

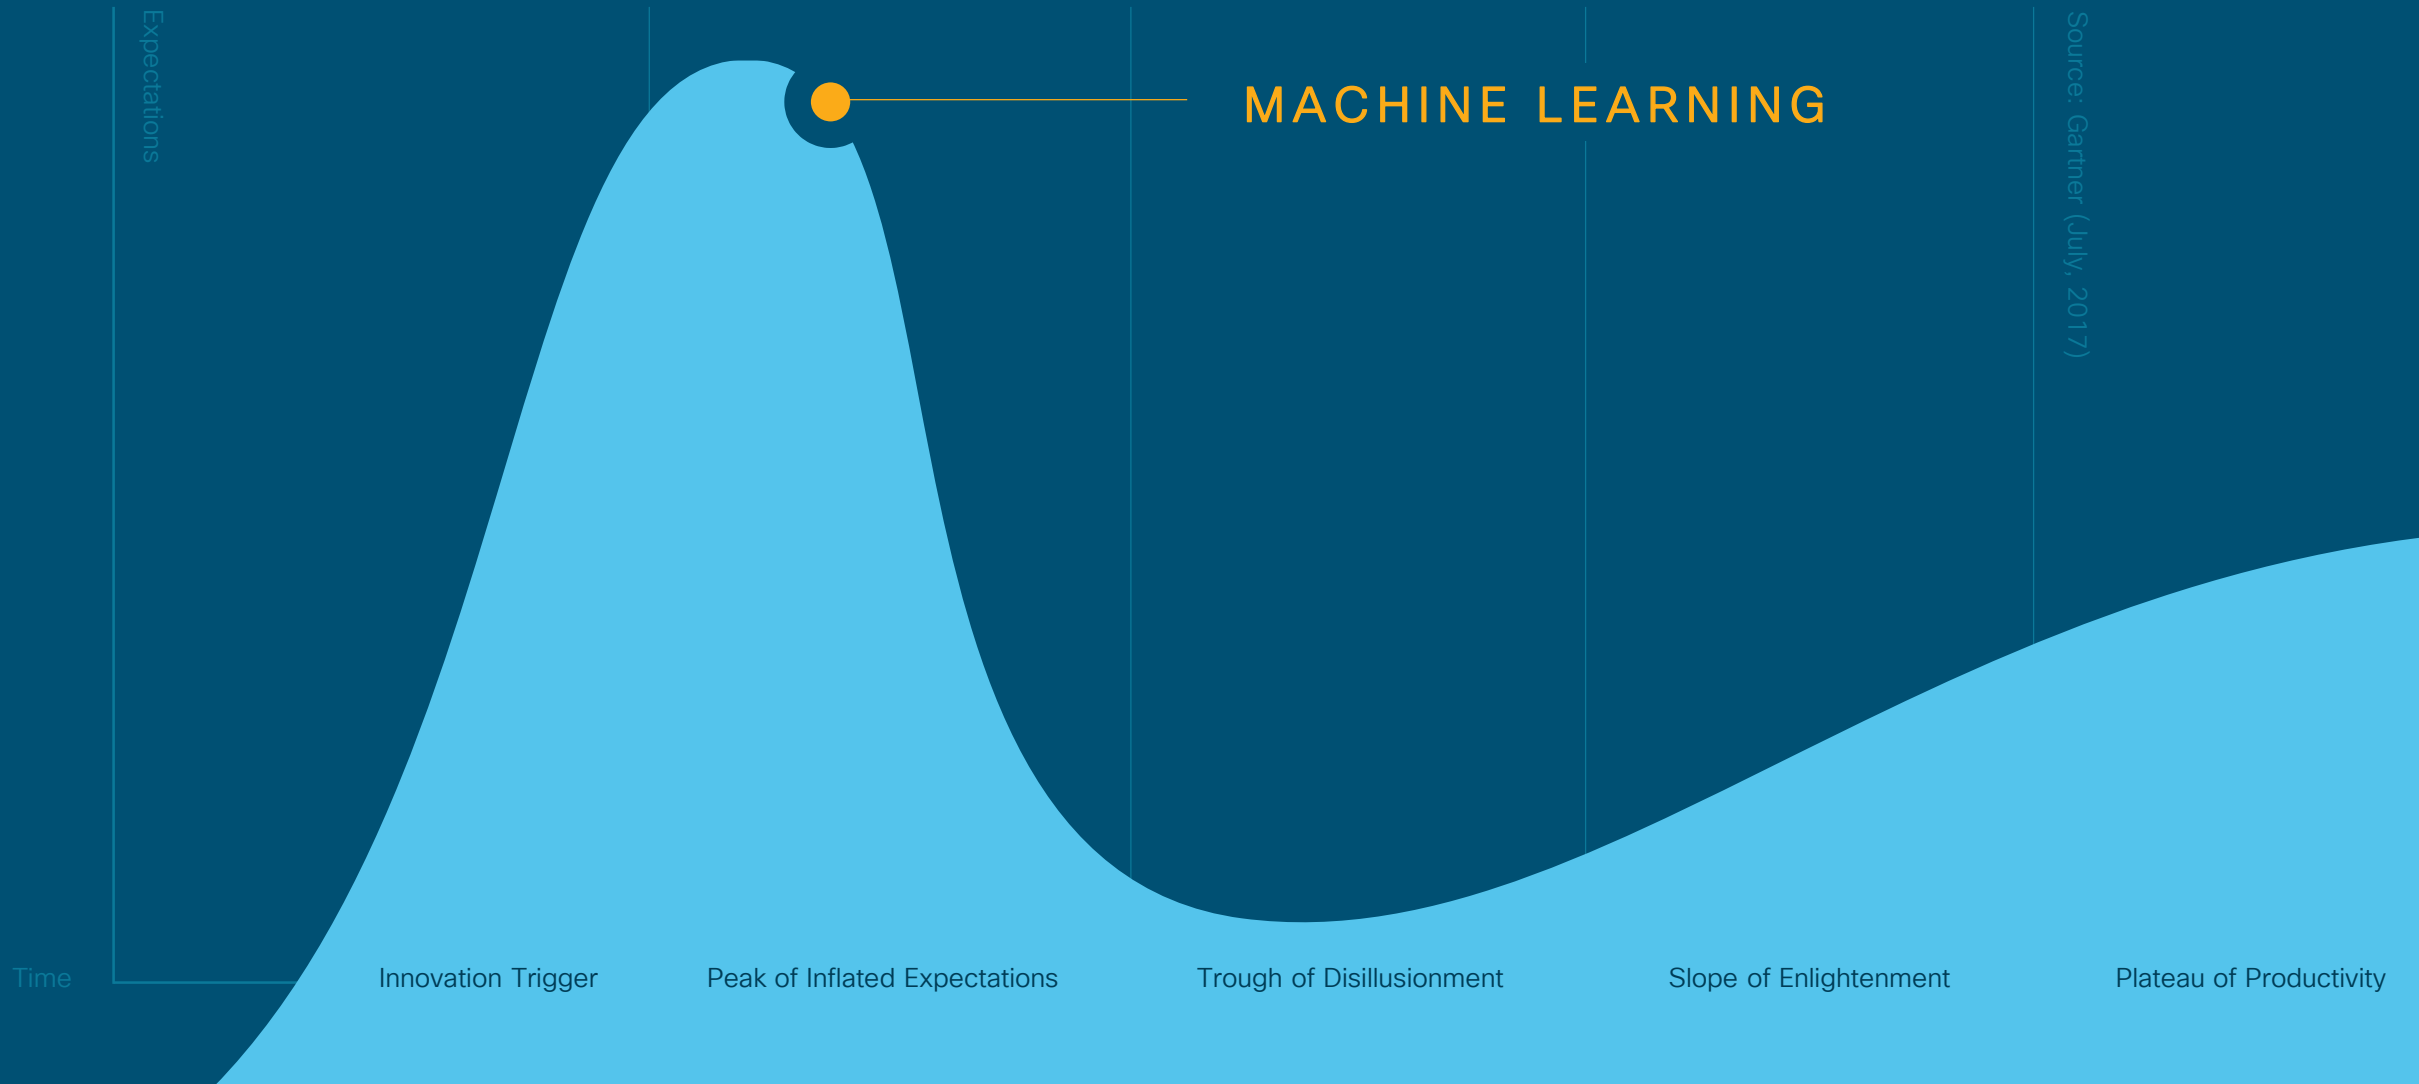


Machine Learning

And the Peak of Inflated Expectations

TK Keanini
Distinguished Engineer
June 2018

Gartner Hype Cycle for Emerging Technologies | 2017



Vendors Got Us Here



“Advanced Threats
are no match for A.I.”

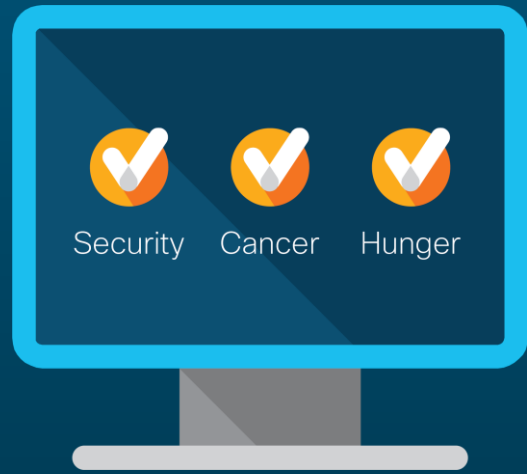


“Our machines detect
threats others cannot”



“100% predictive”

How We Disservice Machine Learning



Silver Bullet
Marketing



No Explanation
or Discussion



Limited Guidance



MACHINE LEARNING

What it is

“Field of study that gives computers the ability to learn without being explicitly programmed.”

Arthur Samuel's definition of machine learning in 1959

instance based
regularization
clustering
ensemble
bayesian
rule system
ground truth
machine learning algorithms
classifier
regression
deep learning
neural network
decision tree
dimensionality reduction


 **NERD ALERT**

Let's define the helpful data science terms



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



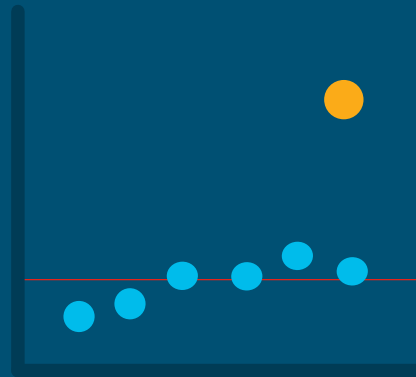
Other
(Reinforcement Learning, etc.)

What did we do before Machine Learning?

Use in combination with Machine Learning



Simple Pattern
Matching



Statistical Methods



Rules and First
Order Logic (FoL)

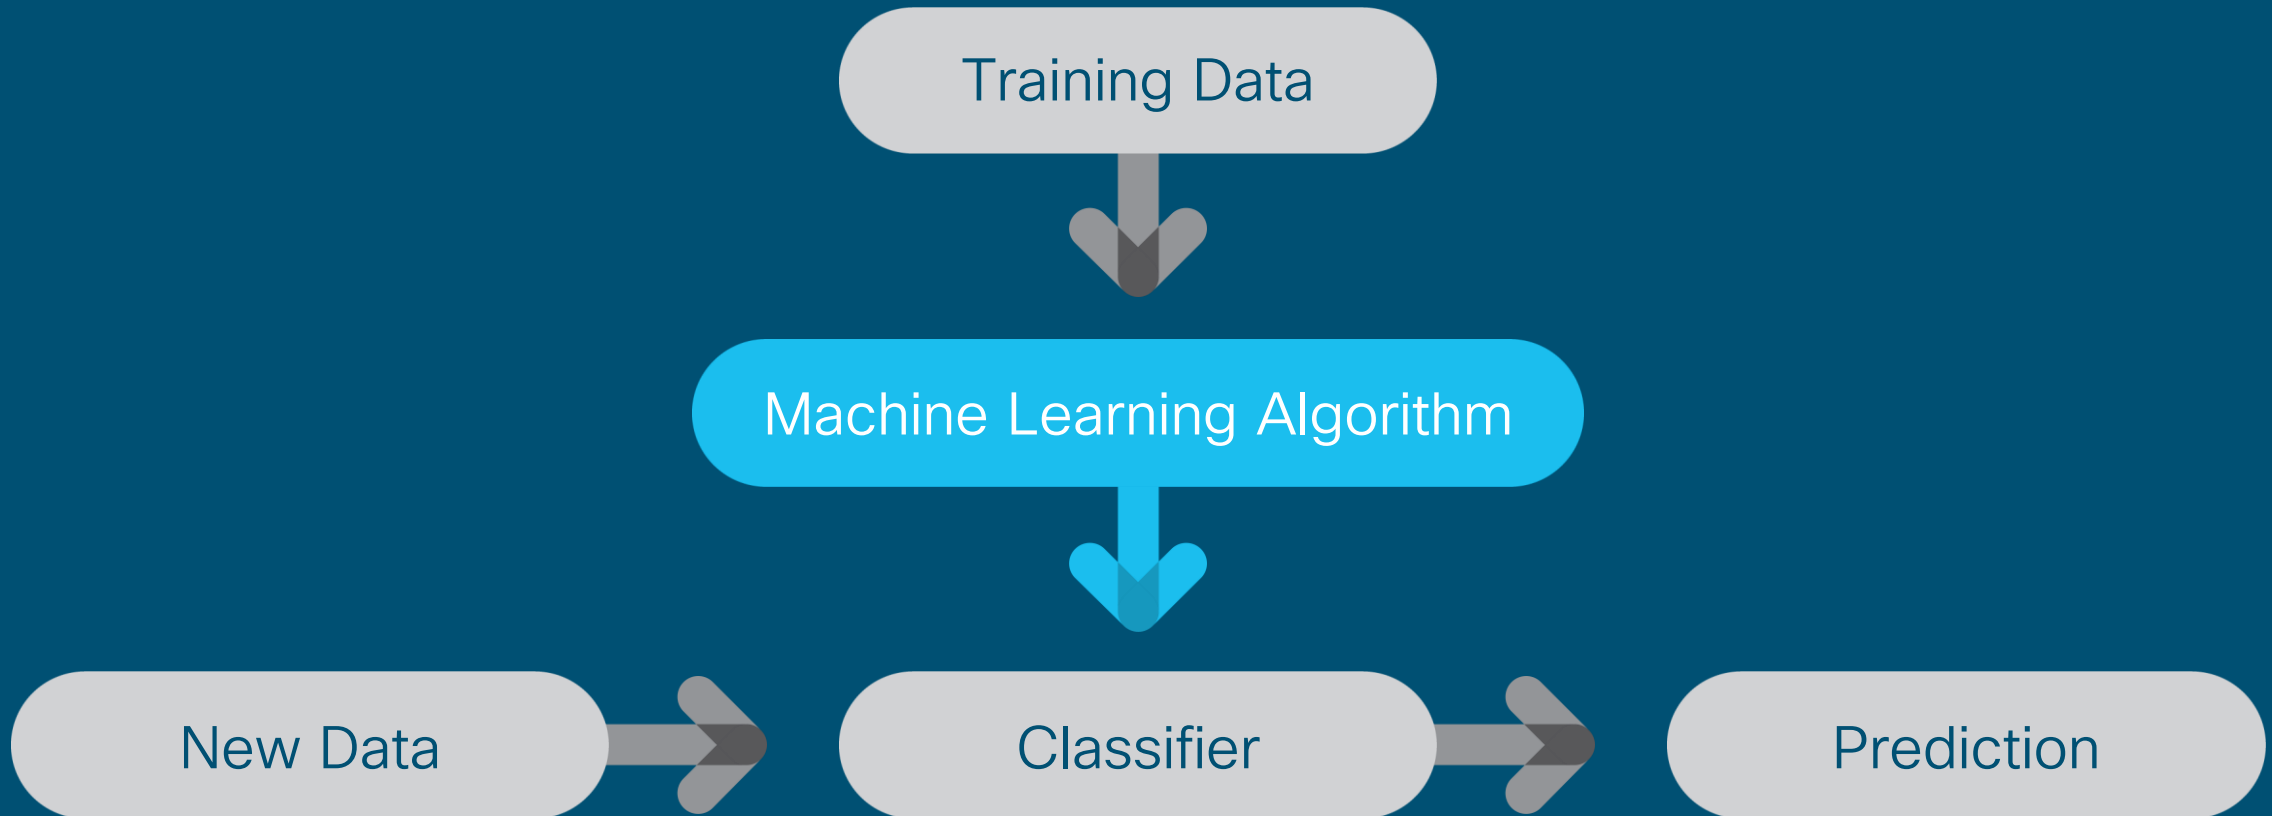
MACHINE LEARNING
Techniques

“Field of study that gives computers the ability to learn without being explicitly programmed.”

Translation

*“Field of study that gives computers **the ability to be implicitly programmed.**”*

Training Classifiers





Ground Truth Used in Supervised Learning

The '**Ground Truth**' is the pairing of example questions and answers.

If you can phrase a problem as *'we know this is right, learn a way to answer more questions of this type'*.

Success depends greatly on the dataset expressing the Question -> Answer mapping.

MACHINE LEARNING

Pitfalls

One Size Does Not Fit All

Other ML Application \neq Security



 NERD ALERT

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



What Is At Stake Matters

Because you watched Deadpool, you might like...



Deadpool



X-Men: First Class



The Flash




Captain America:
The First Avenger



How did you come to that conclusion?

“The Explainability Problem”



Normal Workflow

CFO daily calendar

Irregular Activity

ML detects “suspicious” activity and suggests remediation

Quarantined

However, ML cannot articulate *why* it wants to remediate

Loss of time and resources

MACHINE LEARNING
For Security

How We Know Machine Learning is Working



Accuracy

How often does my classifier give me the correct answer?



Precision

When my classifier predicts an instance in a certain class, how often does the instance belong to that class?



NERD ALERT

Root mean square error & Logical Regression

Translation: On average, how far away are my predictions from what we later know to be true values?



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

Insider Threats and Behavioral Security Analytics



Attackers

They're not breaking in,
they are logging in



Detecting

Through novelty and outliers



Events

Turn weak signals into a
strong one

Classify the Observable World and Infer the Rest



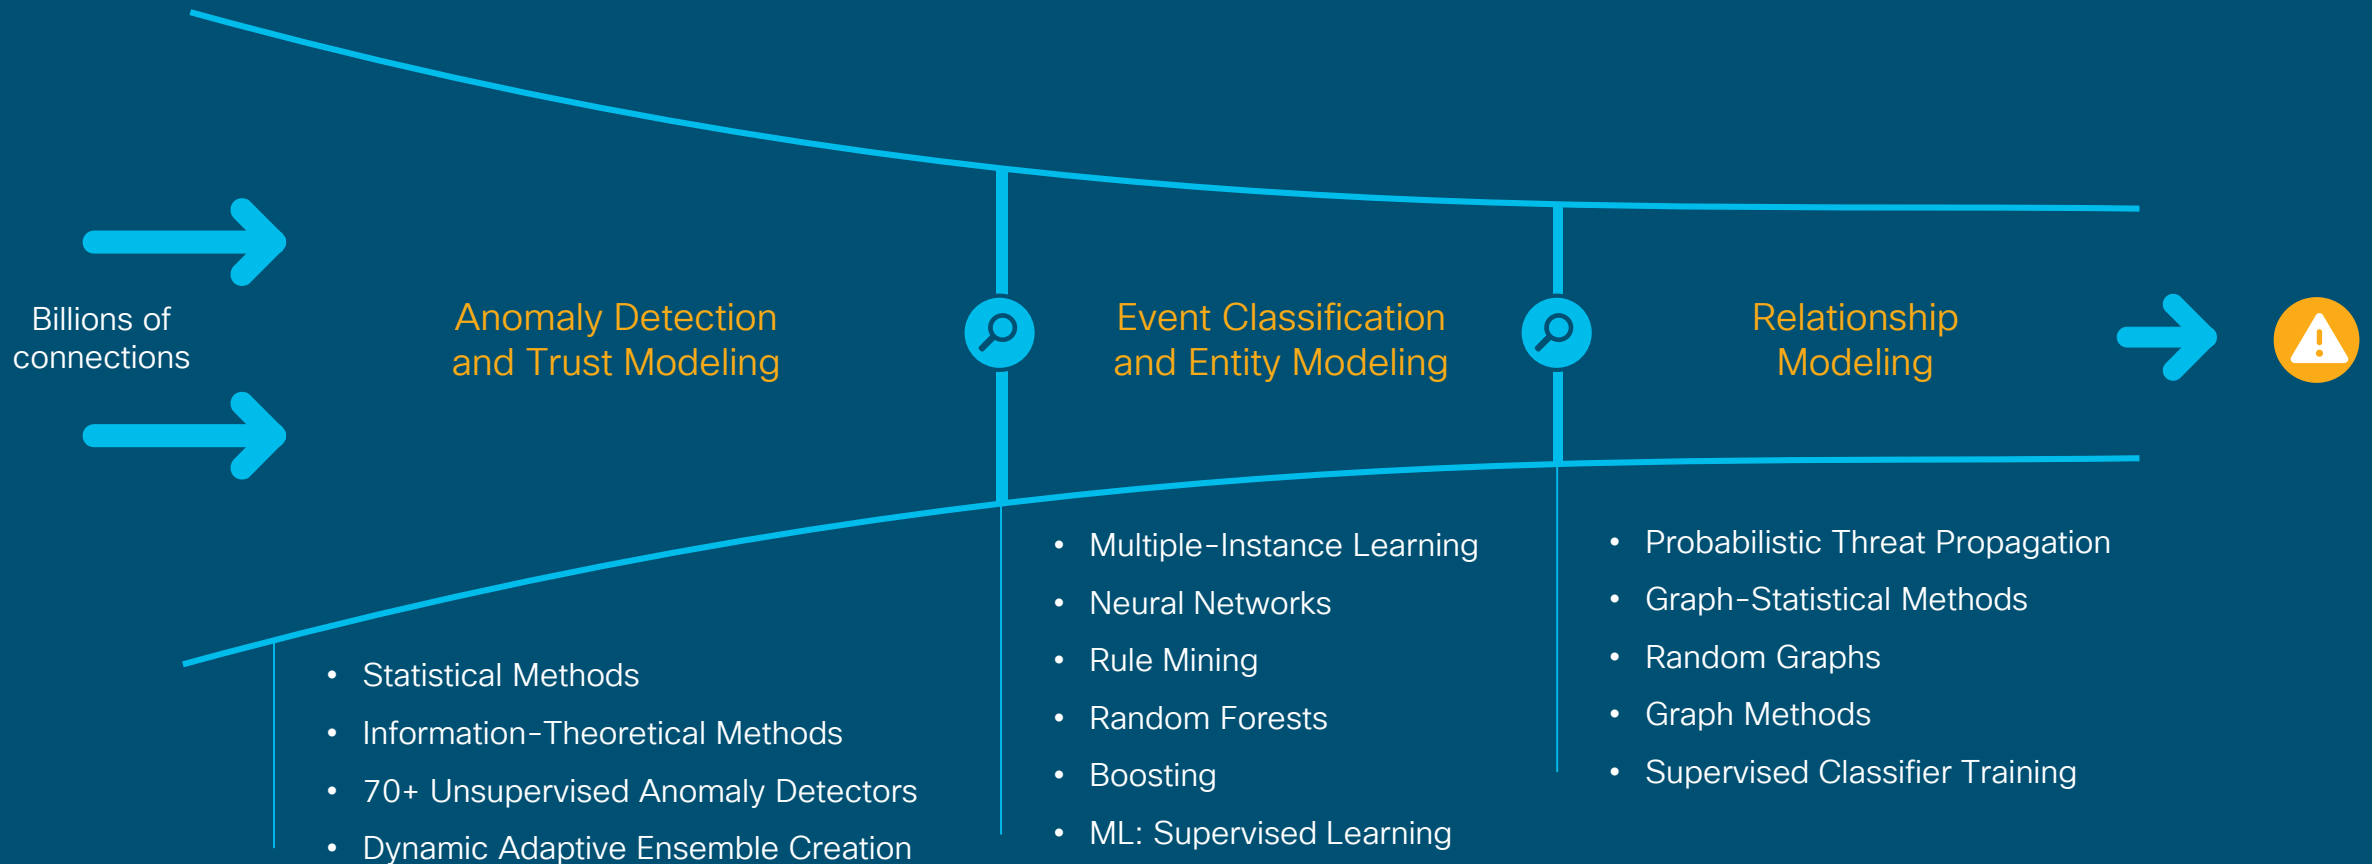
Threat Actor
Activity

Weird Stuff
(but not threat related)

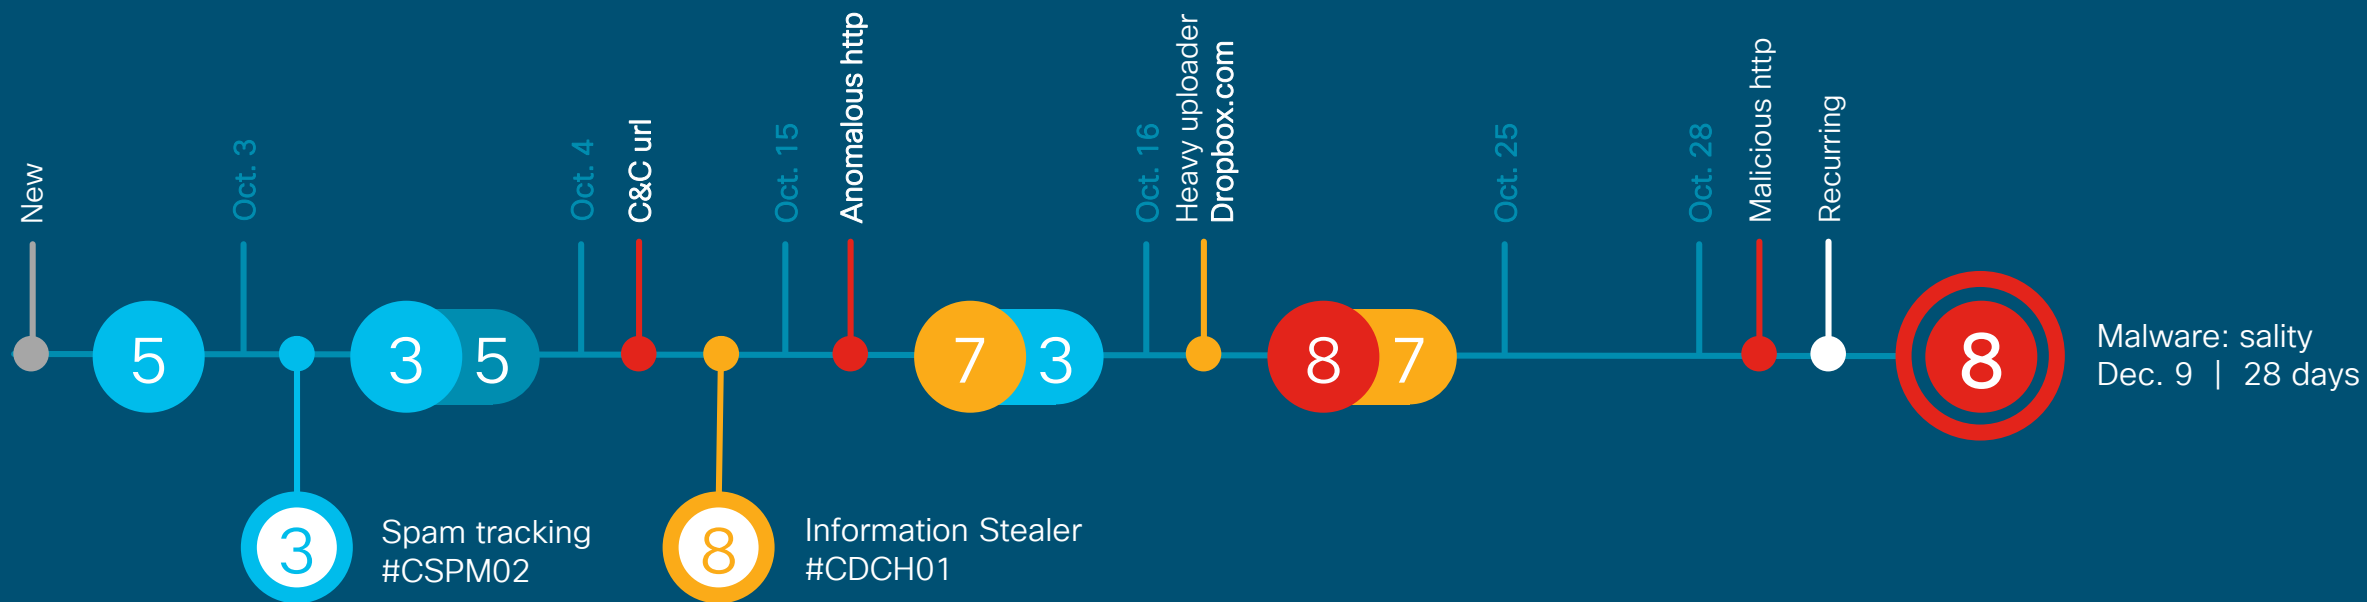
Normal Activity

Multi-layer Analytical Pipeline

Cascade of specialized layers of Machine Learning algorithms



Security that Shows its Work



Measure the Right Things

Efficacy of the Assertions

True/False Positive

True/False Negative

Overfitting/Undefitting

Root Mean Squared Error

Measure the Right Things

Feedback ✕

Was this alert helpful?

This provides feedback to us. It doesn't directly change our alerting criteria.

Conclusion

What to Ask Your Vendor



How are you applying Machine Learning in your product and why?

How do you measure its effectiveness?



Regarding supervised learning, what are you using for 'ground truth'?

What non-machine learning are you using and why?



What papers or open-source have you published regarding your analytics?

For the ML based assertions, what entailments are provided?

A Good Machine Learning Approach



Be Pragmatic



Entailments




Analytical pipeline, over single technique



Success is Domain Specific



Measure helpfulness, not mathematical accuracy



 NERD ALERT

Thank you!

