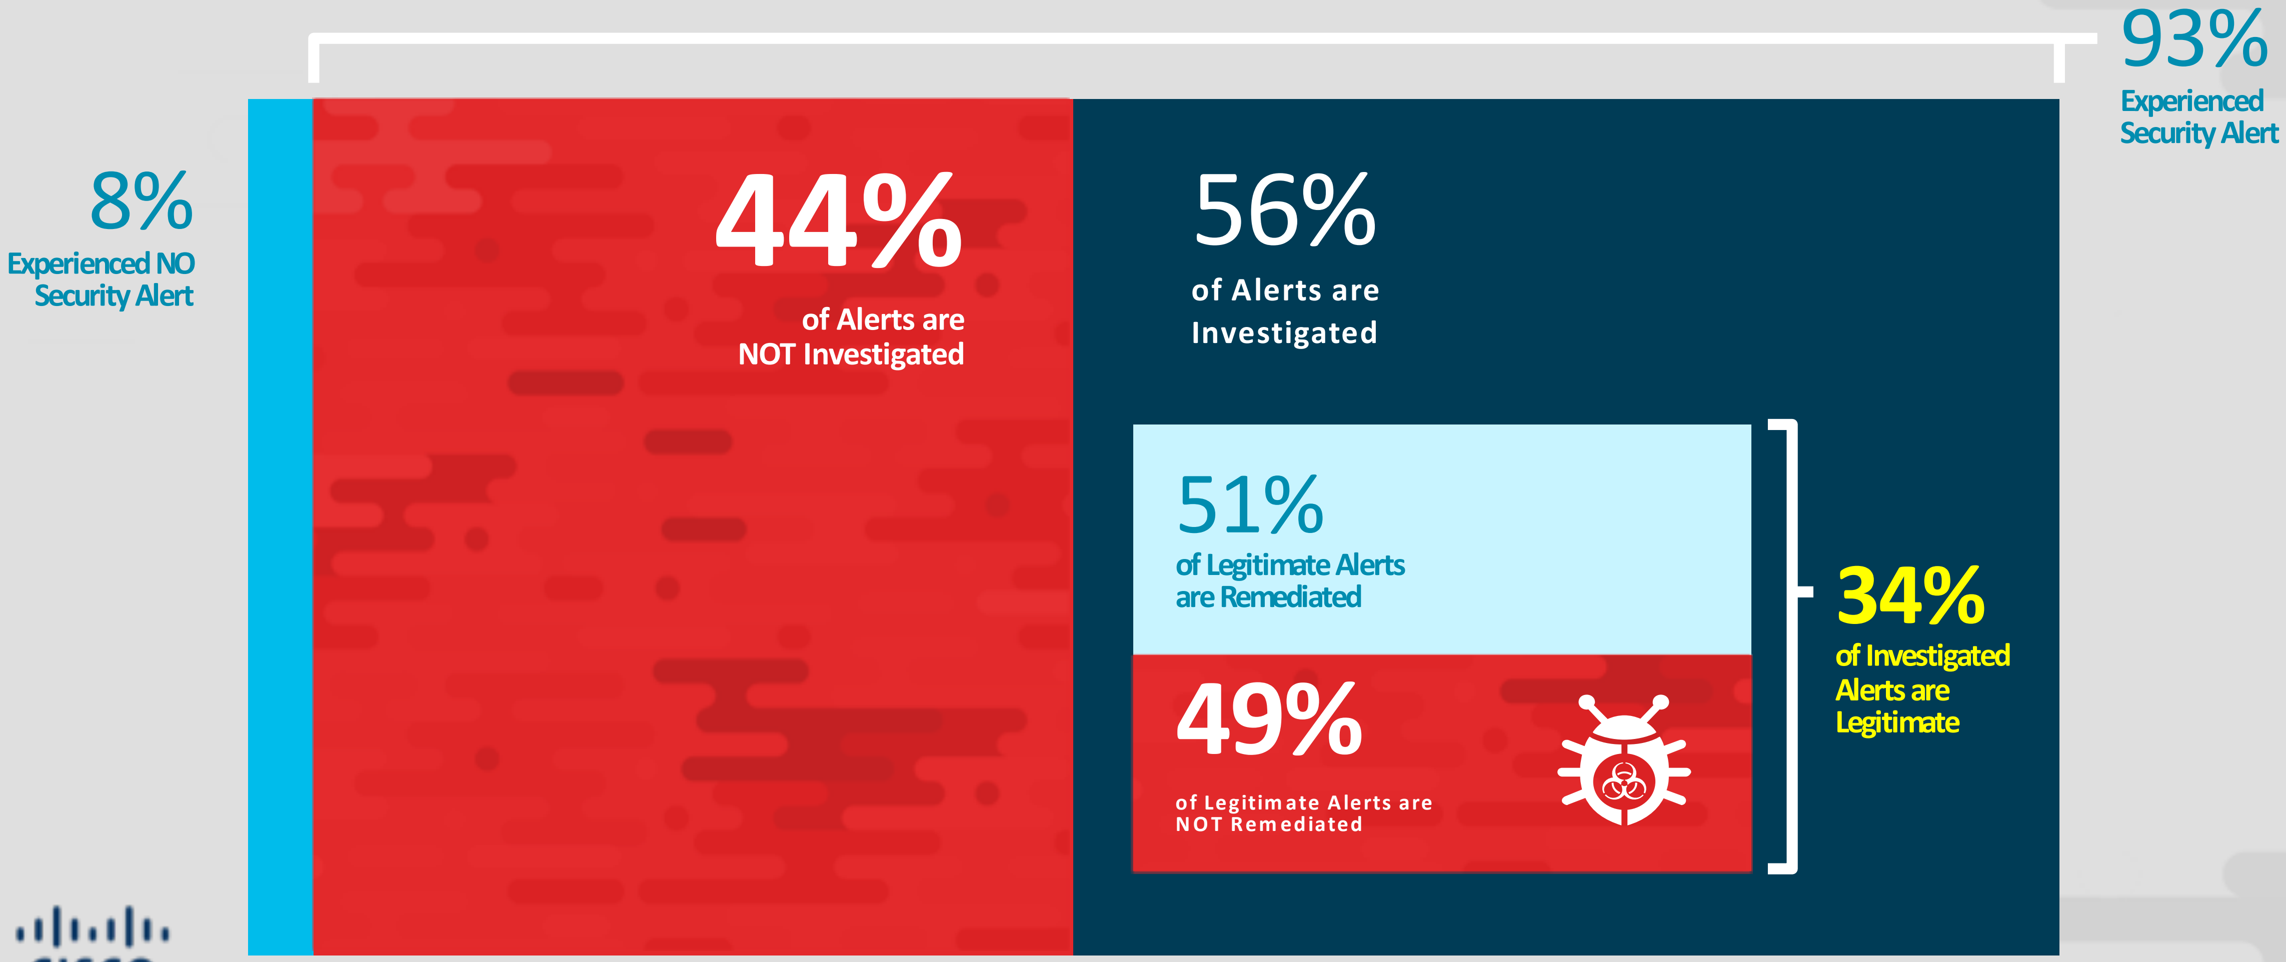

93%
Experienced Security Alert

8%
Experienced NO Security Alert

44%
of Alerts are NOT Investigated

56%
of Alerts are Investigated

51%
of Legitimate Alerts are Remediated

49%
of Legitimate Alerts are NOT Remediated

34%
of Investigated Alerts are Legitimate

MACHINE LEARNING

What it is

What did we do before Machine Learning?

Use in combination with Machine Learning

Simple Pattern Matching (signatures, IoCs…)

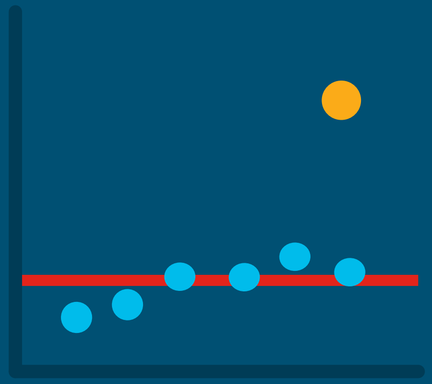

Statistical Methods

Rules and First Order Logic (FoL)

The threat funnel

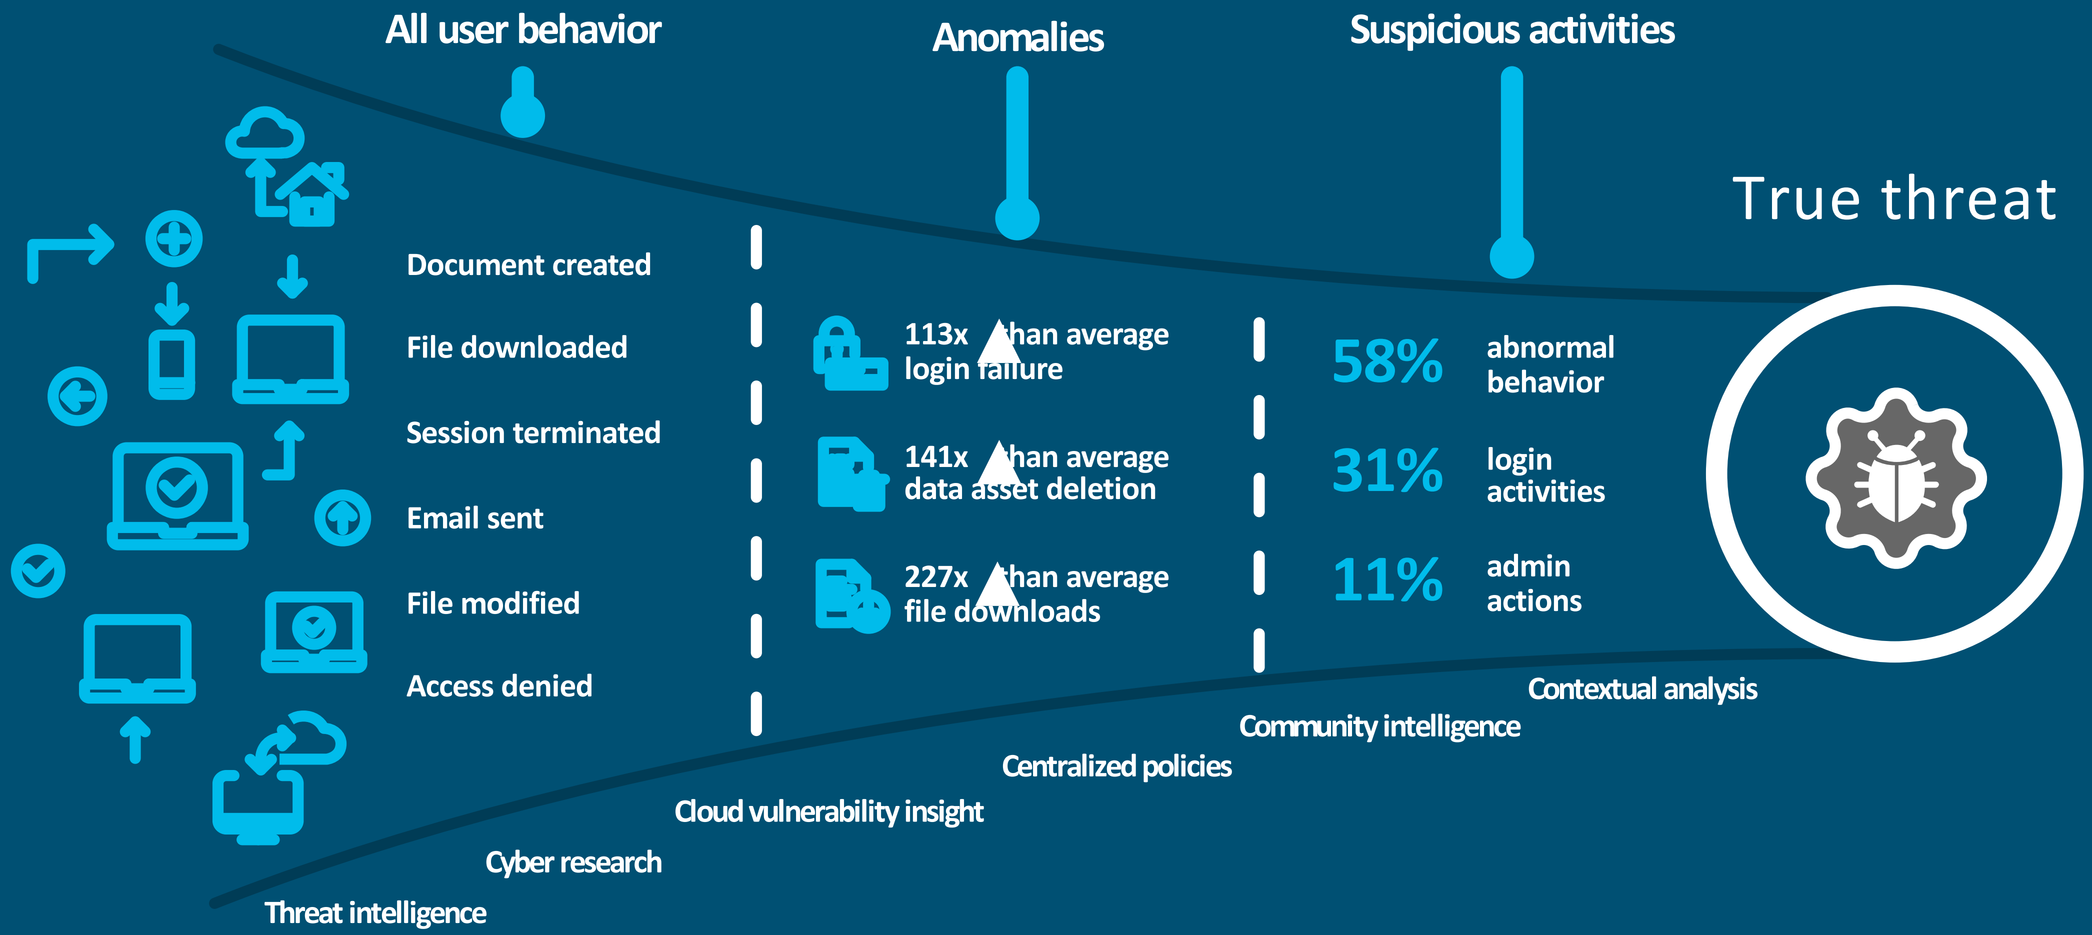

All user behavior

Document created

File downloaded

Session terminated

Email sent

File modified

Access denied

Anomalies

113x ▲ than average login failure

141x ▲ than average data asset deletion

227x ▲ than average file downloads

Suspicious activities

58% abnormal behavior

31% login activities

11% admin actions

True threat

Threat intelligence

Cyber research

Cloud vulnerability insight

Centralized policies

Community intelligence

Contextual analysis

Source: Cloudlock CyberLab

instance based

clustering

regularization

ensemble

bayesian

rule system

ground truth

machine learning algorithms

classifier

regression

decision tree

deep learning

neural network

dimensionality reduction

⚠ NERD ALERT

Machine learnings comes with it a lot of terms that are incredibly confusing

Machine Learning

The Big Picture

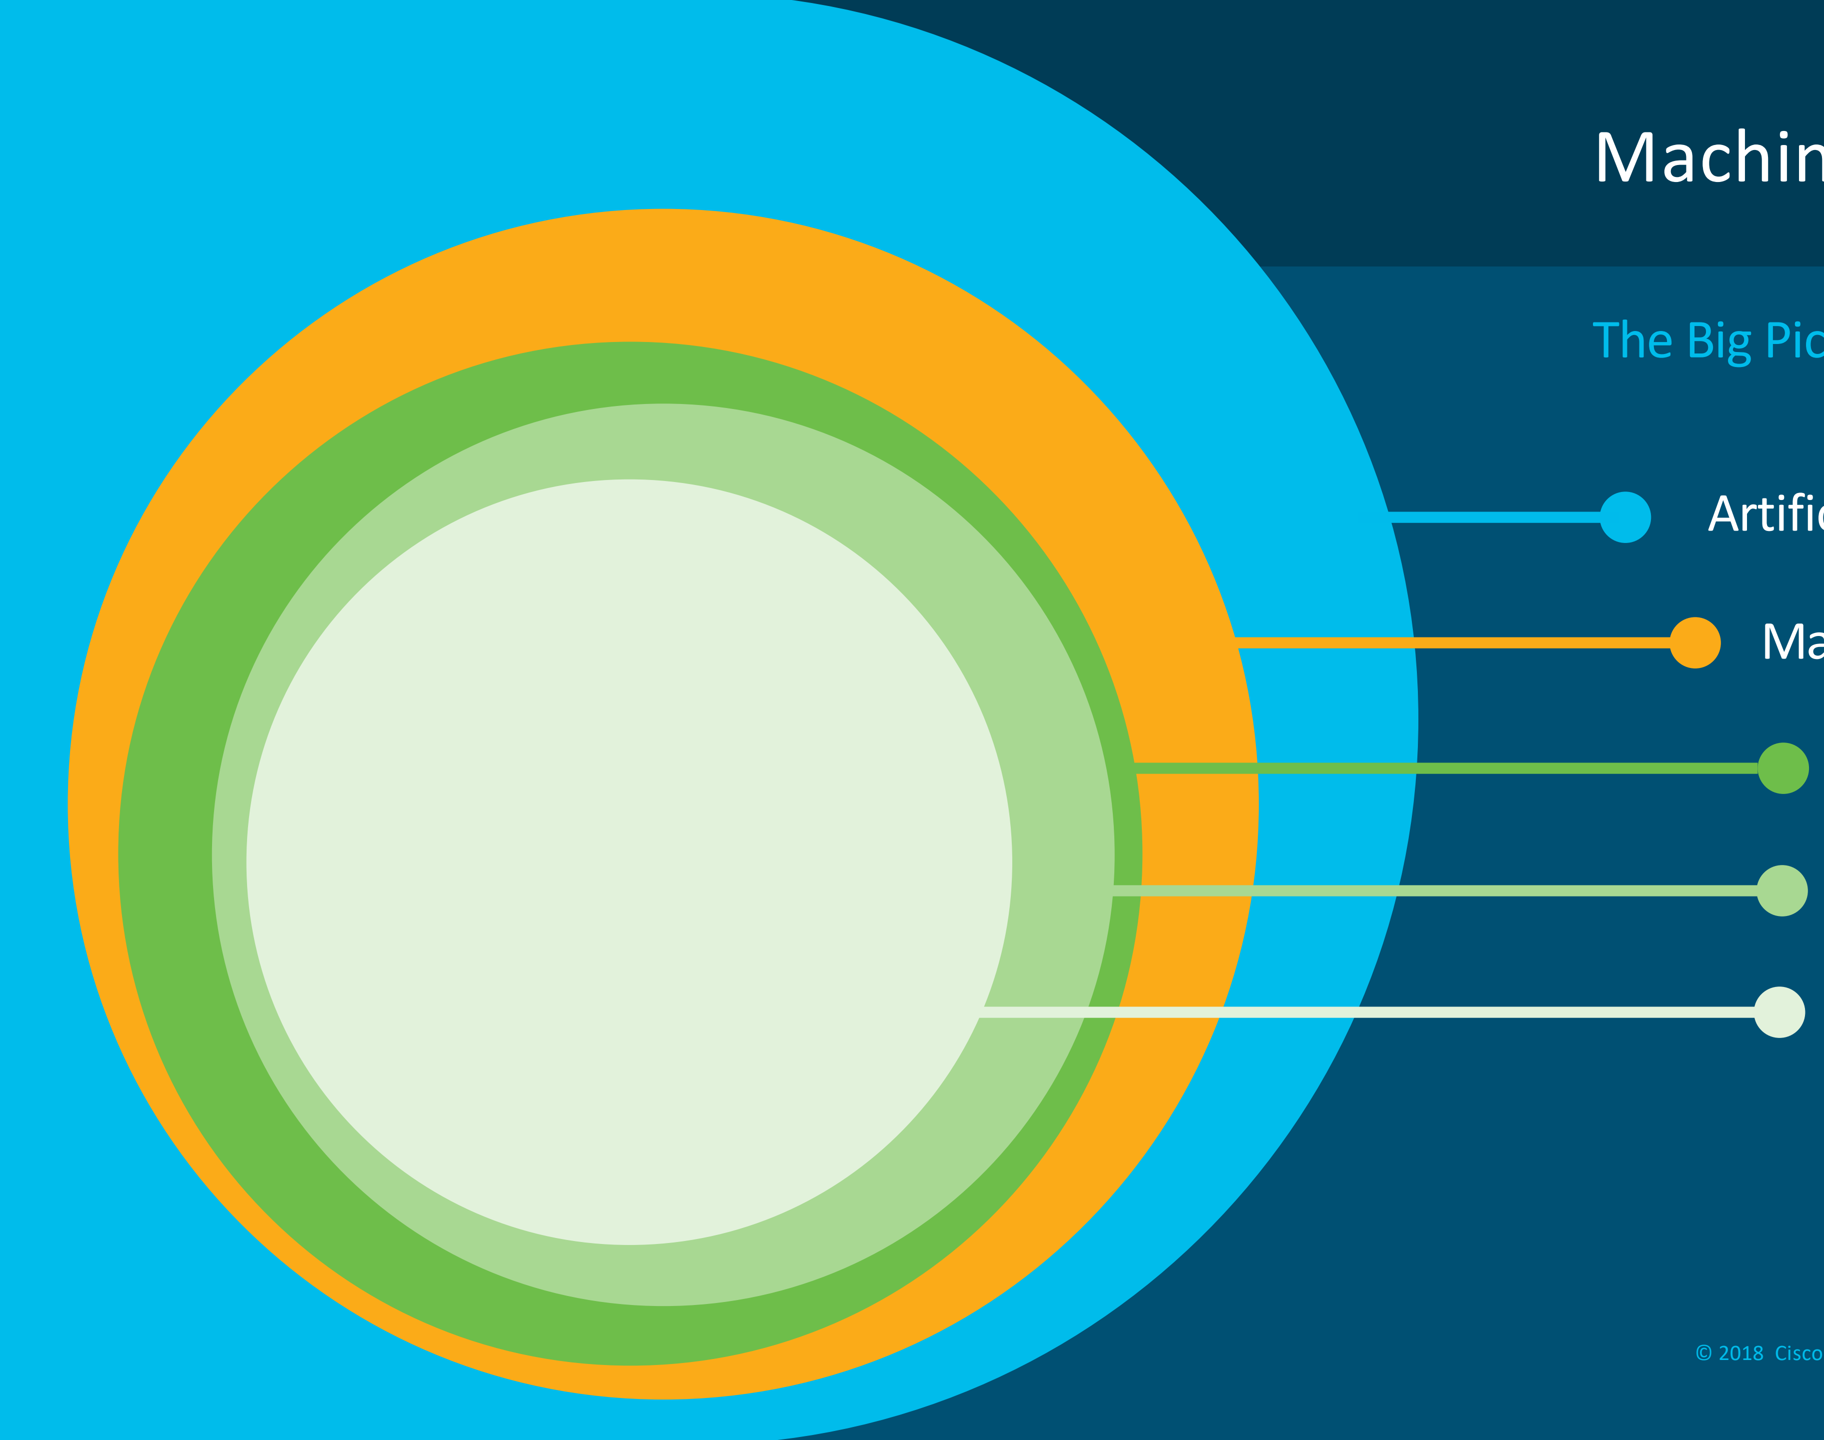

- Artificial Intelligence
- Machine Learning
- Supervised Learning
- Unsupervised Learning
- Reinforcement Learning

Machine Learning

Common Techniques

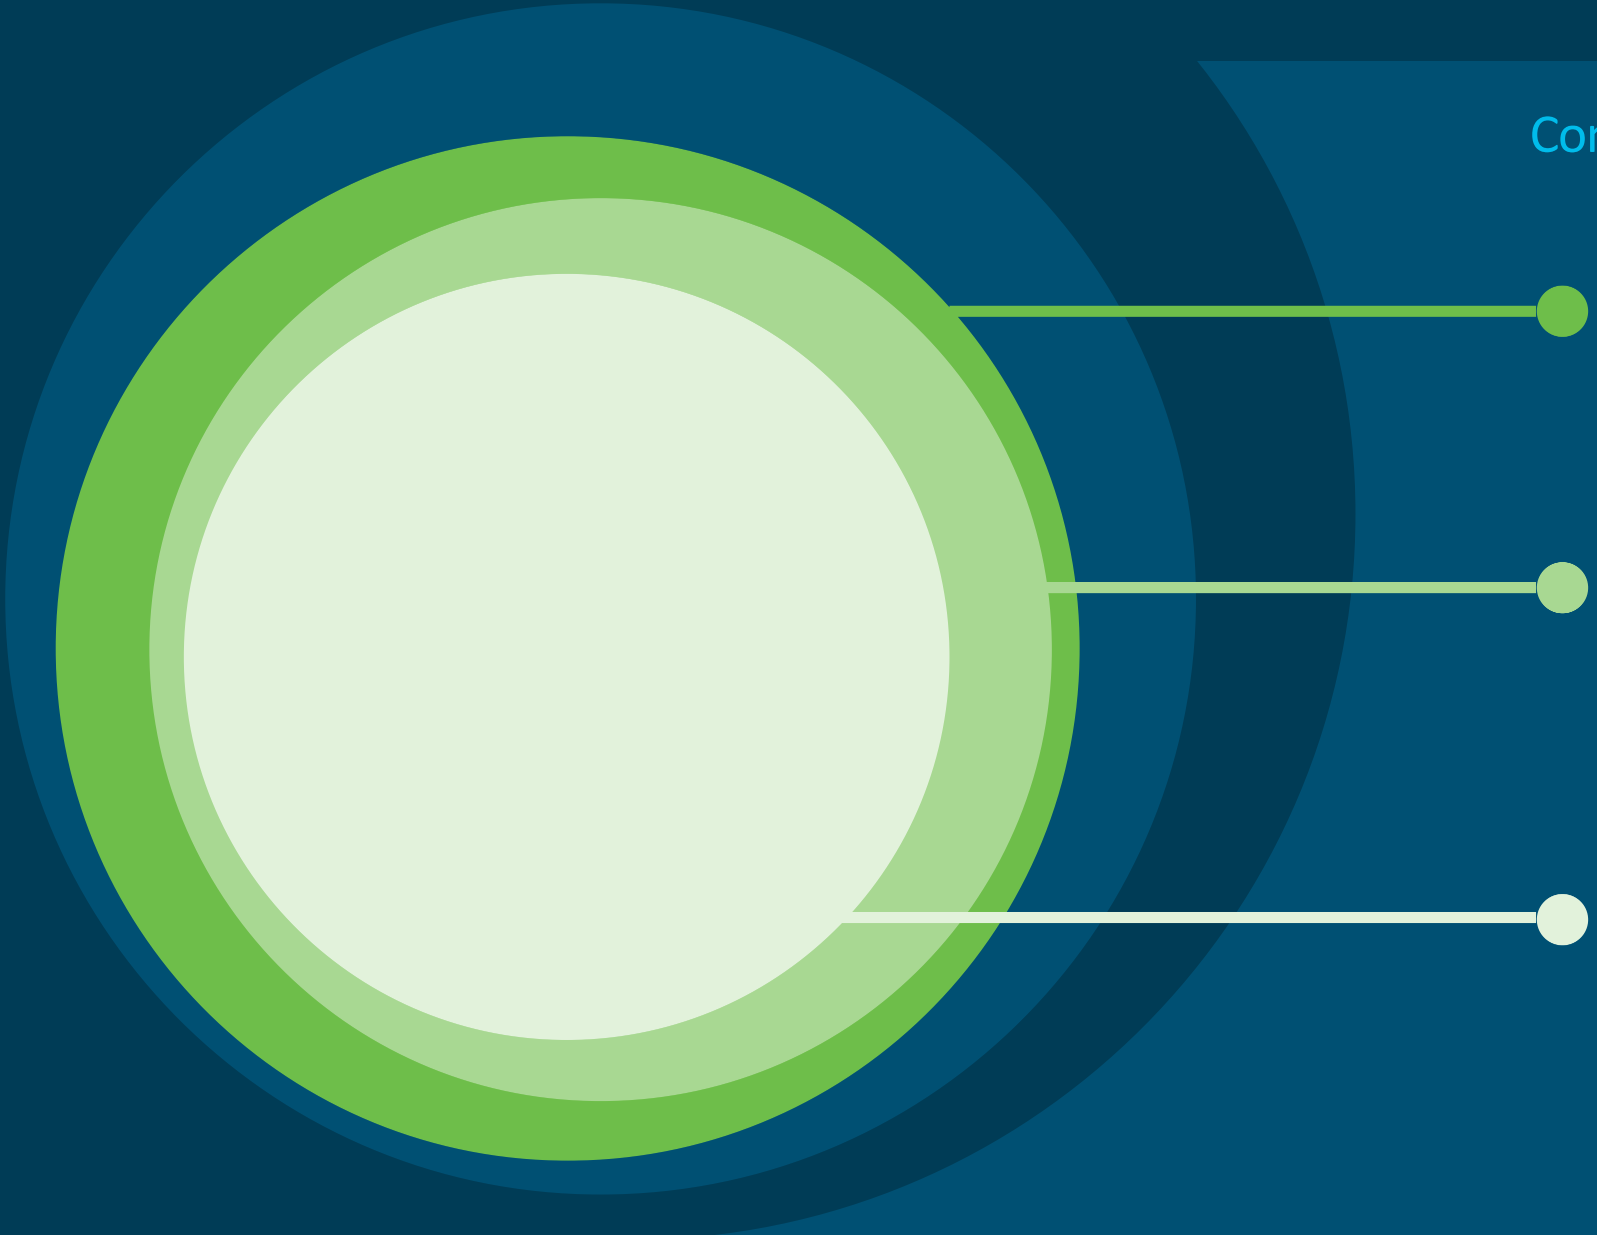

Supervised Learning

When you know the question you are trying to ask and have examples of it being asked and answered correction

Unsupervised Learning

You don't have answers and may not fully know the questions

Reinforcement Learning

“The other” category
Trial and error behavior effective in game scenarios

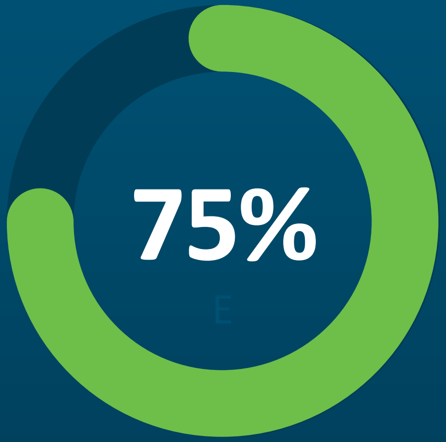

Supervised Learning

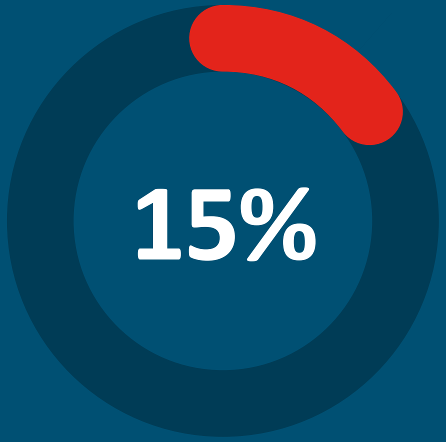

Unsupervised Learning

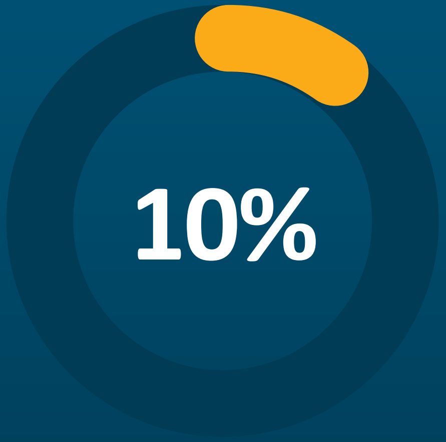

Other

(Reinforcement Learning, etc.)

MACHINE LEARNING

Techniques

Training Classifiers

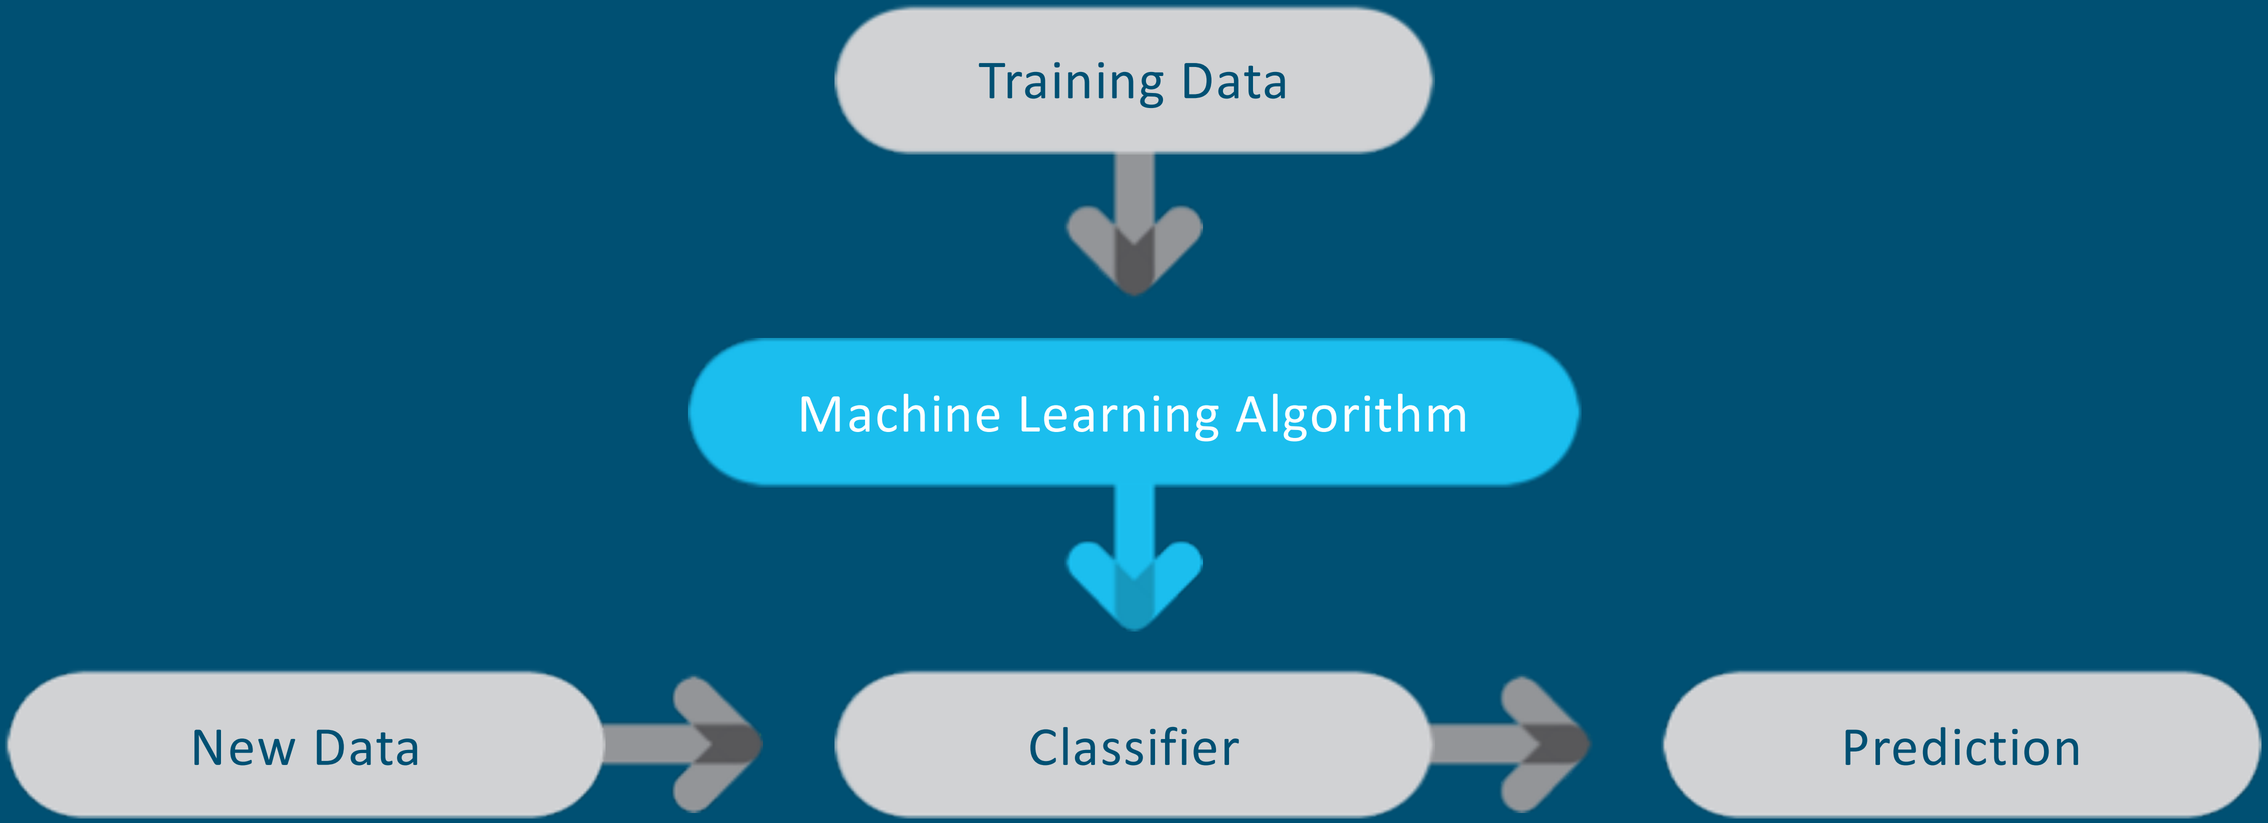

Training Data

Machine Learning Algorithm

New Data

Classifier

Prediction

Training Data from Cisco Talos

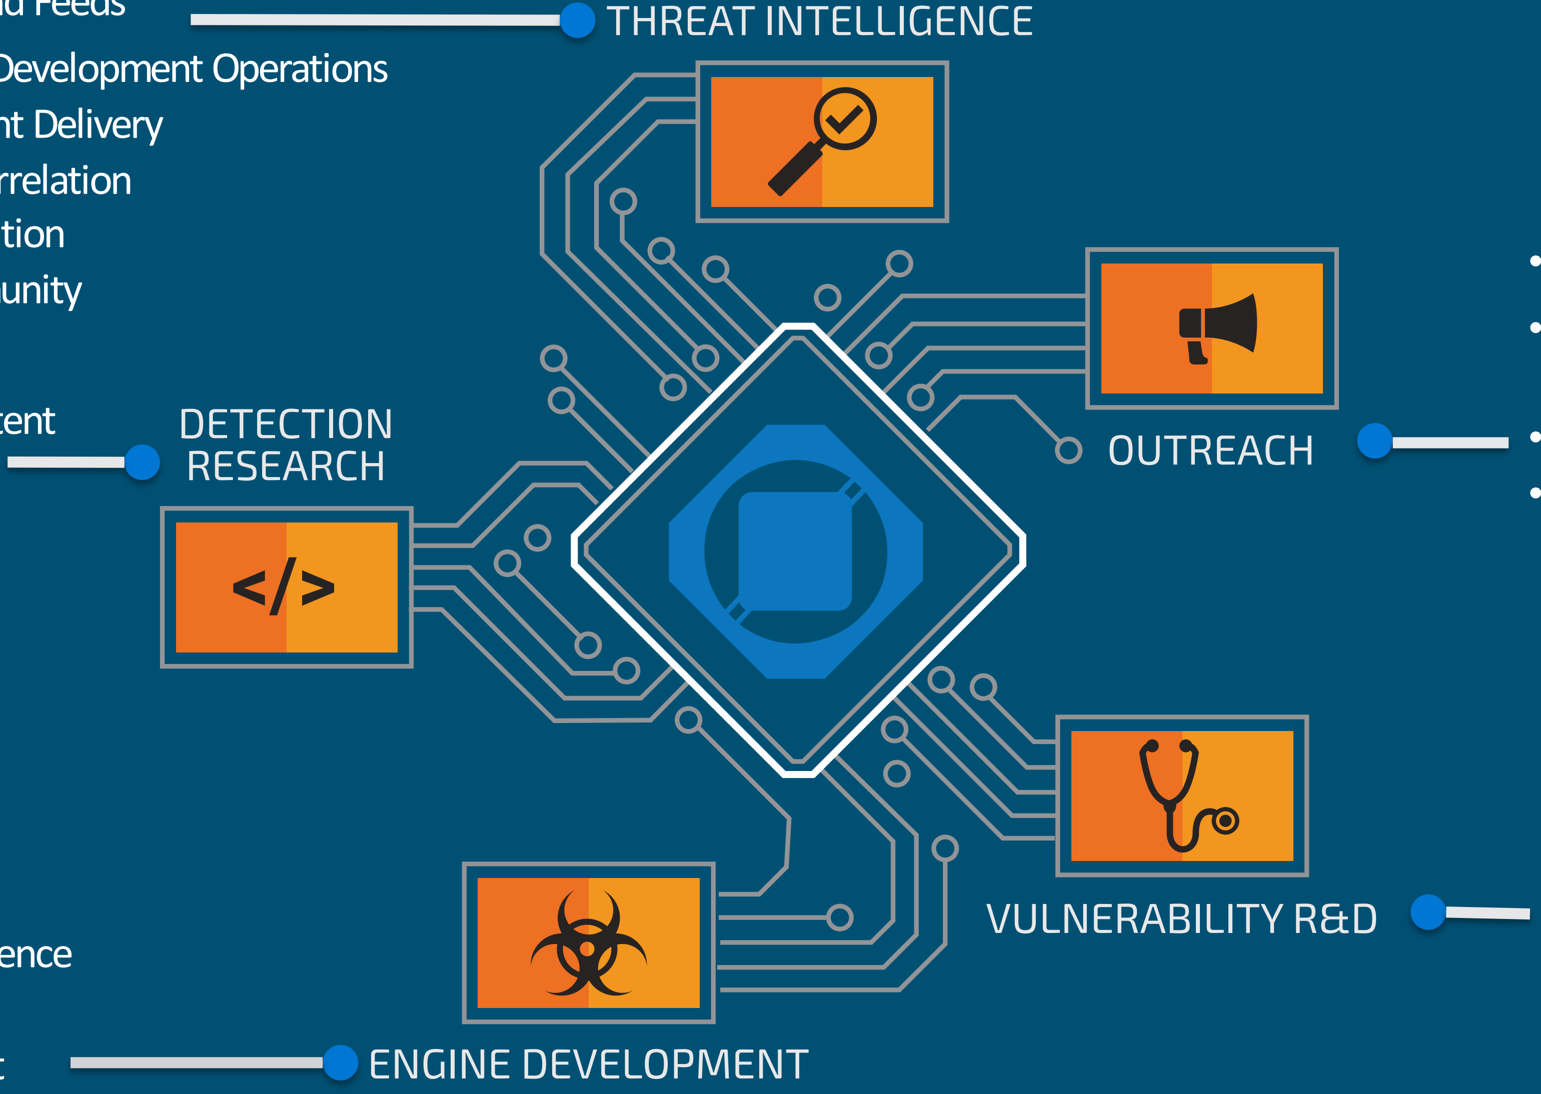

- Inbound & Outbound Feeds
- Internal Systems & Development Operations
- All Detection Content Delivery
- Data Analytics & Correlation
- Threat Actor Attribution
- Open Source Community

- Detection & Prevention Content
- Vulnerability Research
- Malware Research
- Detection Research
- Policy Improvements

- Thought Leadership
- Consistent, Repeatable Security Messaging
- Threat Reports
- Media Relations

- Discovery
- Triage
- Exploit Development
- Mitigations

- Intelligence Systems
- Web & Email Intelligence
- Sandbox
- Engine Development
- ClamAV Development

Real Cisco Big Data for Security Training Set

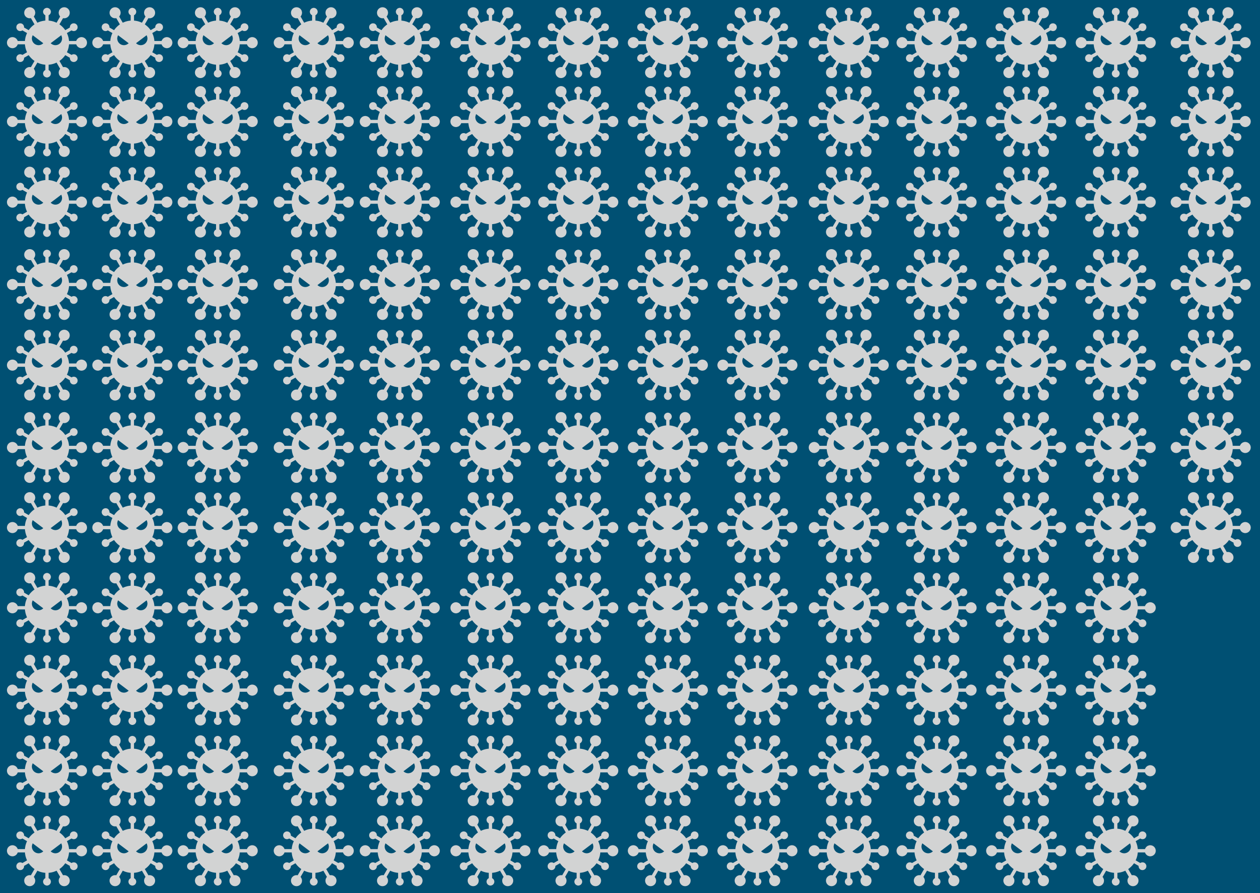

= 10,000

1.5 Million

Malware Samples

DAILY

Google: facts and numbers

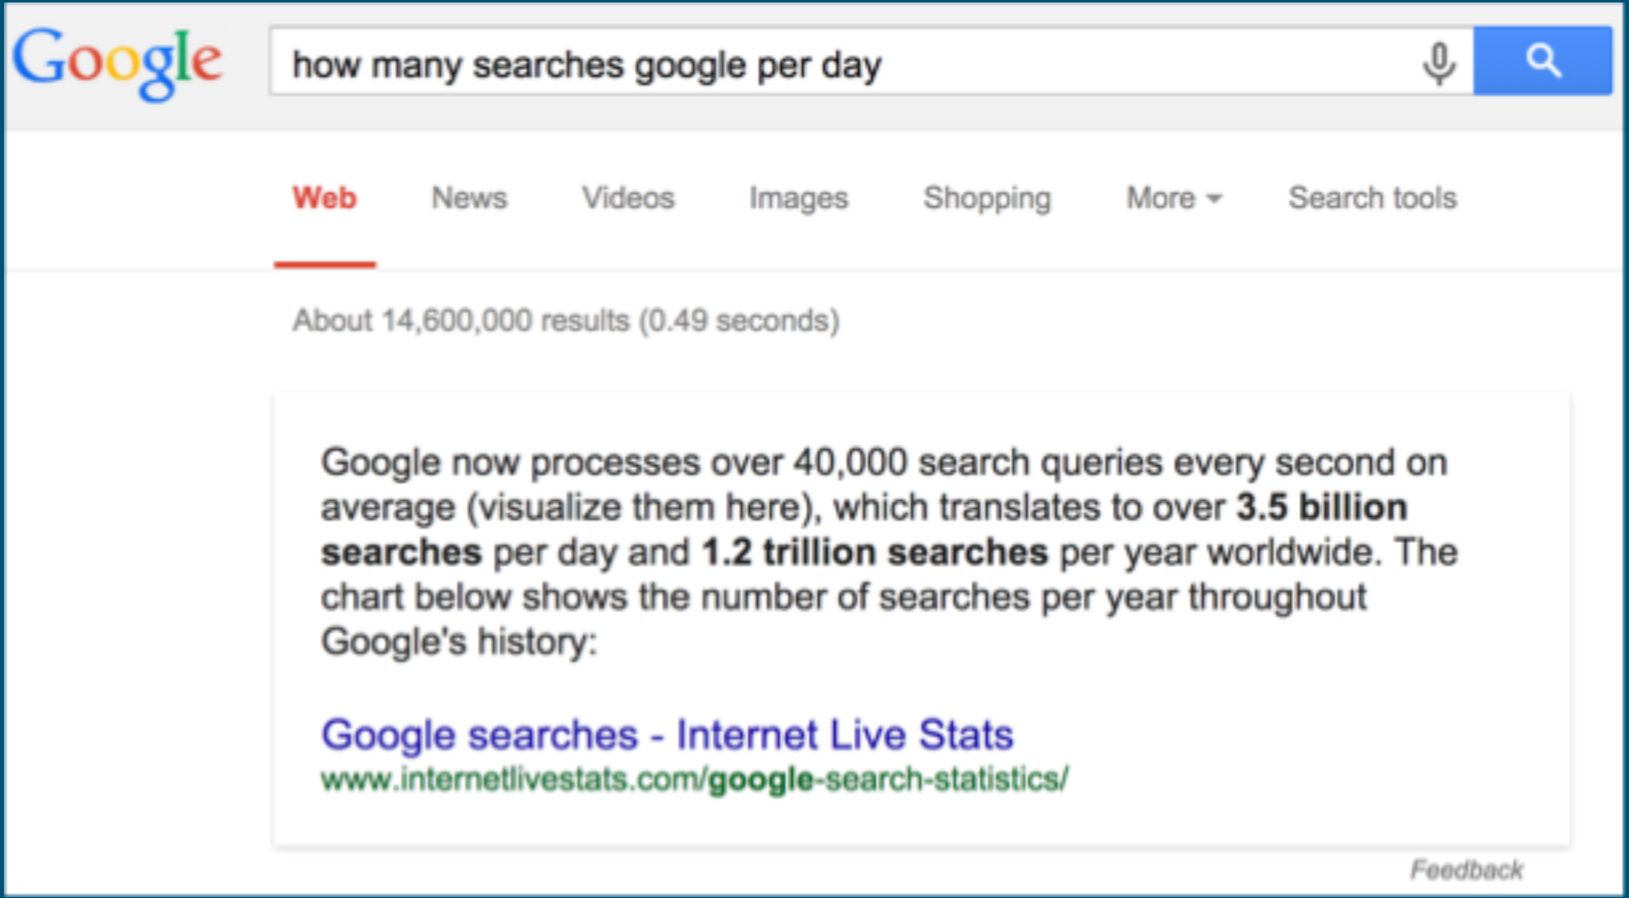

About 14,600,000 results (0.49 seconds)

Google now processes over 40,000 search queries every second on average (visualize them here), which translates to over **3.5 billion searches** per day and **1.2 trillion searches** per year worldwide. The chart below shows the number of searches per year throughout Google's history:

Google searches - Internet Live Stats
www.internetlivestats.com/**google**-search-statistics/

3.5 Billion

searches a day

1.2 Trillion

searches a year

Real Cisco Big Data for Security Training Set

19.7 Billion

TOTAL THREAT BLOCKS

DAILY

=

7.2 Trillion

YEARLY

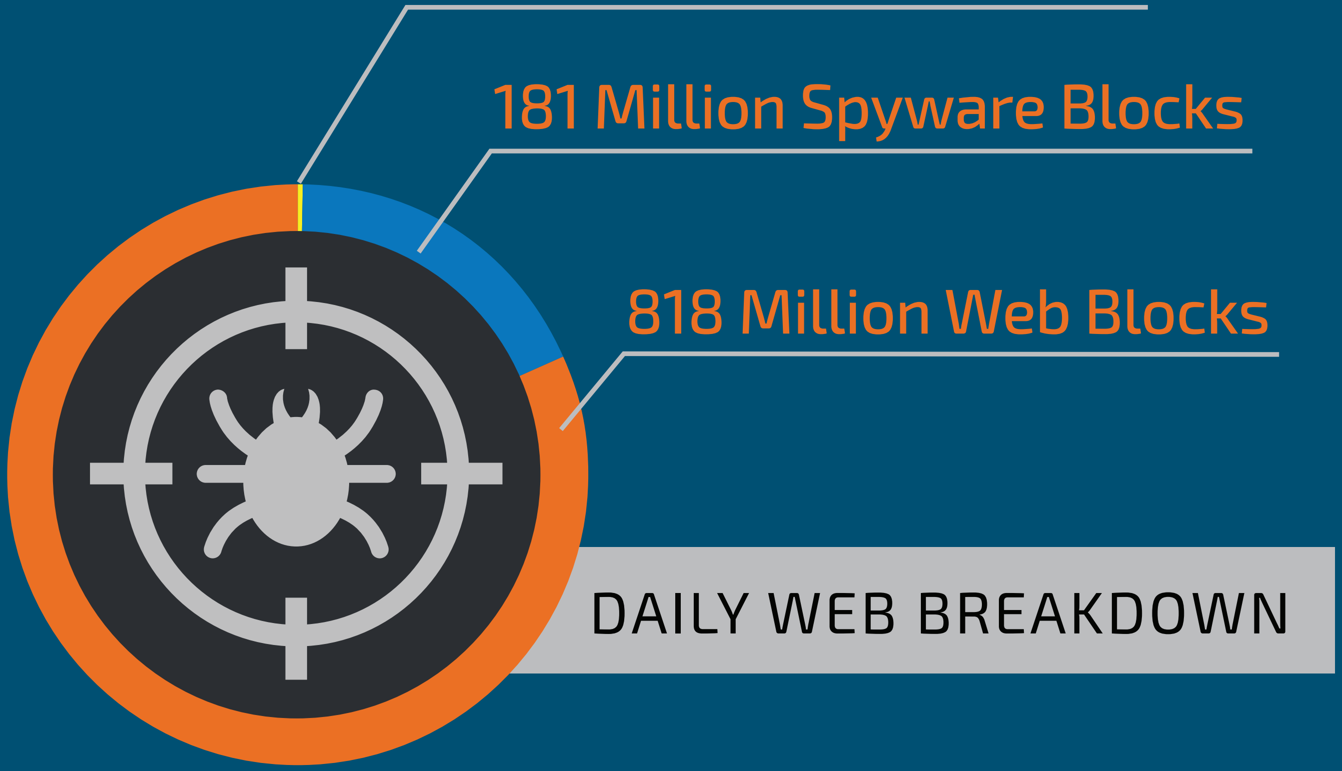

82 Thousand Virus Blocks

181 Million Spyware Blocks

818 Million Web Blocks

DAILY WEB BREAKDOWN

MACHINE LEARNING

For Security

Why is Machine Learning so useful in Security?

Static

With limited variability or is well-understood

Evolving Security

The security domain is always evolving,

has a large amount of variability,

and is not well-understood

One Size Does Not Fit All

Other ML Application ≠ Security

⚠ NERD ALERT

Warning: Success in one domain does not guarantee success in another

Multi-layer Analytical Pipeline

Cascade of specialized layers of Machine Learning algorithms

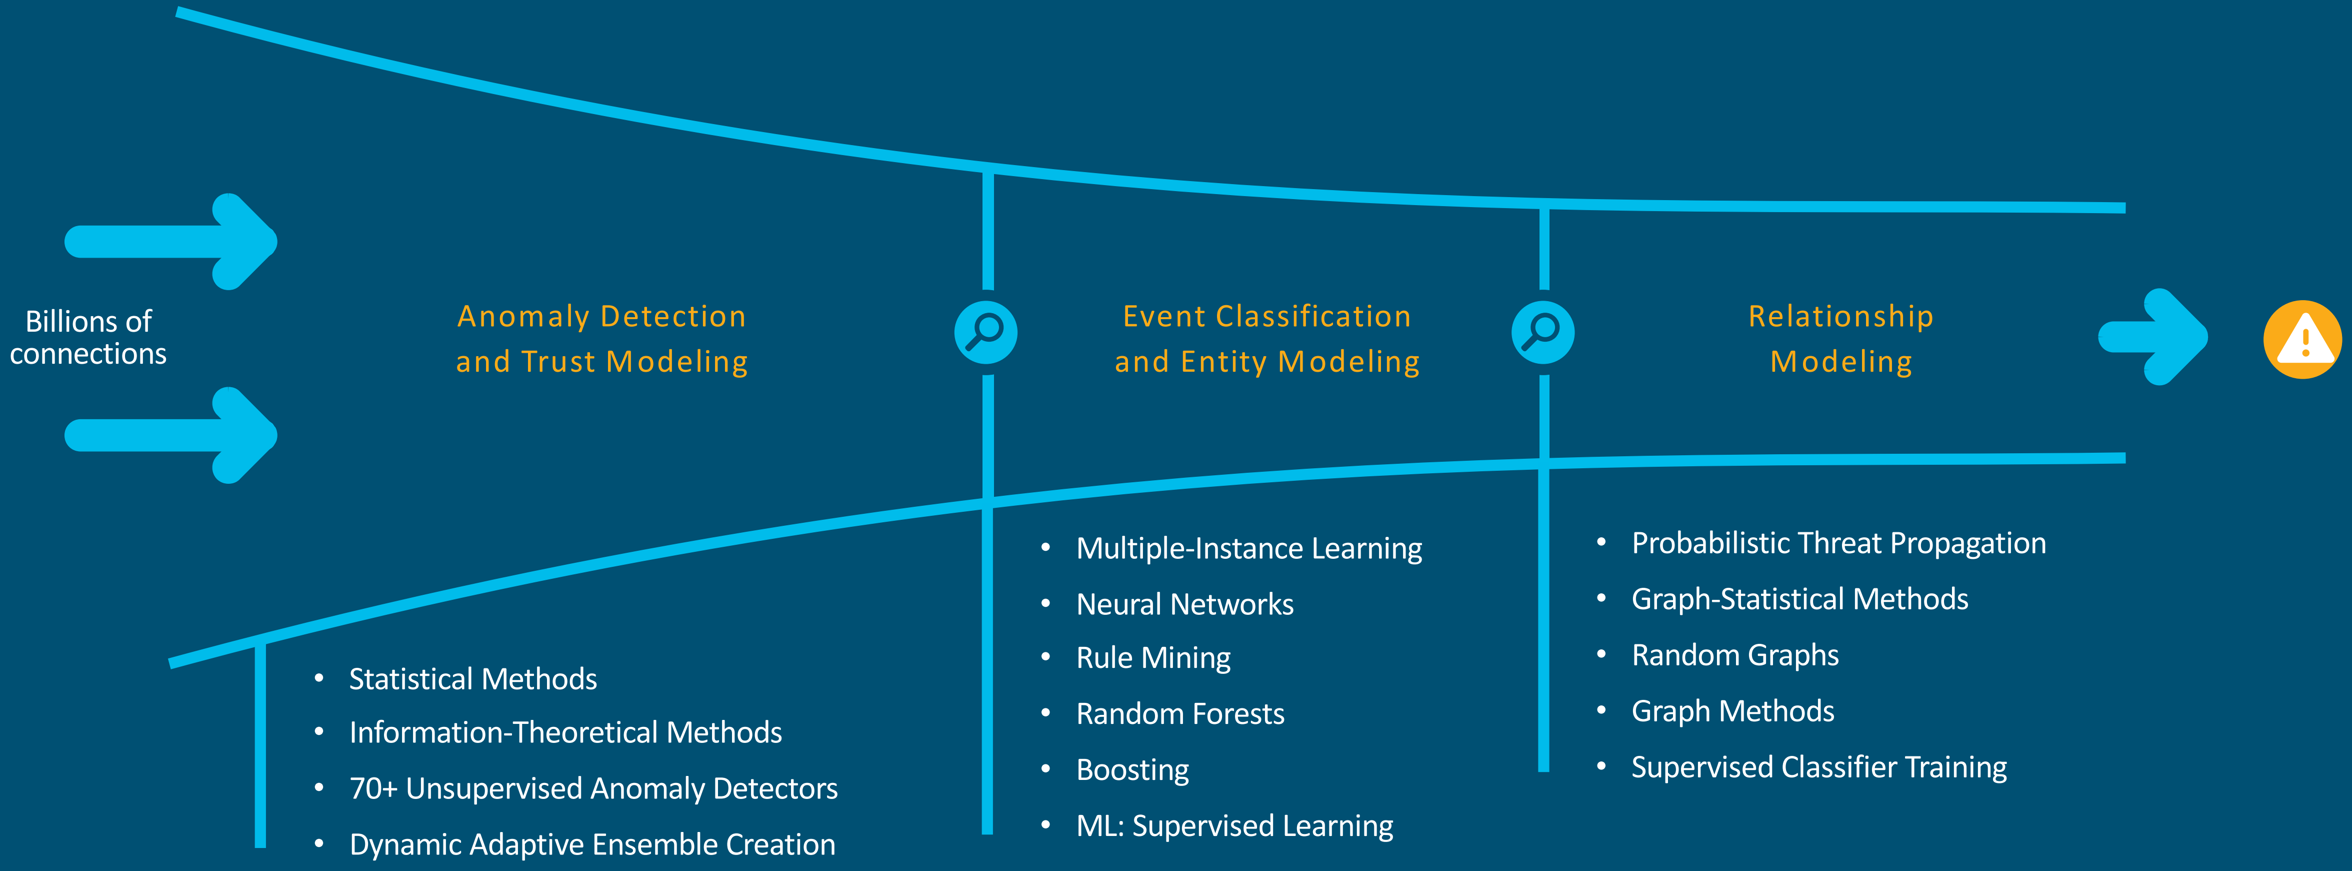

Billions of connections

Anomaly Detection and Trust Modeling

- Statistical Methods
- Information-Theoretical Methods
- 70+ Unsupervised Anomaly Detectors
- Dynamic Adaptive Ensemble Creation

Event Classification and Entity Modeling

- Multiple-Instance Learning
- Neural Networks
- Rule Mining
- Random Forests
- Boosting
- ML: Supervised Learning

Relationship Modeling

- Probabilistic Threat Propagation
- Graph-Statistical Methods
- Random Graphs
- Graph Methods
- Supervised Classifier Training

MACHINE LEARNING vs HACKERS

Real examples

Malicious Activity and Encryption

Attackers embrace encryption to conceal command-and-control activity

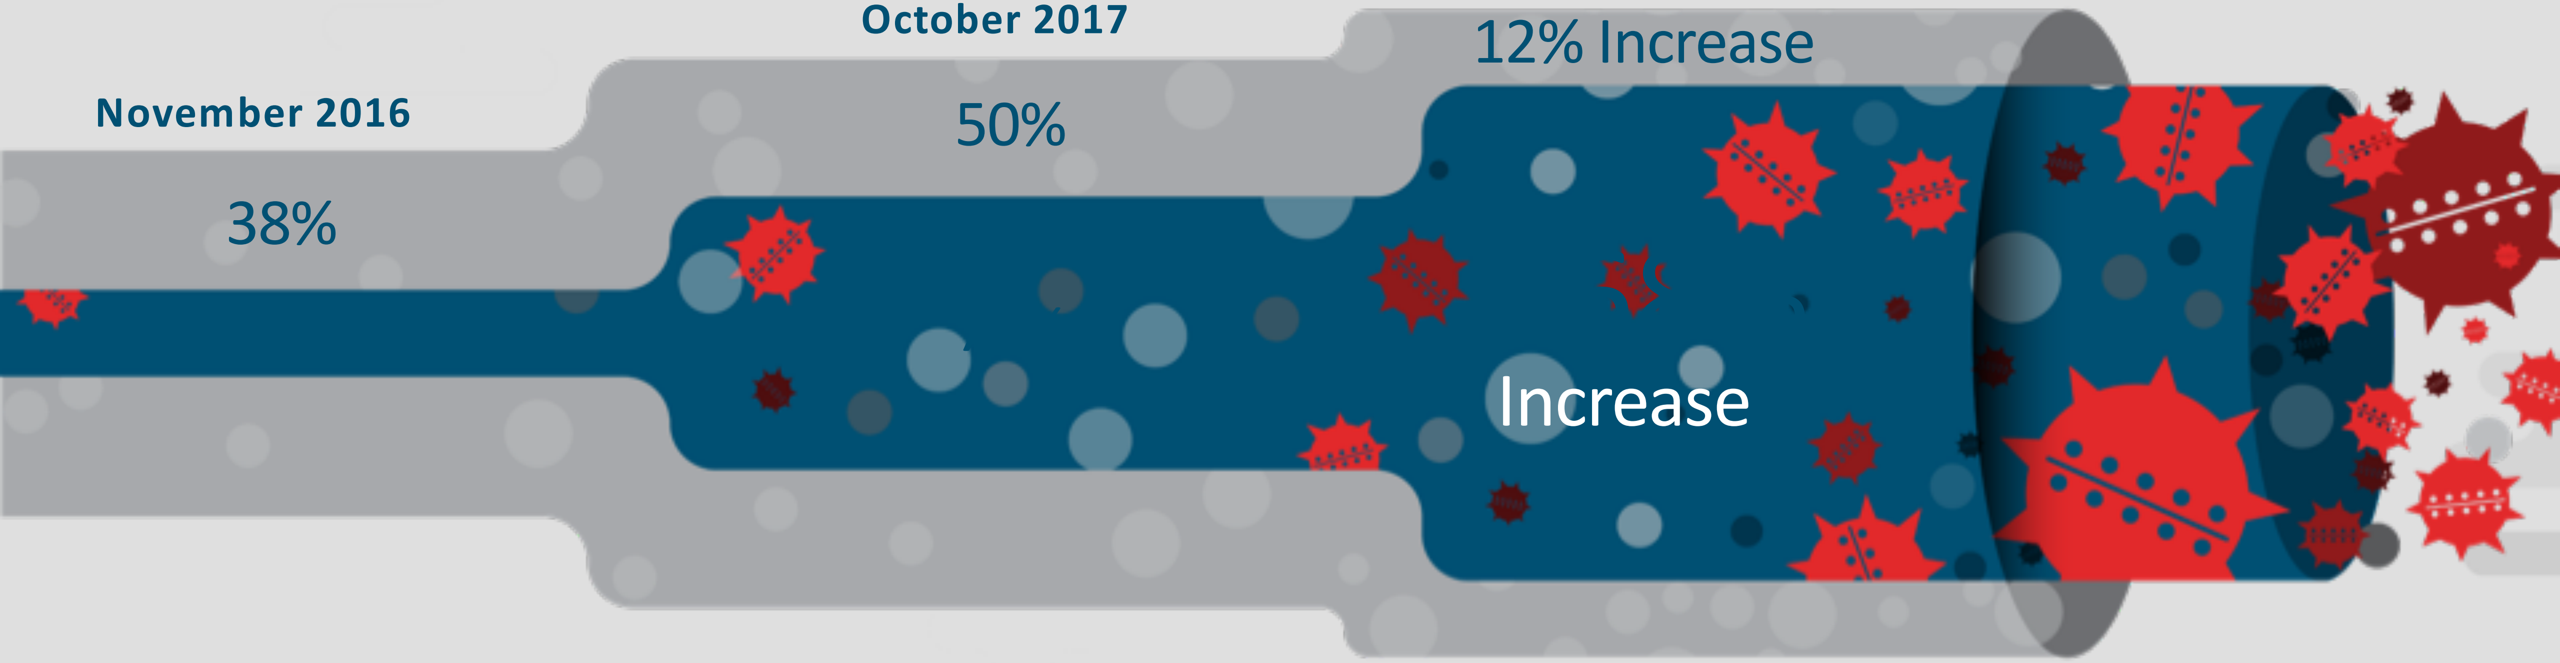

● Global Encrypted Web Traffic

● Malicious Sandbox Binaries with Encryption

Finding Malicious Activity in Encrypted Traffic

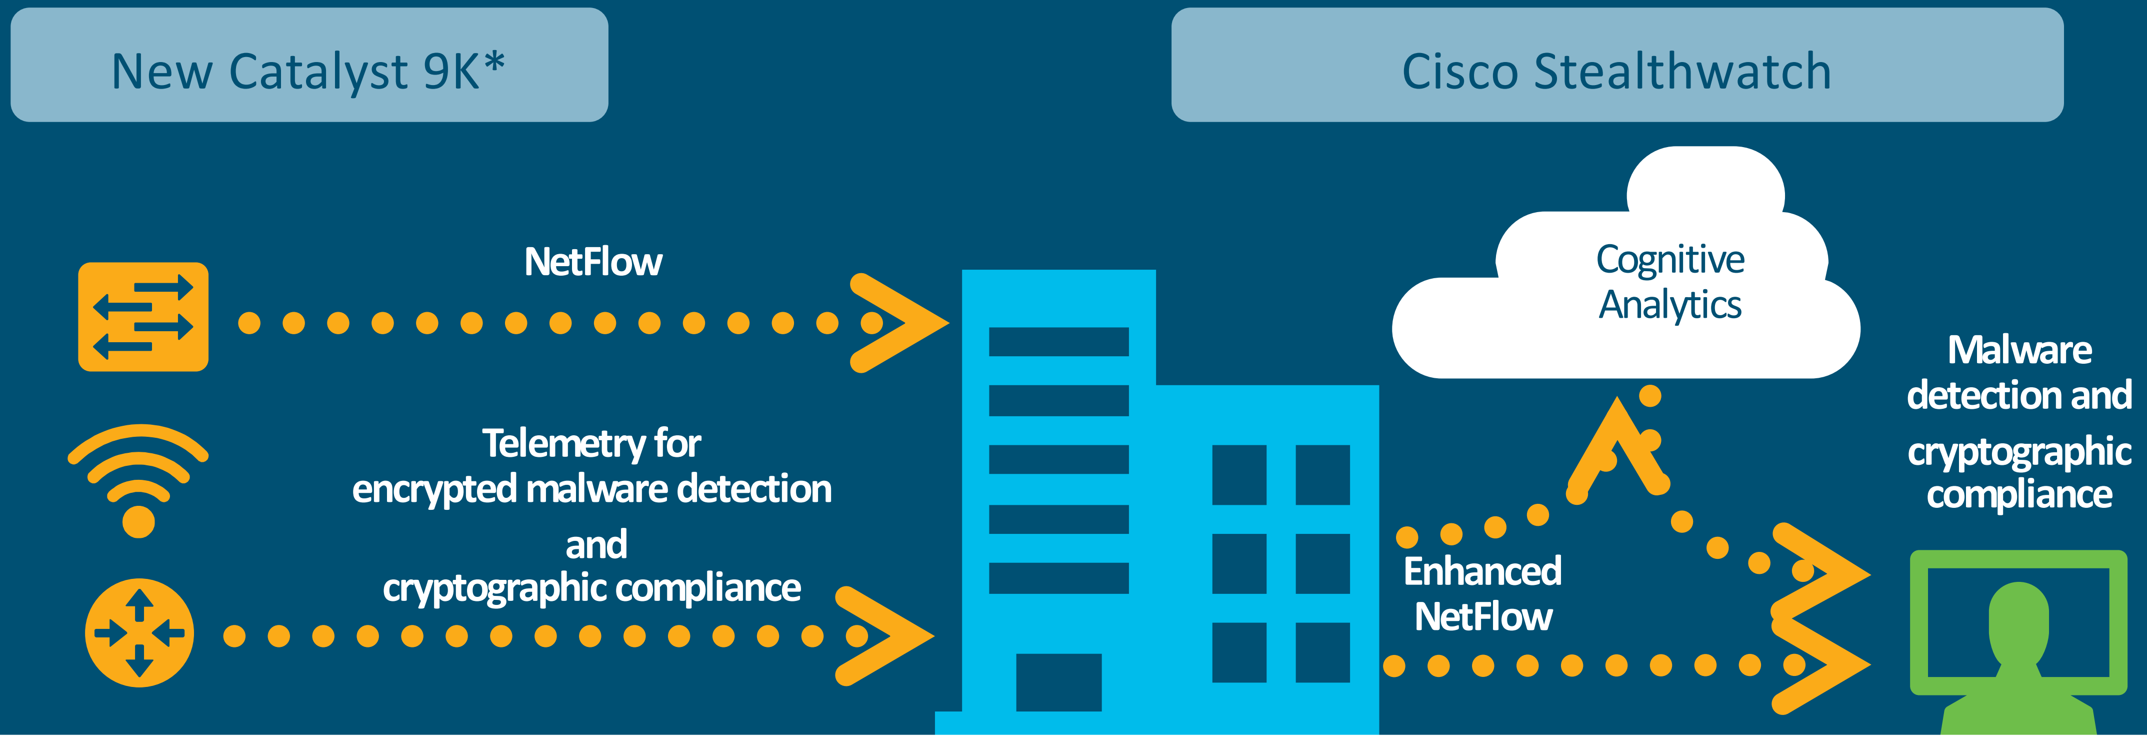
New Catalyst 9K*
Cisco Stealthwatch
NetFlow
Telemetry for encrypted malware detection and cryptographic compliance
Cognitive Analytics
Enhanced NetFlow
Malware detection and cryptographic compliance

Leveraged network	Faster investigation	Higher precision	Stronger protection
Enhanced NetFlow from Cisco’s newest switches and routers	Enhanced analytics and machine learning	Global-to-local knowledge correlation	Continuous enterprise-wide compliance

* Future support coming soon for ISR and ASR systems

Encrypted Traffic Analytics Efficiency

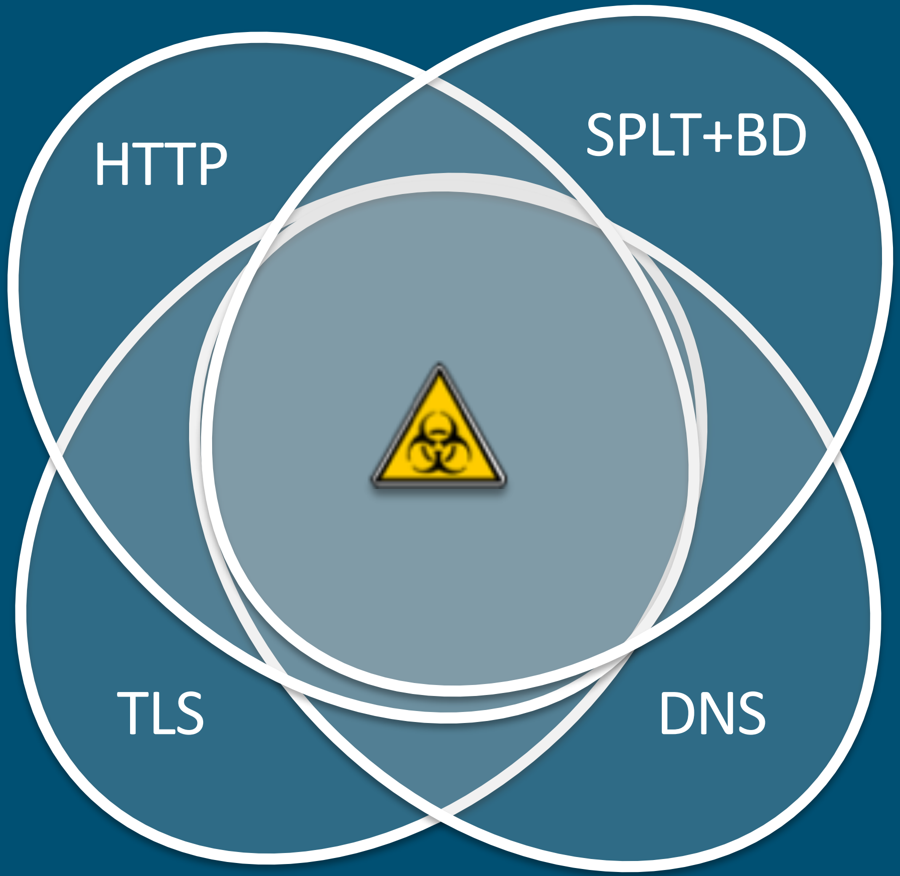

	Acc.	FDR
SPLT+BD+TLS+HTTP+DNS	99.993%	99.978%
SPLT+BD+TLS+HTTP	99.983%	99.956%
SPLT+BD+TLS+DNS	99.968%	98.043%
SPLT+BD+TLS	99.933%	70.351%
HTTP+DNS	99.985%	99.956%
TLS+HTTP	99.955%	99.660%
TLS+DNS	99.883%	96.551%
HTTP	99.945%	98.996%
DNS	99.496%	94.654%
TLS	94.836%	50.406%

Malicious Use of Legitimate Resources

Cybercriminals are adopting command-and-control channels that rely on legitimate Internet services, making malware traffic almost impossible to shut down

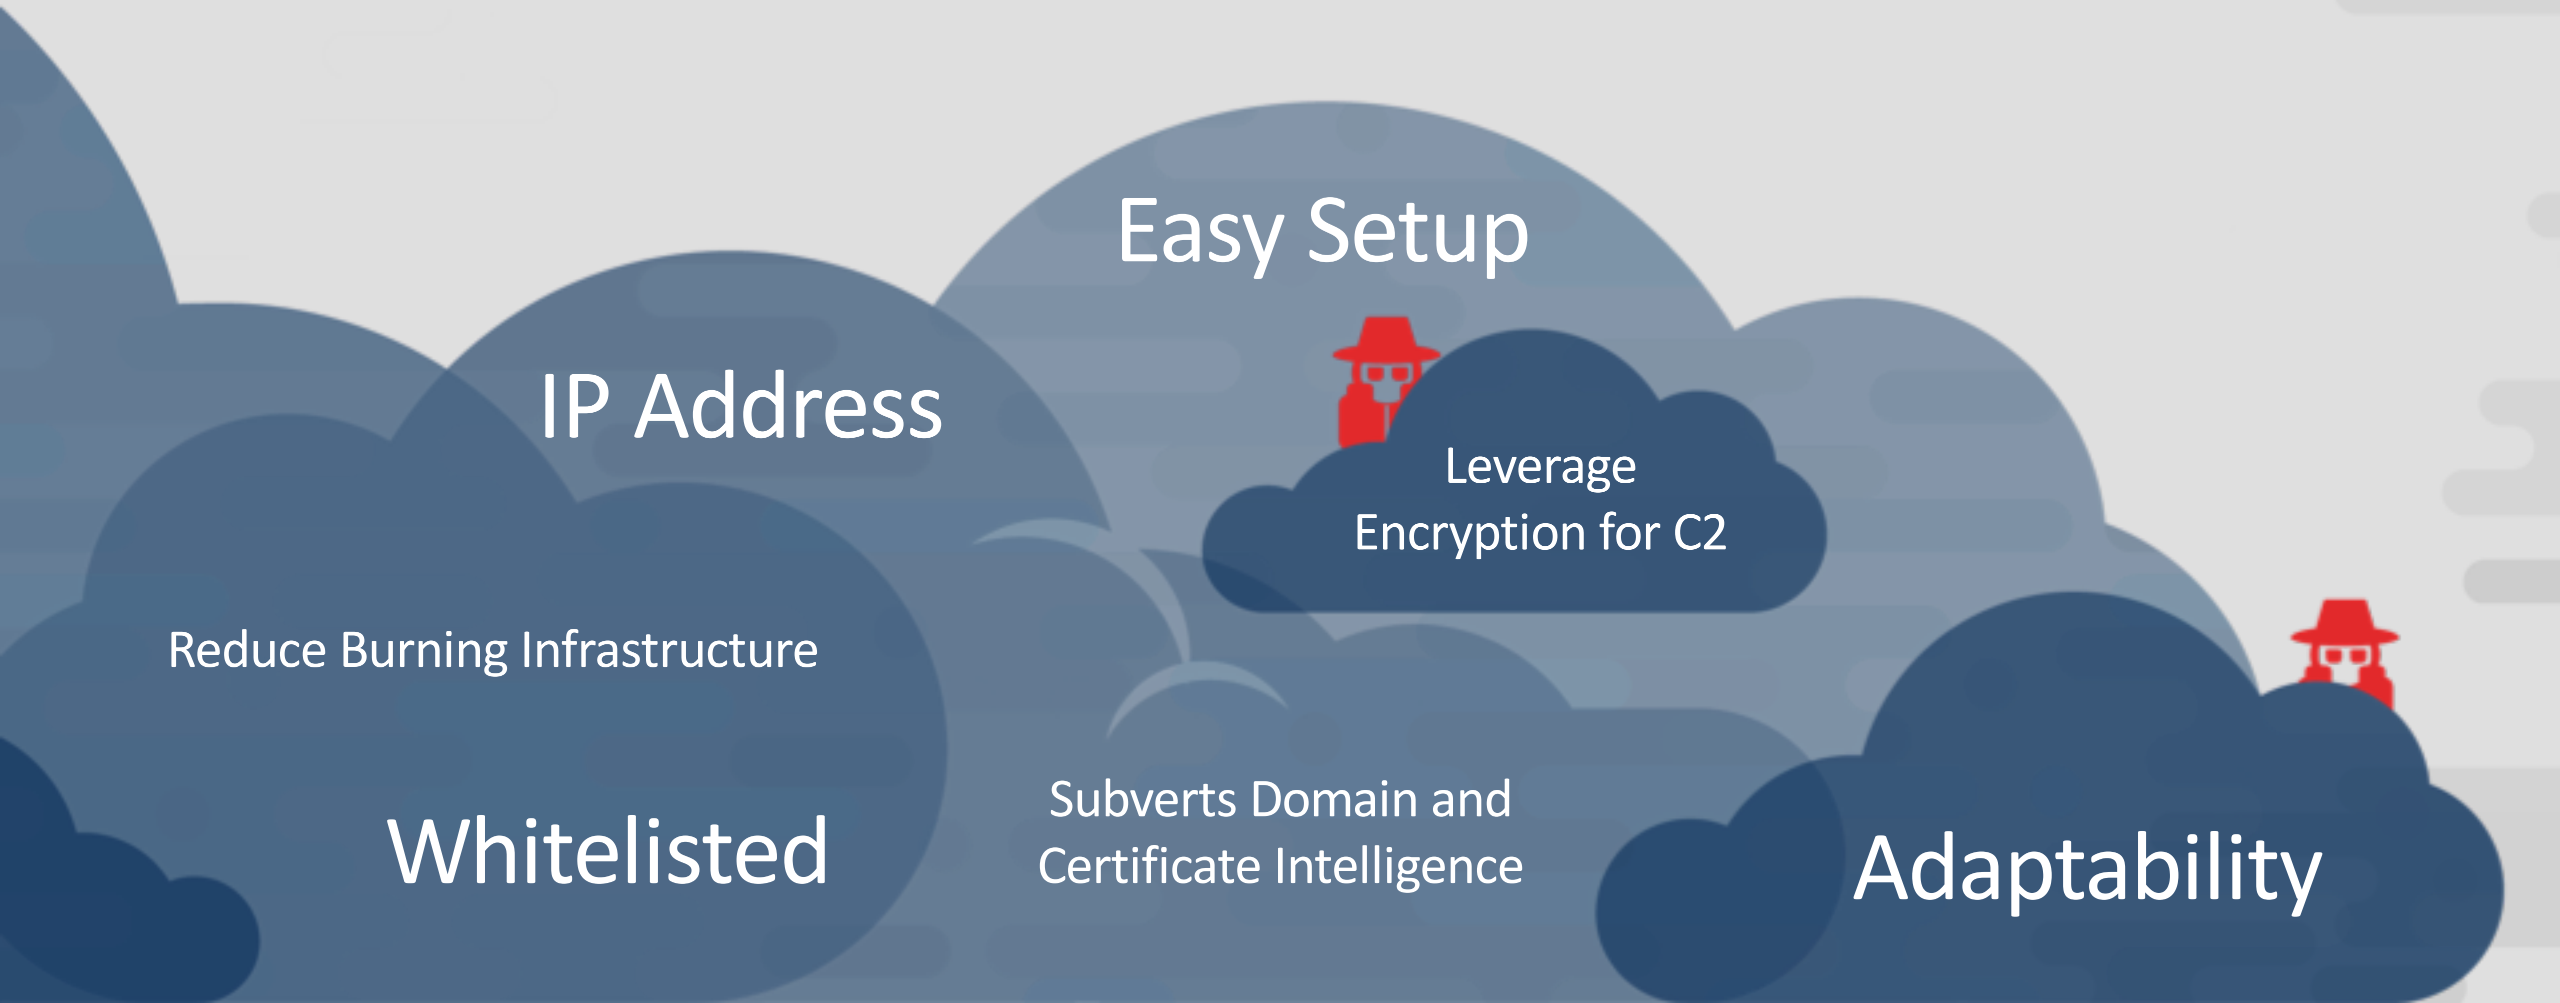

Easy Setup

IP Address

Leverage Encryption for C2

Reduce Burning Infrastructure

Whitelisted

Subverts Domain and Certificate Intelligence

Adaptability

Hackers don’t think about that

~600 features per single web request

- Generic – lengths, status codes, mime types
- HTTP – URLs, referrers, character distribution
- HTTPS – anomaly values, timings, context
- Global – domain/AS popularity
- External – whois, TLS certificates

Cisco Cognitive Threat Analytics

Early Detection with Artificial Intelligence

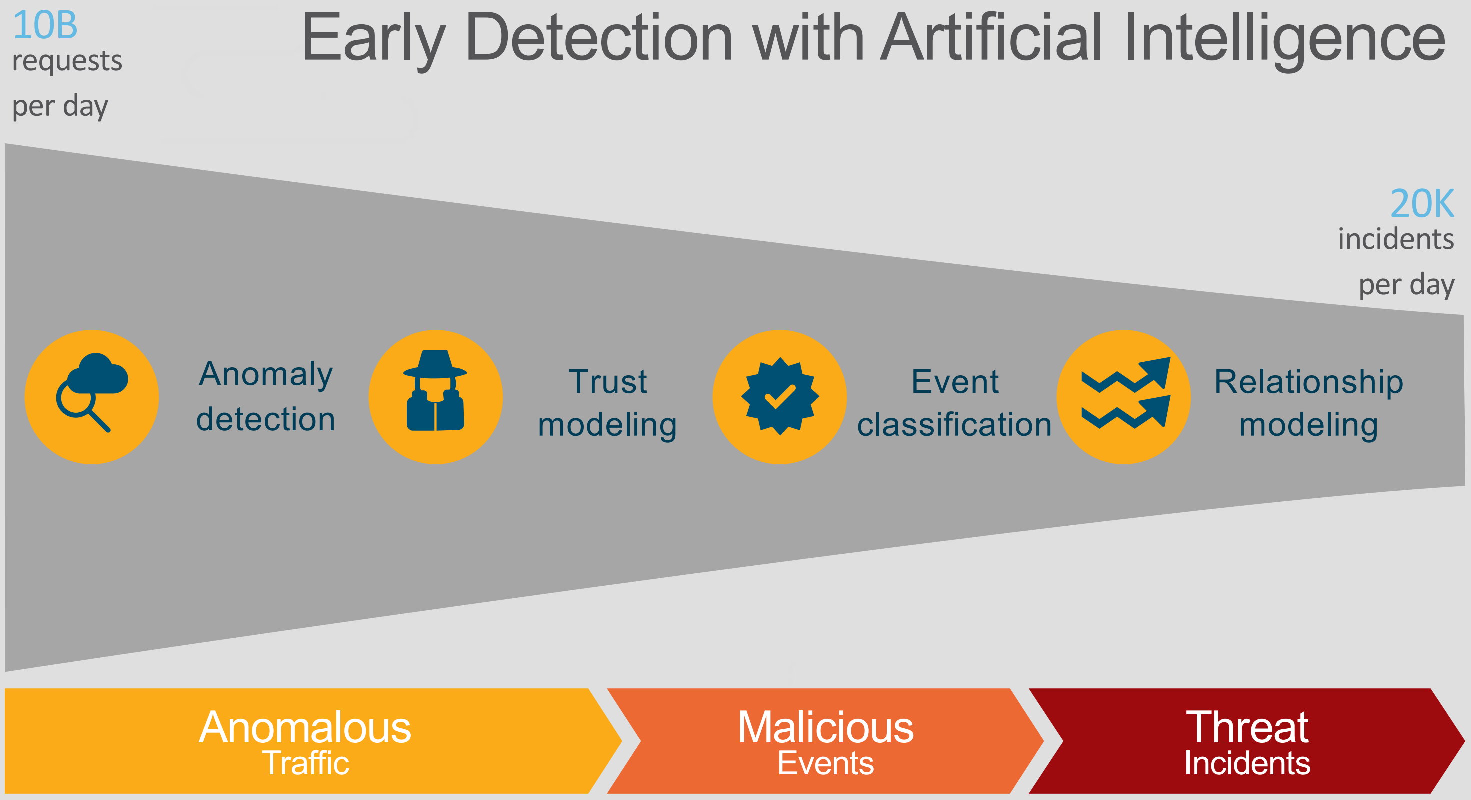

10B requests per day

Anomaly detection

Trust modeling

Event classification

Relationship modeling

20K incidents per day

Anomalous Traffic

Malicious Events

Threat Incidents

What Does CTA Typically Detect

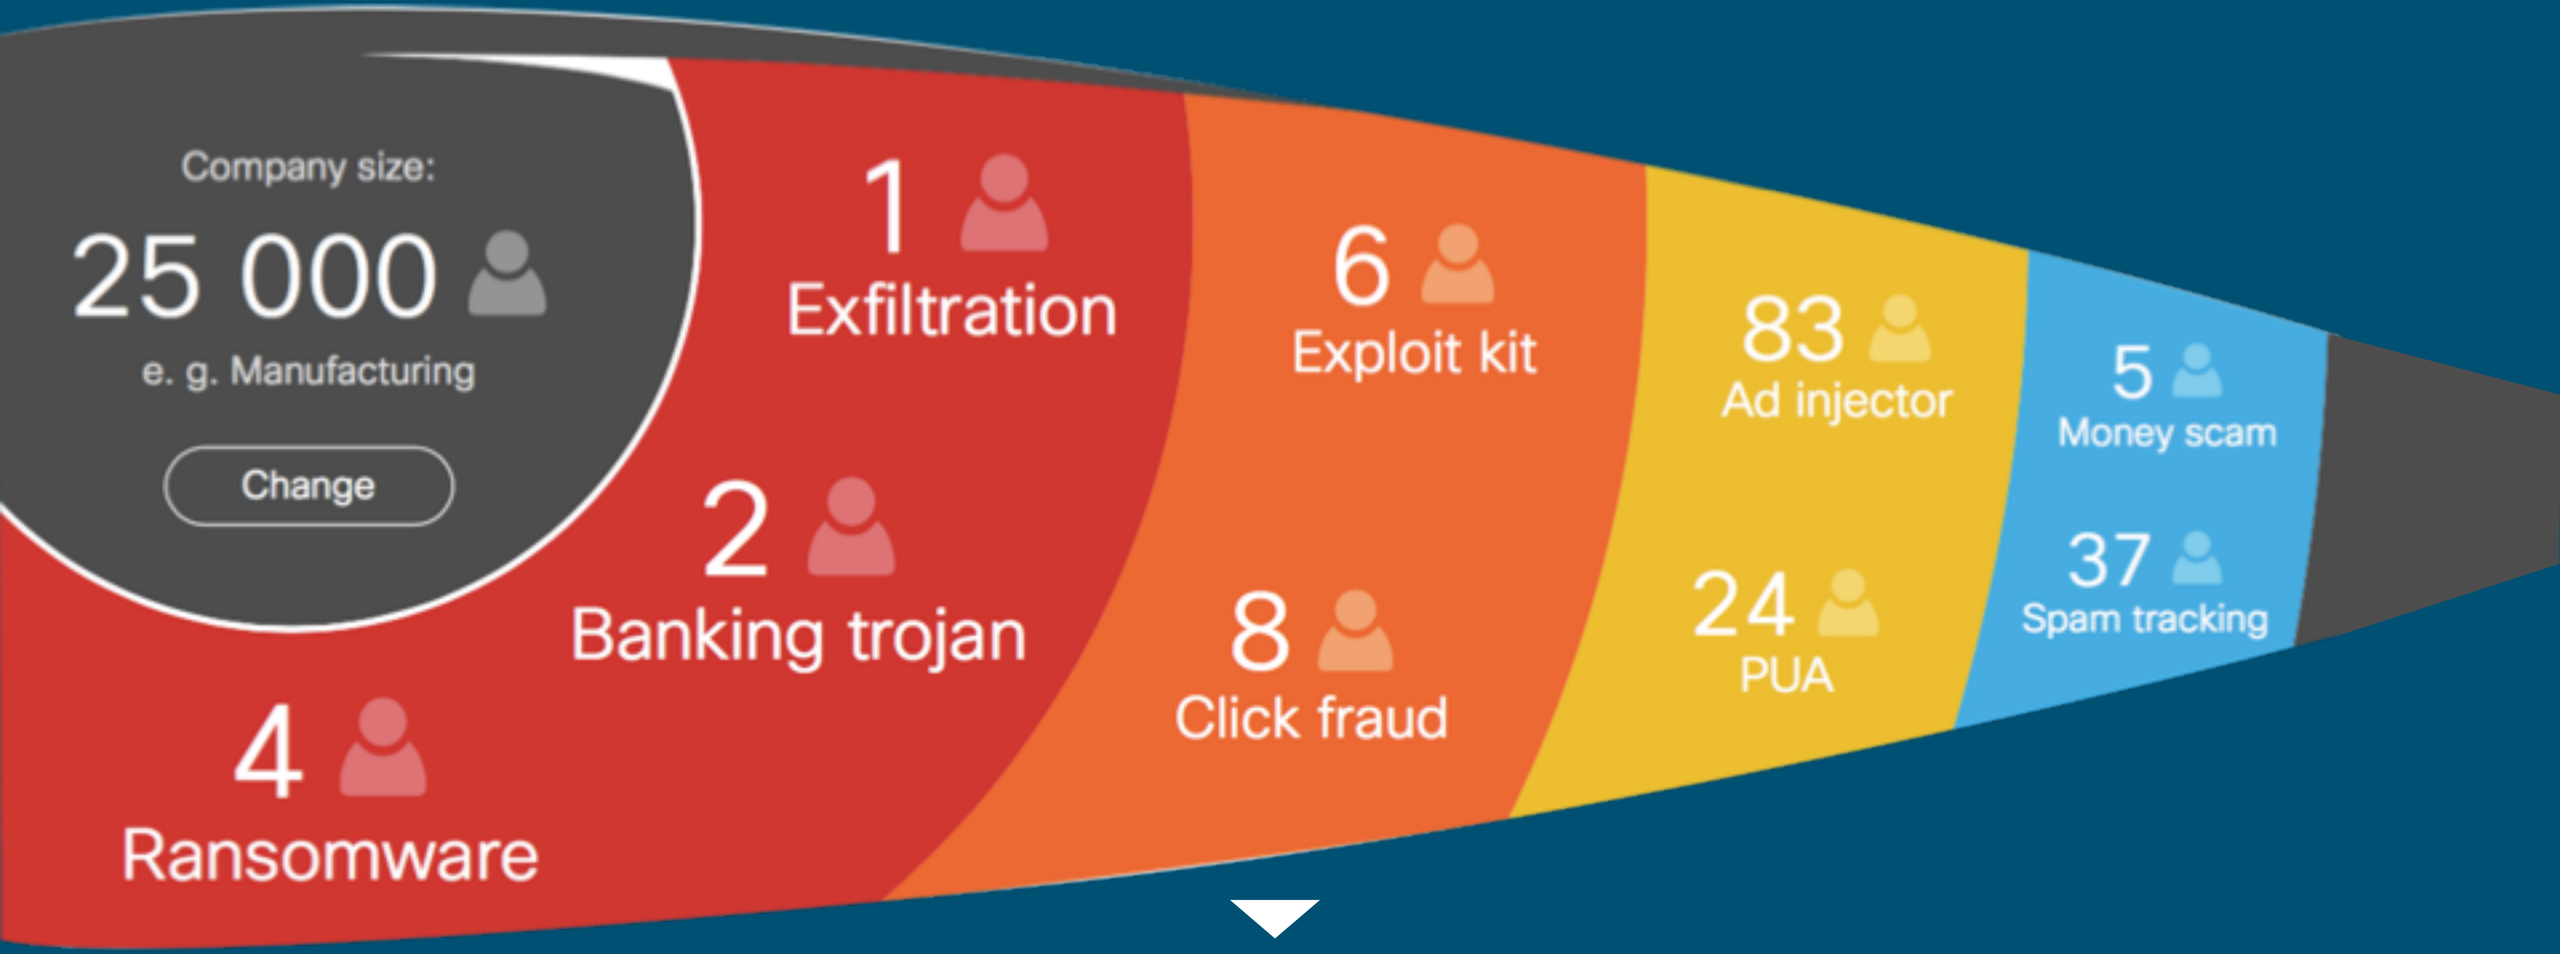

Company size:

25 000

e. g. Manufacturing

Change

1 Exfiltration

2 Banking trojan

4 Ransomware

6 Exploit kit

8 Click fraud

83 Ad injector

24 PUA

5 Money scam

37 Spam tracking

Sample report demonstrating an advanced threat visibility gap: http://cognitive.cisco.com/preview

Insider Threat

Machine learning algorithms can greatly help detect internal malicious actors

5200
docs per user

“Data”
was the most popular keyword in doc titles

PDFs
were the most common file type

62%
occur outside of normal work hours

High*
accuracy of malicious activity detection since June 2017

Compromised Cloud Account Detection by CloudLock

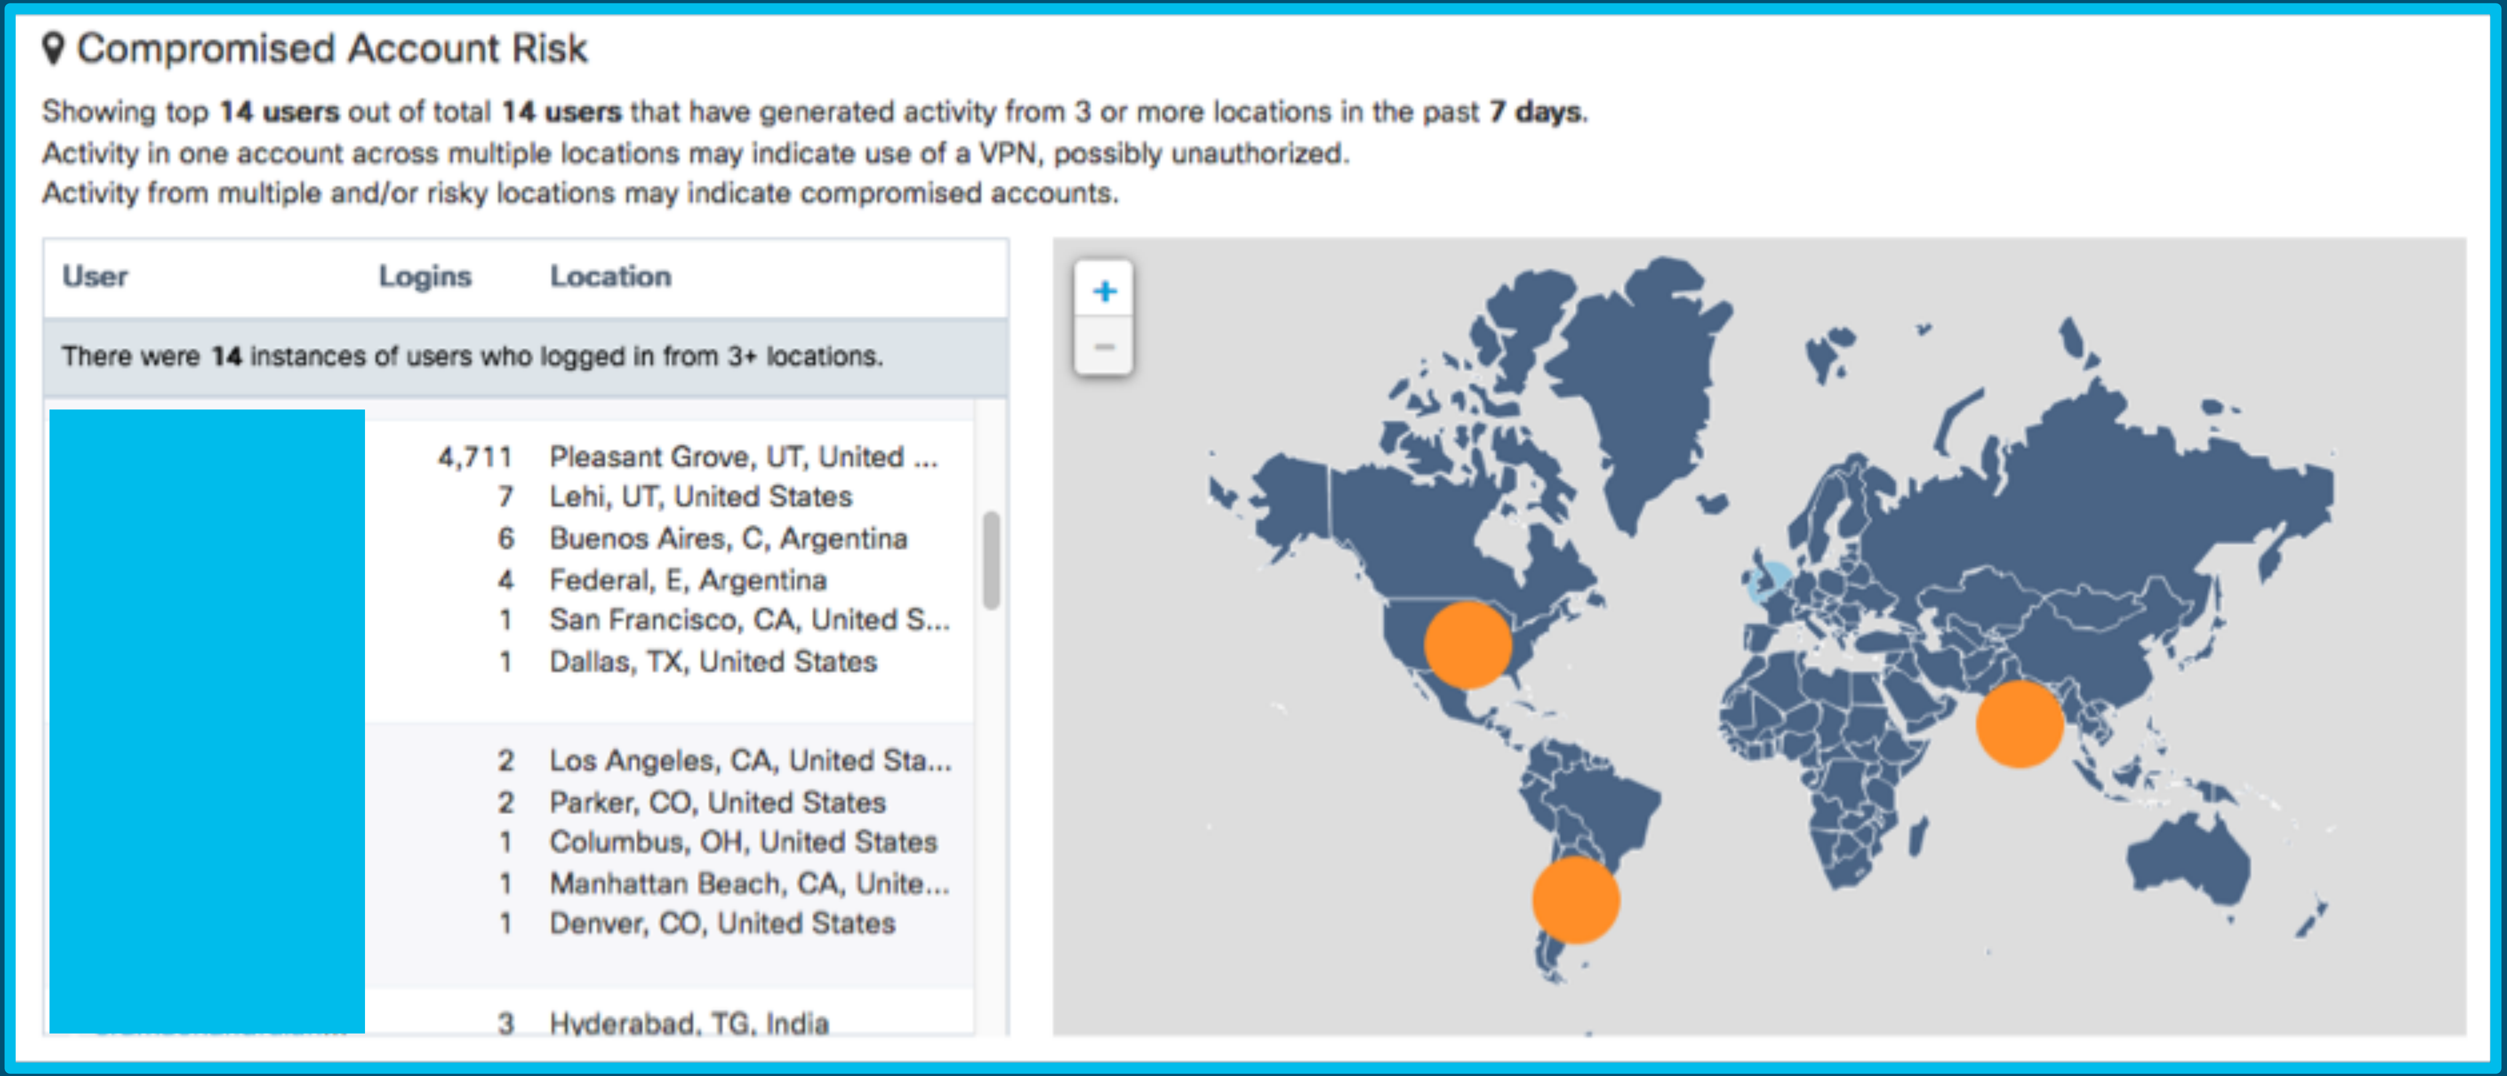

Compromised Account Risk

Showing top **14 users** out of total **14 users** that have generated activity from 3 or more locations in the past **7 days**.
Activity in one account across multiple locations may indicate use of a VPN, possibly unauthorized.
Activity from multiple and/or risky locations may indicate compromised accounts.

There were **14** instances of users who logged in from 3+ locations.

User	Logins	Location
[redacted]	4,711	Pleasant Grove, UT, United ...
	7	Lehi, UT, United States
	6	Buenos Aires, C, Argentina
	4	Federal, E, Argentina
	1	San Francisco, CA, United S...
	1	Dallas, TX, United States
[redacted]	2	Los Angeles, CA, United Sta...
	2	Parker, CO, United States
	1	Columbus, OH, United States
	1	Manhattan Beach, CA, Unite...
	1	Denver, CO, United States
[redacted]	3	Hyderabad, TG, India

How Malicious Actors Leverage Domains

Type of Attack

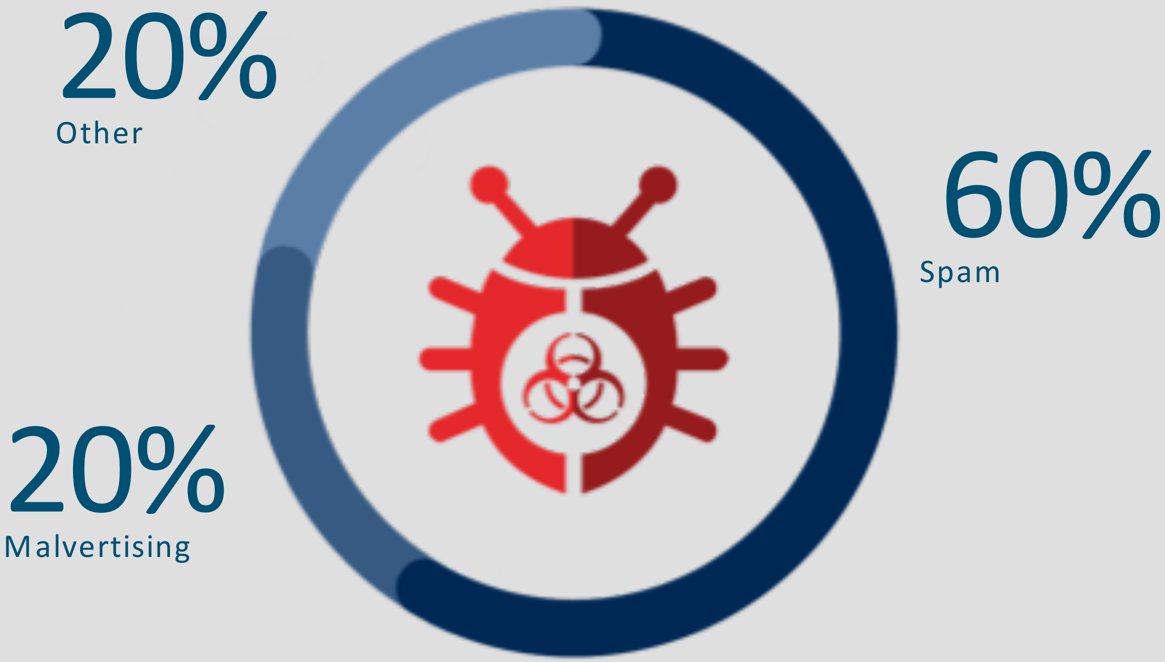

RLD Registered Times

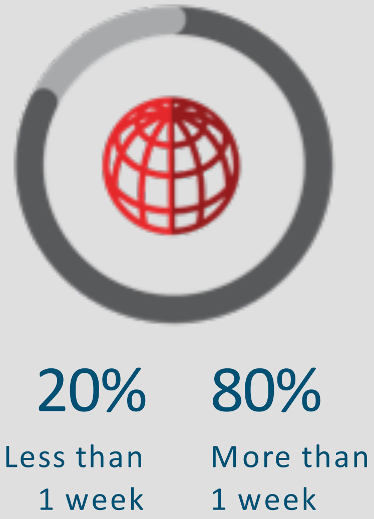

New or Reused Domains

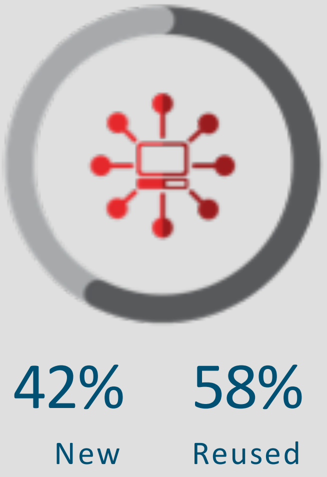

Organizations need to minimize access to malicious domains

Umbrella predictive models

2M+ live events per second

11B+ historical events

Guilt by inference

- Co-occurrence model
- Sender rank model
- Secure rank model

Guilt by association

- Predictive IP Space Modeling
- Passive DNS and WHOIS Correlation

Patterns of guilt

- Spike rank model
- Natural Language Processing rank model
- Live DGA prediction

Suspicious events in internal network

Source or target of malicious behavior	Reconnaissance	Command and Control	DDoS Activity	Insider threats
Scanning, excessive network activity such as file copying or transfer, policy violation, etc.	Port scanning for vulnerabilities or running services	Communication back to an external remote controlling server through malware	Sending or receiving SYN flood and other types of data floods	Data hoarding and data exfiltration

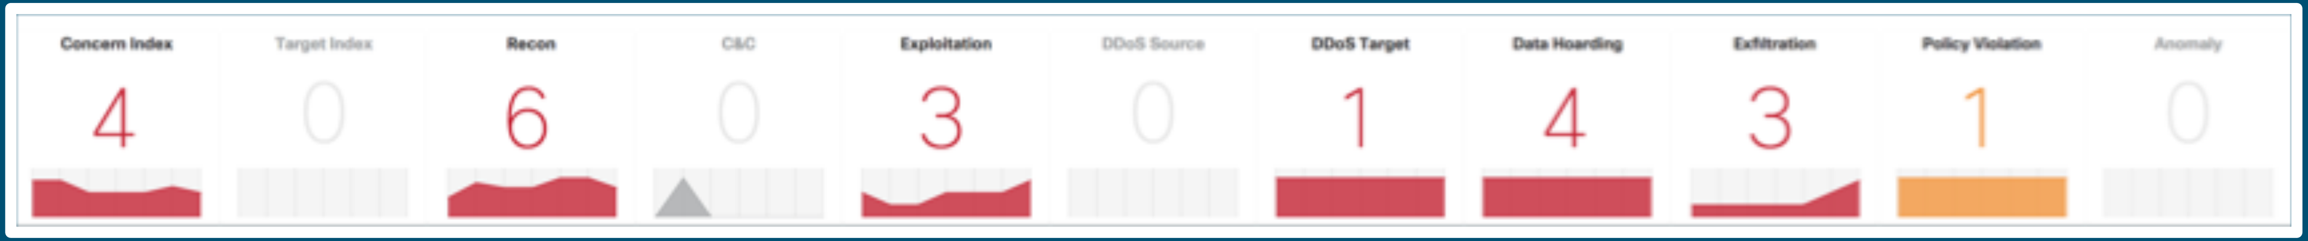

Concern Index	Target Index	Recon	C&C	Exploitation	DDoS Source	DDoS Target	Data Hoarding	Exfiltration	Policy Violation	Anomaly
4	0	6	0	3	0	1	4	3	1	0

Stealthwatch = netflow security brain

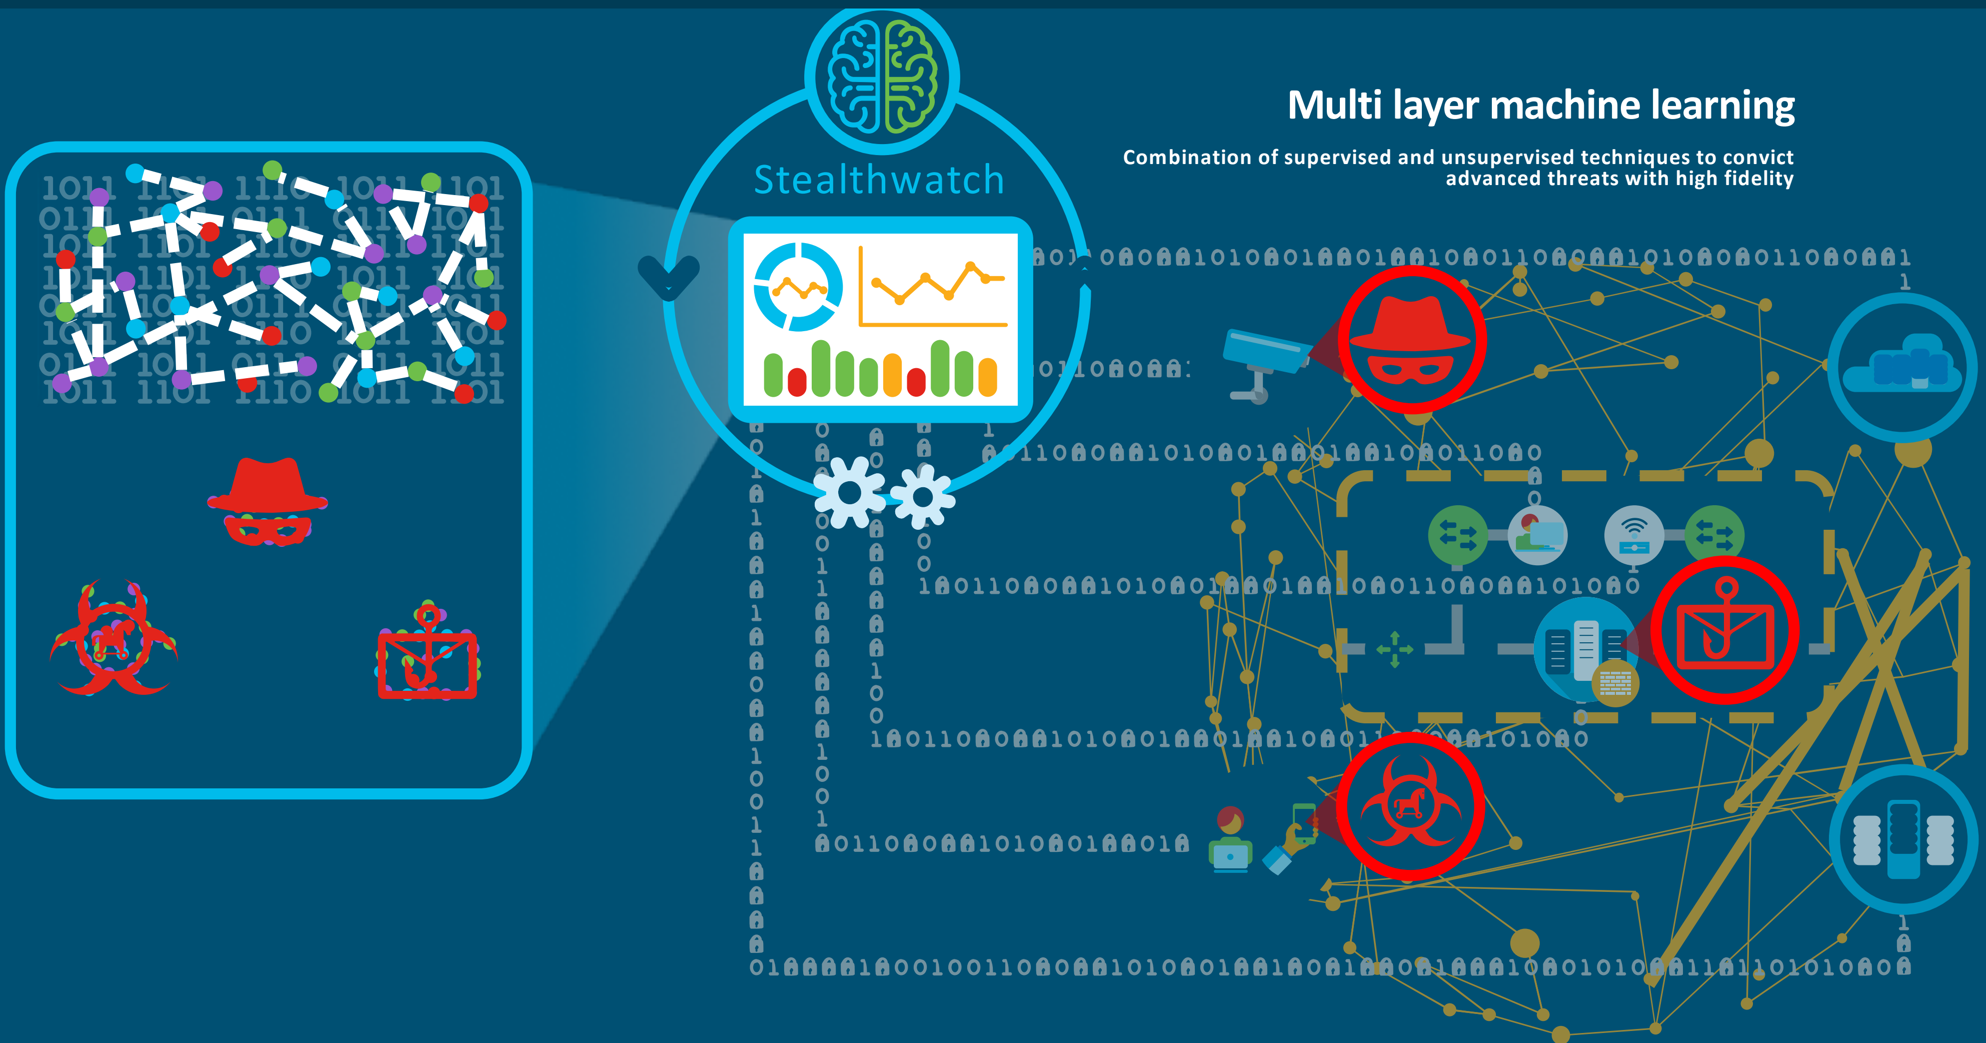

Multi layer machine learning

Combination of supervised and unsupervised techniques to convict advanced threats with high fidelity

Stealthwatch = netflow security brain

Global Threat Intelligence

TALOS

Stealthwatch

Global Risk Map

Intelligence of global threat campaigns mapped to local alarms for faster mitigation

Advanced detection using entity modeling

Comprehensive data Set

Optimized to remove redundancies and improve performance

Netflow, IPFIX, sFlow as well as other layer 7 protocols

Security events

~100 heuristics to detect anomalies and known bad behavior

Addr_Scan, Beaconing Host, Brute Force Login, Max Flows Initiated, Suspect Data Hoarding, Suspect Data Loss

Alarm categories

High-risk, low-noise alerts for faster response

Concern, Recon, C&C, Exploitation, DDoS, Data Hoarding, Exfiltration, Policy Violation

Power of multi-layer machine learning

Increase fidelity of detection using best-in-class security analytics

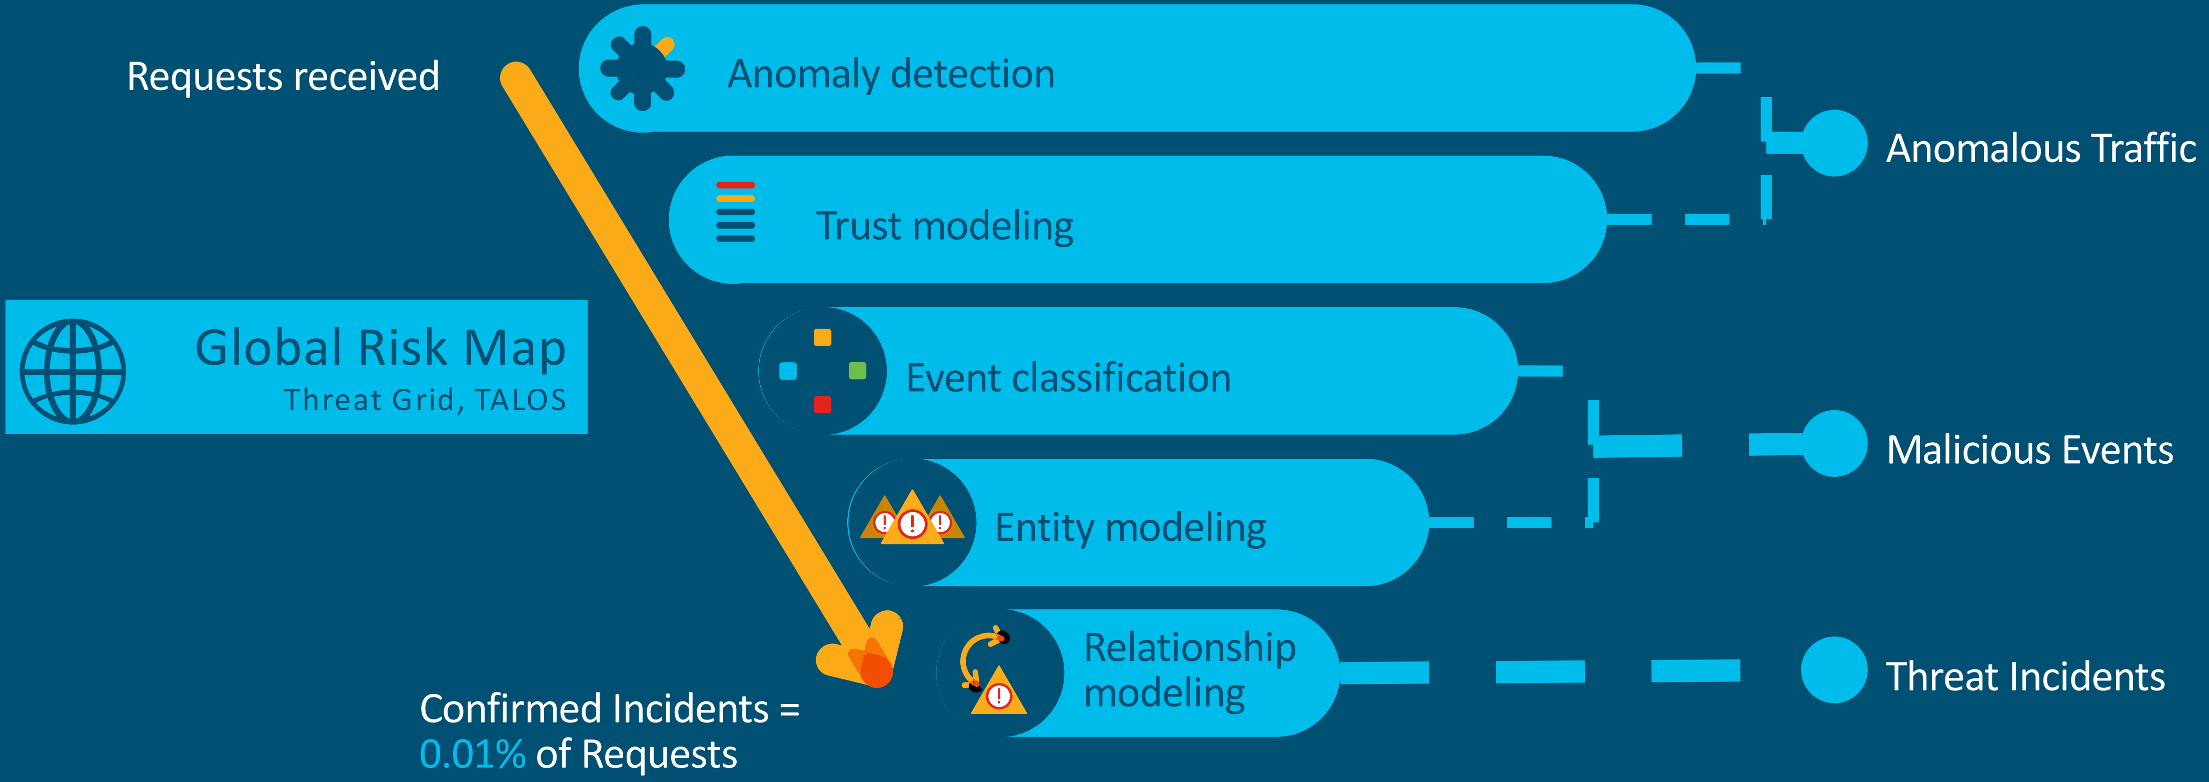

Requests received

Anomaly detection

Trust modeling

Event classification

Entity modeling

Relationship modeling

Anomalous Traffic

Malicious Events

Threat Incidents

Global Risk Map

Threat Grid, TALOS

Confirmed Incidents = 0.01% of Requests

Endpoints continue to be the primary point of entry for breaches

70% of breaches start on endpoint devices

WHY?

Gaps in protection

65%

of organizations say attacks evaded existing preventative tools

User error

48%

of attackers bypass endpoint defenses because of user error

Gaps in visibility

55%

of organizations are unable to determine cause of breach

100

DAYS

industry average time to detection

The AMP Cloud Prevention Framework

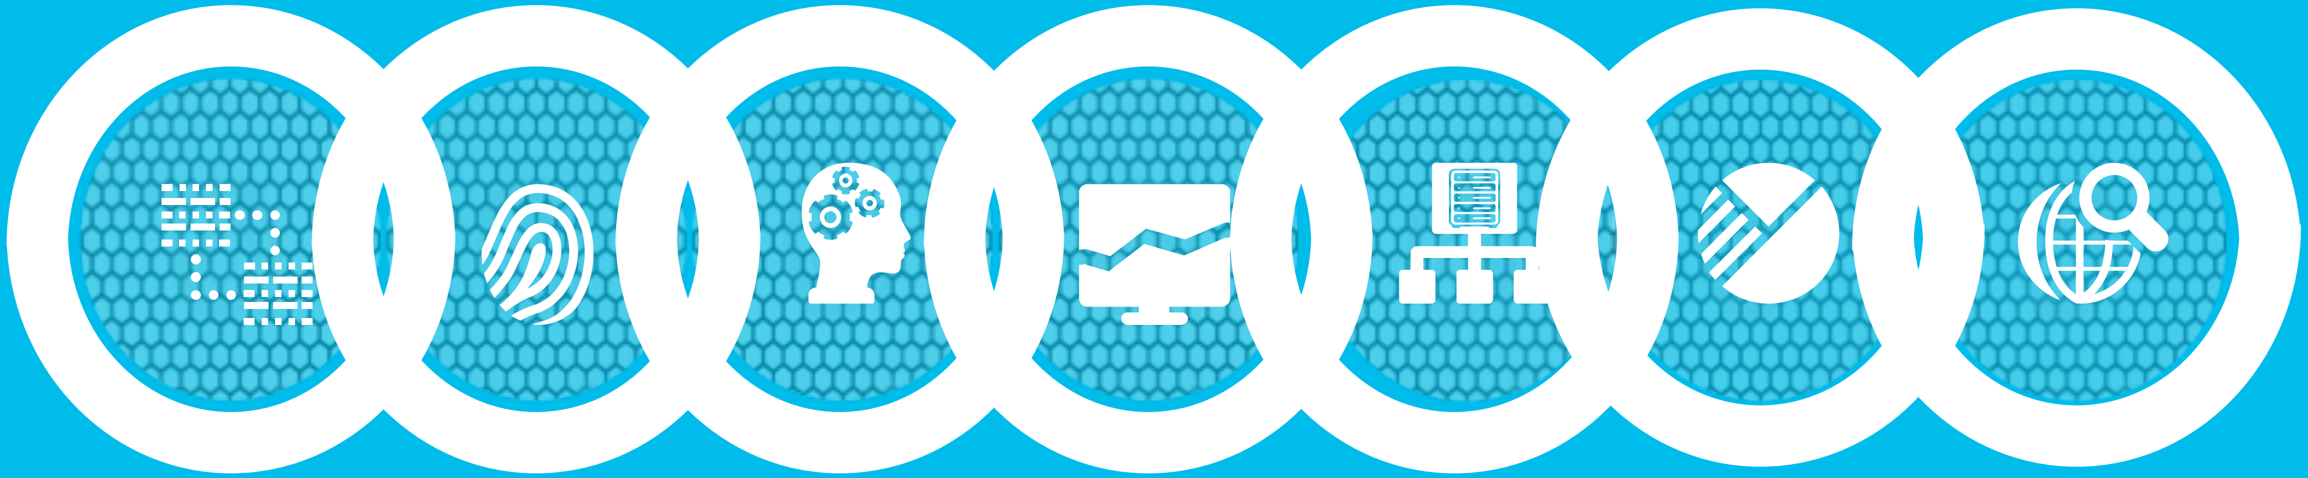

1-to-1 Signatures

Ethos

Spero

IOCs

Device Flow Correlation

Advanced Analytics

Dynamic Analysis

Spero Engine in Cisco Advanced Malware Protection

- Machine Learning
  - Automatically constructs a framework
  - Needs data to learn/adjust
  - Requires large sets of good data
- Behaviour modeling
  - Discover patterns better than human analysts
- 0-day insight is the goal

Spero: A machine-learning based technology that proactively identifies threats that were previously unknown

Uses active heuristics to gather execution attributes

Needs good data in large sets to tune

Built to identify *new malware*

Conclusion

Market Expectations: Threat Landscape

The threat landscape to remain complex and challenging

- Few predict radically new threats on the horizon, but they see more capable and more diabolical bad actors
- Believe they'll need ever more sophisticated security arsenals to keep they at bay

Market Expectations: Modern Workplace

The modern workplace will continue to create conditions that favor the attackers

- The footprint security executives must secure continues to expand
- Employees increasingly carry their work (and the company's data) with them wherever they go—a well-documented source of exposure
- Clients, partners and suppliers all need secure access to corporate resources
- With the increasing deployment of IoT sensors, etc., companies' interfaces to the internet will multiply dramatically

Effective security depends on total visibility

KNOW
every host

SEE
every conversation

Understand what is
NORMAL

Be alerted to
CHANGE

Respond to
THREATS quickly

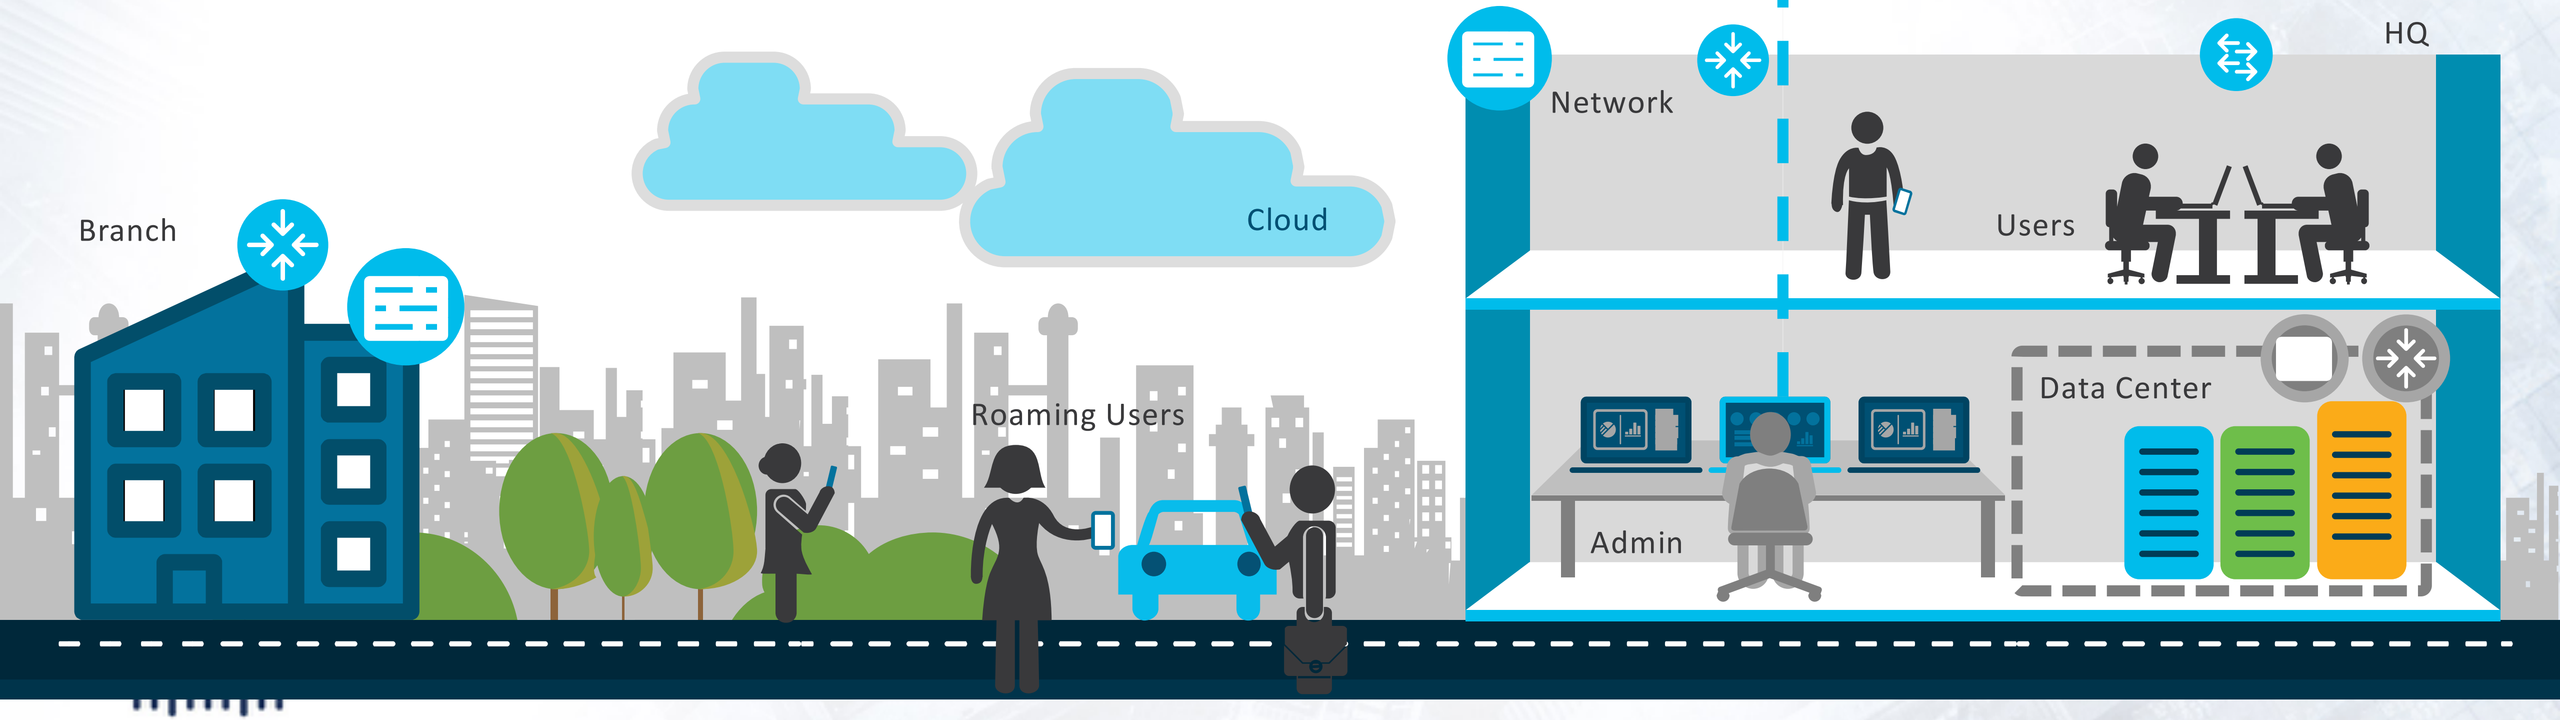

Market Expectations: AI and Machine Learning

More spending on AI/ML capabilities

- AI, ML and automation increasingly desired and expected
- 83%: Reliant on automation to reduce level of effort to secure the organization
- 74%: Reliant on AI to reduce level of effort to secure the organization
- CISOs expect to take increasing advantage of AI and robotics
- 92% of security professionals say behavior analytics tools work well in identifying bad actors

AI in cyber security isn’t panacea but future

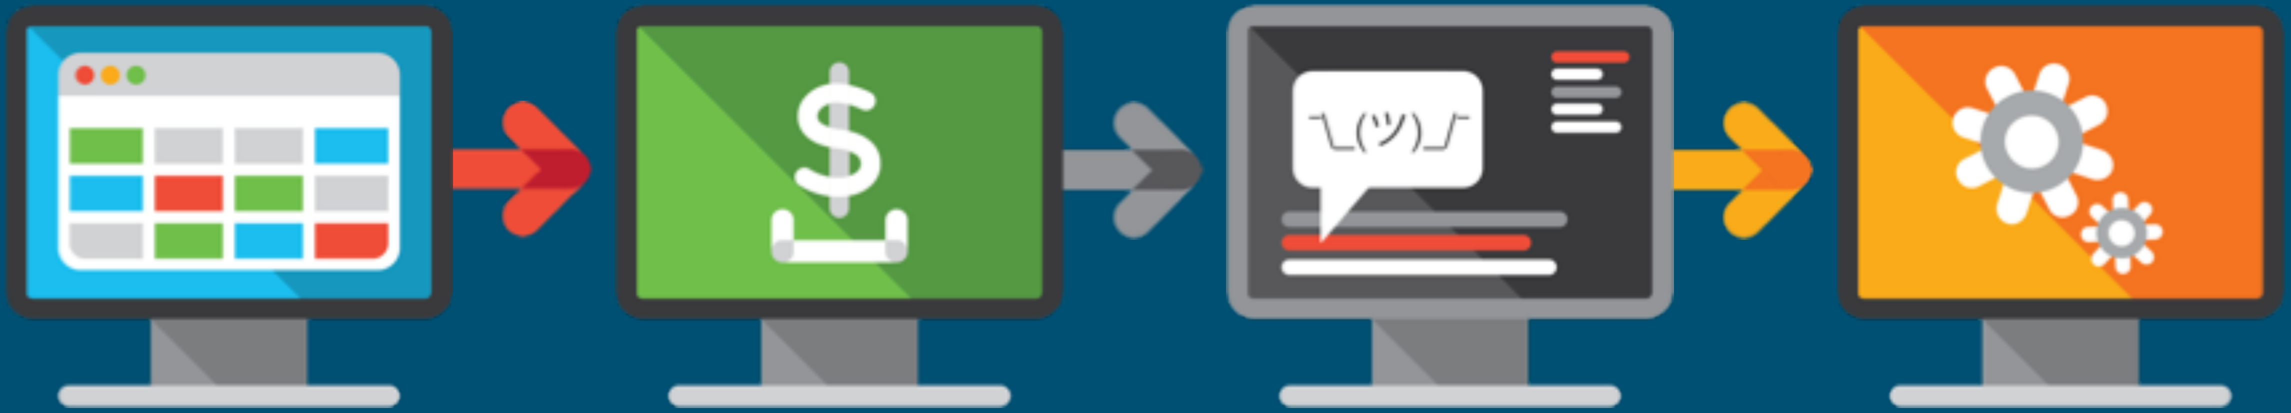

Signatures and IoC

IDS, AV, NGIPS, EDR, TIP

Rules

NGFW, WSA, SIEM, ESA

Statistical models

Netflow

AI algorithms

How We Disservice AI in Cyber Security

Silver Bullet Marketing

No Explanation or Discussion

Limited Guidance

Cisco Internal Infosec AI-based solutions

Cisco iCAM – Intelligent Context Aware Monitoring (UEBA + DLP)

40 Billion
Cisco files being protected

16,000+
servers are monitored

10 seconds
to detect risk

Users-to-Ops
100,000 : 1

Cisco TIP – internal Big Data Threat Intelligence and Security Analytics Platform

2,2 PB
Hadoop Cluster

2848
cores

27 TB
Memory

200 TB
ElasticSearch Cluster

⚠ NERD ALERT

We can’t sell these solutions!

References for Cisco Cyber Security & Machine Learning

- https://www.cisco.com/go/security
- https://www.talosintelligence.com
- https://blogs.cisco.com/tag/machine-learning
- http://www.cisco-ai.com
- You can test all of our Cisco Security Solutions

NERD ALERT

Thank you!