A **New** Application of **ML** for **Cybersecurity**: Attack Surface Management

Pamela Toman
Sr. Principal Machine Learning Engineer
Cortex Xpanse

March 2023

Cybersecurity is mission + AI

ASM is impossible without AI

AI for ASM means solving cool, hard problems

Nice to meet you

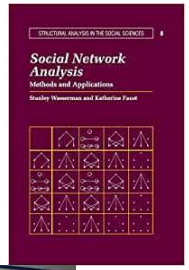

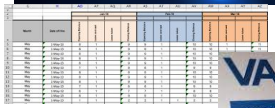

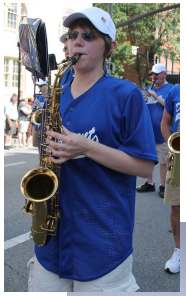

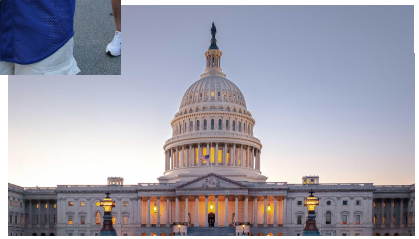

Stanford
ARTIFICIAL
INTELLIGENCE

QADIUM

EXPANSE

CORTEX XPANSE
BY PALO ALTO NETWORKS

Cybersecurity is mission + AI

ASM is impossible without AI

AI for ASM means solving cool, hard problems

Bad actors attack everyone all the time

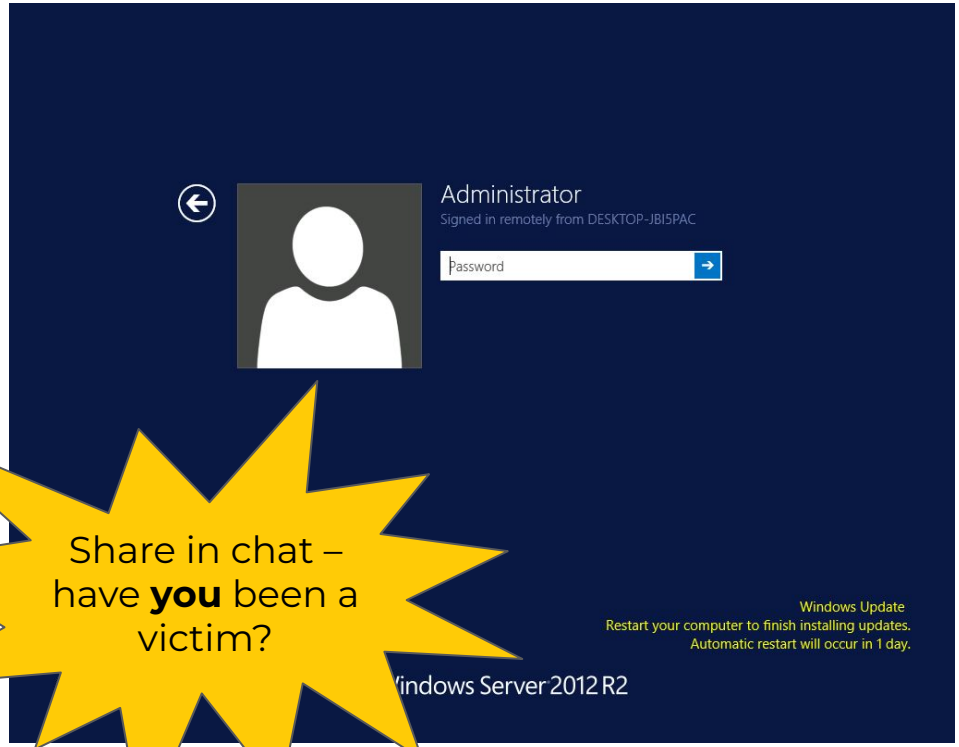

our screenshot of publicly accessible system
(most dangerous systems aren't screenshot-able)

#	Lowercase	Mixed Case	Mixed Case + Numbers
1	0 $	0 $	0 $
2	0 $	0 $	0 $
3	0 $	0 $	0 $
4	0 $	0 $	0 $
5	0 $	0 $	0 $
6	0 $	0 $	0 $
7	0 $	0 $	2 $
8	0 $	6 $	155 $
9	1 $	315 $	12,118 $
10	16 $	16,391 $	945,165 $
11	416 $	852,312 $	73.7 M$
12	10,820 $	44.3 M$	6 B$
13	281,330 $	2 B$	449 B$
14	7.3 M$	120 B$	34 T$
15	190.2 M$	6.2 T$	-
16	5 B$	324 T$	-
17	129 B$	-	-
18	3.3 T$	-	-

brute forcing passwords is easy
turingpoint.de, February 2020

Attacks produce money and power

- Sell data
- Hold data for ransom

- Rent the machine
- Mine cryptocurrency
- Drop malware installer

- Download non-public info
- Break infrastructure

Increased internet connectedness makes it worse

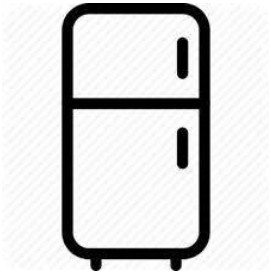

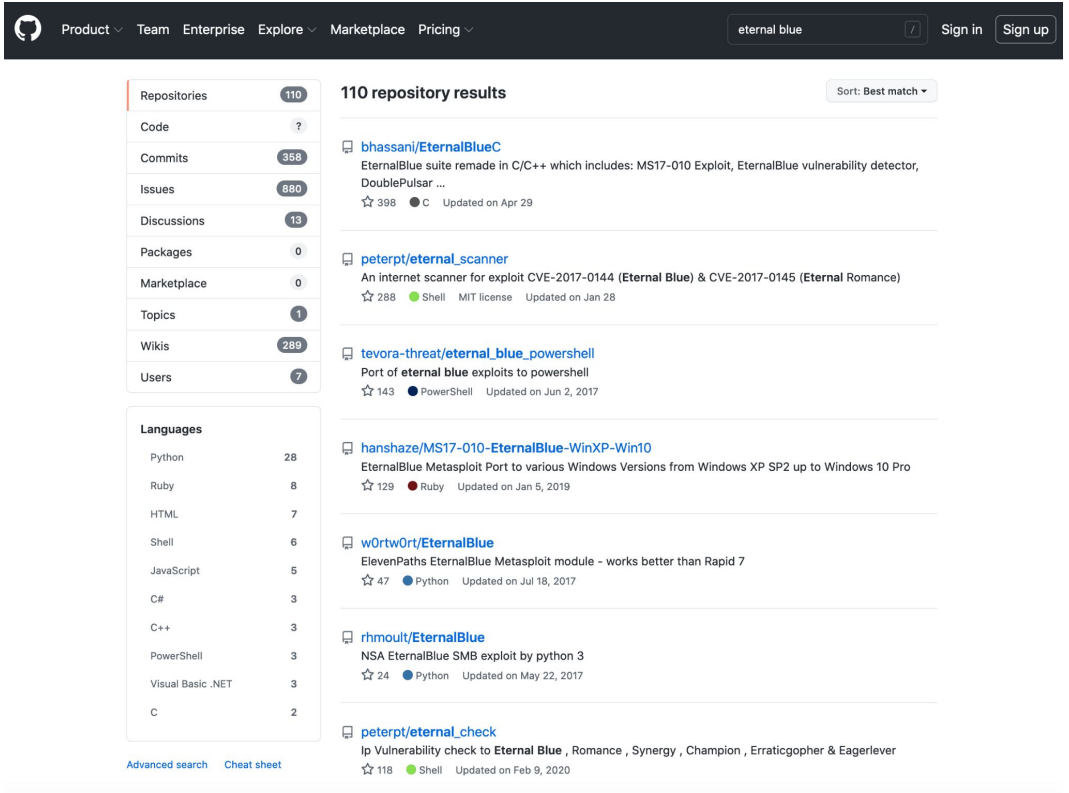

Without defense innocent people will fall

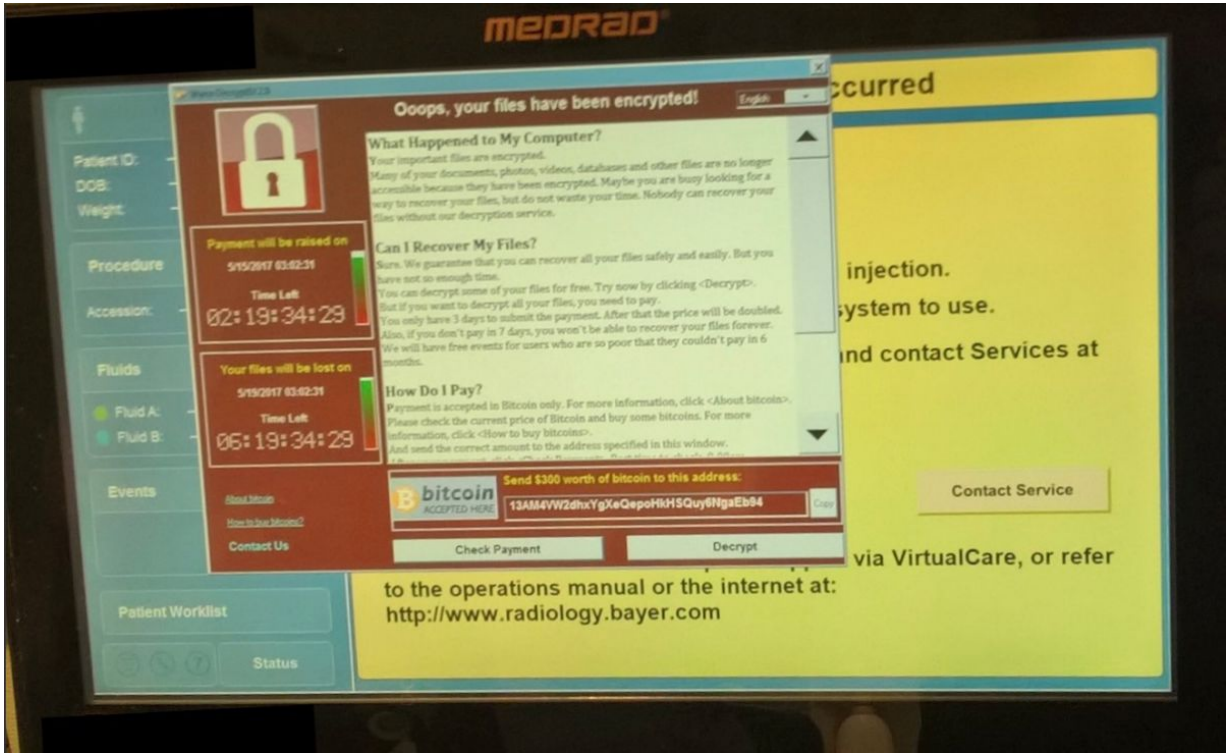

leaked

Without defense innocent people will fall

The settlement, if approved by a judge, would end a seven-year legal effort to win compensation for more than 21 million current and former federal employees who were victims of the hack of the Office of Personnel Management (OPM) in

information was intensely personal. It included background check forms that delved into victims' financial and romantic lives as well as Social Security numbers and — in a subset of about 5.6 million cases — fingerprint information.

OPM victims have faced a number of hurdles, including legal precedents that make it difficult or impossible to win compensation from data breaches that don't create direct economic loss. That's a high bar for OPM victims because the breach appears to have been for espionage purposes and there's no definitive evidence any of the stolen data

The Washington Post
Democracy Dies in Darkness

THE CYBERSECURITY 202

Lawyers are nearing a settlement deal for the infamous 2015 OPM hack

Analysis by Joseph Marks and Aaron Schaffer
May 9, 2022 at 7:28 a.m. EDT

Attorneys are closing in on a settlement deal that could deliver up to $63 million to some victims of one of the most cataclysmic data breaches in history.

The settlement, if approved by a judge, would end a seven-year legal effort to win compensation for more than 21 million current and former federal employees who were victims of the hack of the Office of Personnel Management (OPM) in 2015, which intelligence officials say was almost certainly perpetrated by the Chinese government.

The OPM breach marked a devastating blow to the U.S. government's reputation for cybersecurity and sparked intense anger among many victims — largely because the breached information was intensely personal. It included background check forms that delved into victims' financial and romantic lives as well as Social Security numbers and — in a subset of about 5.6 million cases — fingerprint information.

OPM victims have faced a number of hurdles, including legal precedents that make it difficult or impossible to win compensation from data breaches that don't create direct economic loss. That's a high bar for OPM victims because the breach appears to have been for espionage purposes and there's no definitive evidence any of the stolen data was ever used for cybercrime.

Our mission: Know what's there so you can defend it

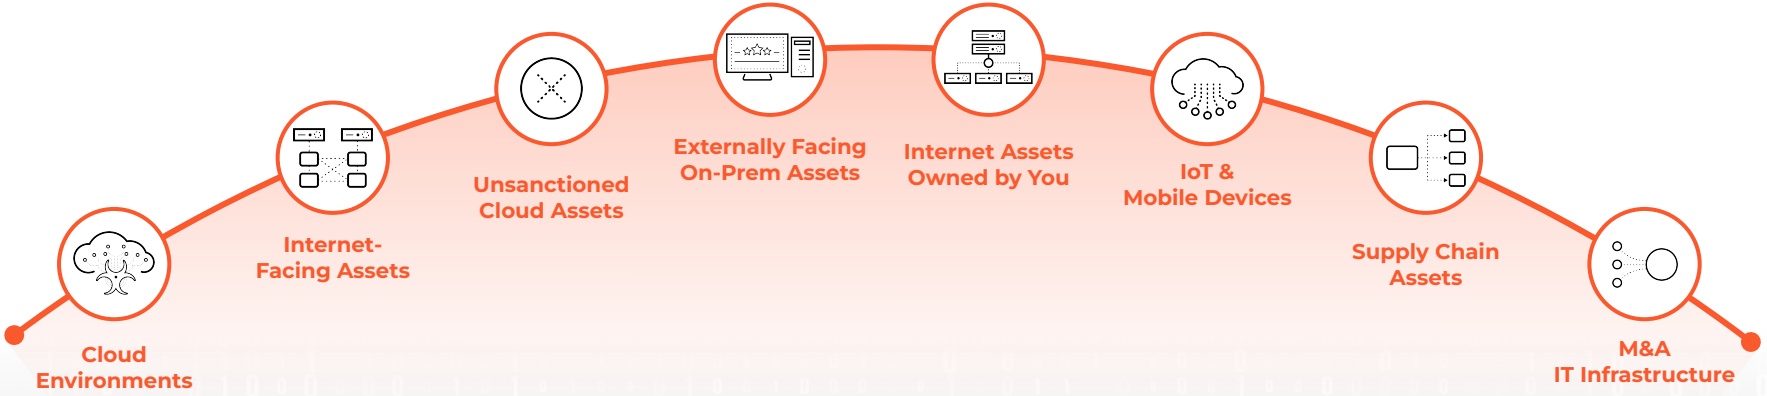

Attack Surface

 |

Cybersecurity is mission + AI

ASM is impossible without AI

AI for ASM means solving
cool, hard problems

ASM is impossible without AI

Observe

Attribute

Act

Observe: We need AI to explore IP space & find “what exists”

Brute force search? Not anymore….

IPv4 (32-bit)

`34.107.151.202`

4,294,967,296

IPv6 (128-bit)

`2607:f8b0:4005:807::200e`

340,282,366,920,938,463,463,374,607,431,768,211,456

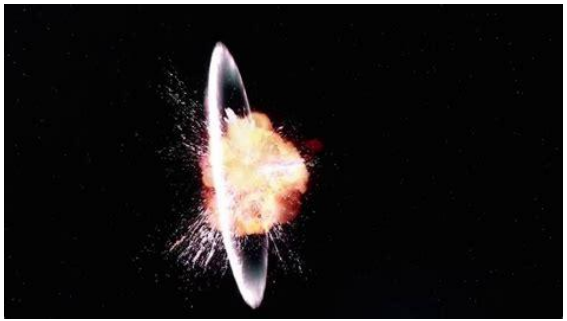

Alderaan explodes… us too…

Attribute: We need AI to identify the owner

IP address
34.107.151.202

domain
pppds.com

certificate
MIIGbDCCBVSg...

?

CVS Health

accenture

 |

Act: We need AI to make remediation possible

55% of enterprises
see more than **10k alerts** a day

79% of security teams
feel **overwhelmed** by the volume of threat alerts

85% of security professionals
think their security team is **understaffed**

sources: SC Media via survey at 2018 RSA; analyst report by Enterprise Management Associates; SOC report by Ponemon Institute

ASM is impossible without AI

Observe

Attribute

Act

Cybersecurity is mission + AI

ASM is impossible without AI

AI for ASM means solving cool, hard problems

 |

Let’s sample cool, hard problems in 4 roles

Who owns what on the internet?

ML
Framing the question

Engineering
Scaling the answer

MLOps
Supporting the lifecycle

Ethics
Monitoring the effects

ML: How should we frame the attribution question?

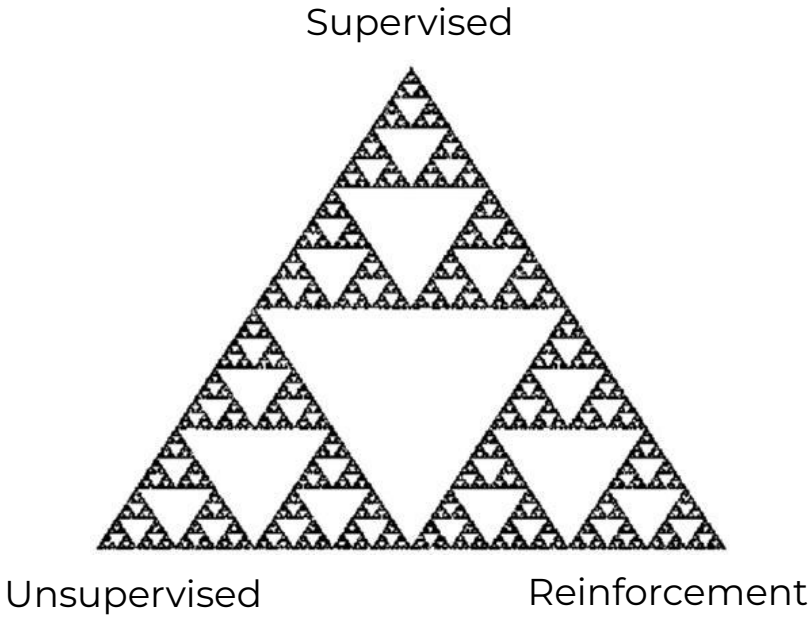

 |

Engineering: How do you attribute ownership at internet scale?

Naively doing inference is too expensive

Direct & indirect field values

Every domain/IP/etc.

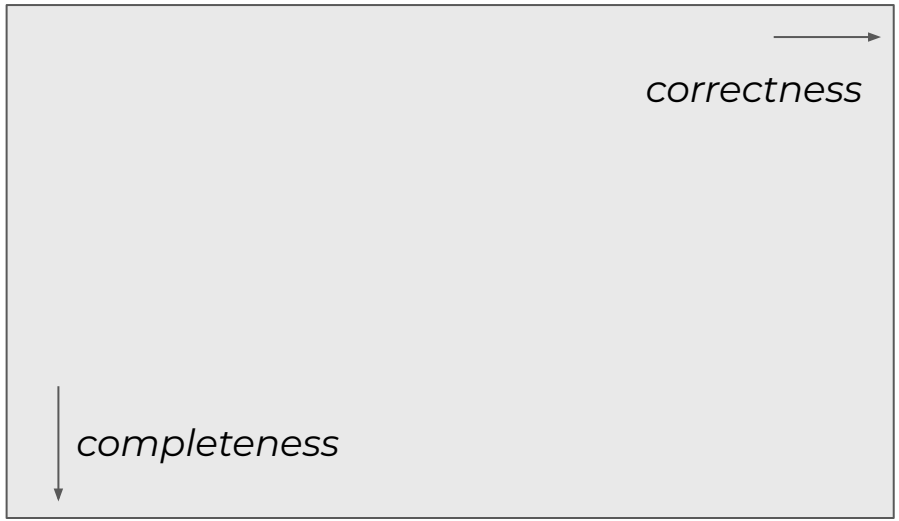

(* every organization)

MLOps: How do you support the ML development lifecycle?

In this paper, we argue that ML systems have a special capacity for incurring technical debt, because they have all of the maintenance problems of traditional code plus an additional set of ML-specific issues. This debt may be difficult to detect because it exists at the *system* level rather than the code level. Traditional abstractions and boundaries may be subtly corrupted or invalidated by the fact that data influences ML system behavior. Typical methods for paying down code level technical debt are not sufficient to address ML-specific technical debt at the system level.

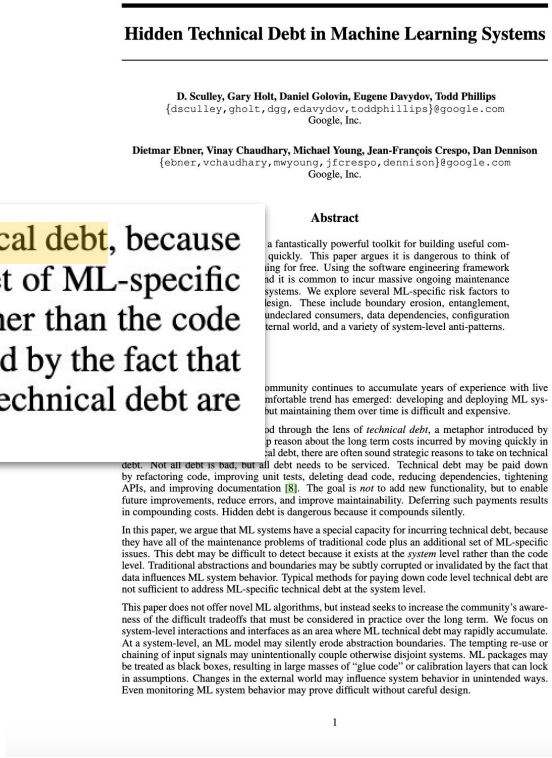

Hidden Technical Debt in Machine Learning Systems

D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips
{dsculley,gholt,dgg,edavydov,toddphillips}@google.com
Google, Inc.

Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-François Crespo, Dan Dennison
{ebner,vchaudhary,mwyoung,jfcrespo,dennison}@google.com
Google, Inc.

Abstract

a fantastically powerful toolkit for building useful com-
quickly. This paper argues it is dangerous to think of
ing for free. Using the software engineering framework
d it is common to incur massive ongoing maintenance
systems. We explore several ML-specific risk factors to
esign. These include boundary erosion, entanglement,
undeclared consumers, data dependencies, configuration
ternal world, and a variety of system-level anti-patterns.

ommunity continues to accumulate years of experience with live
mfortable trend has emerged: developing and deploying ML sys-
but maintaining them over time is difficult and expensive.

od through the lens of *technical debt*, a metaphor introduced by
p reason about the long term costs incurred by moving quickly in
al debt, there are often sound strategic reasons to take on technical
debt. Not all debt is bad, but all debt needs to be serviced. Technical debt may be paid down by refactoring code, improving unit tests, deleting dead code, reducing dependencies, tightening APIs, and improving documentation [8]. The goal is *not* to add new functionality, but to enable future improvements, reduce errors, and improve maintainability. Deferring such payments results in compounding costs. Hidden debt is dangerous because it compounds silently.

In this paper, we argue that ML systems have a special capacity for incurring technical debt, because they have all of the maintenance problems of traditional code plus an additional set of ML-specific issues. This debt may be difficult to detect because it exists at the *system* level rather than the code level. Traditional abstractions and boundaries may be subtly corrupted or invalidated by the fact that data influences ML system behavior. Typical methods for paying down code level technical debt are not sufficient to address ML-specific technical debt at the system level.

This paper does not offer novel ML algorithms, but instead seeks to increase the community's awareness of the difficult tradeoffs that must be considered in practice over the long term. We focus on system-level interactions and interfaces as an area where ML technical debt may rapidly accumulate. At a system-level, an ML model may silently erode abstraction boundaries. The tempting re-use or chaining of input signals may unintentionally couple otherwise disjoint systems. ML packages may be treated as black boxes, resulting in large masses of "glue code" or calibration layers that can lock in assumptions. Changes in the external world may influence system behavior in unintended ways. Even monitoring ML system behavior may prove difficult without careful design.

1

Ethics: How do you monitor for adverse effects?

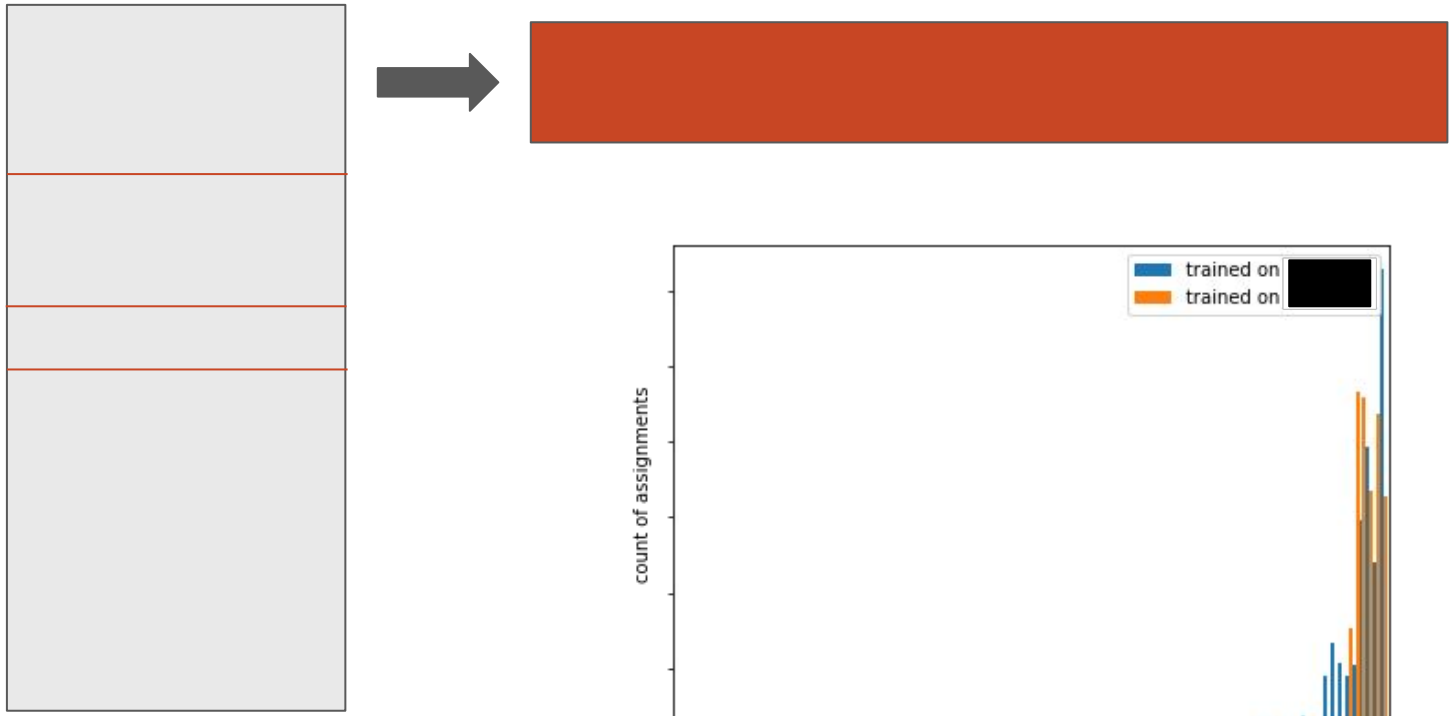

"Data Slices" paper in NeurIps '21
Data-Centric AI workshop

We have cool, hard problems coming out our ears!

Who owns what on the internet?

ML
Framing the question

Engineering
Scaling the answer

MLOps
Supporting the lifecycle

Ethics
Monitoring the effects

Drilling out on Palo Alto Networks

Our leadership drives innovation

“We’re driving an AI-based SOC transformation.”

– CEO Nikesh Arora
in the 23 February 2023 earnings call

We have the data, scope & customers to build big

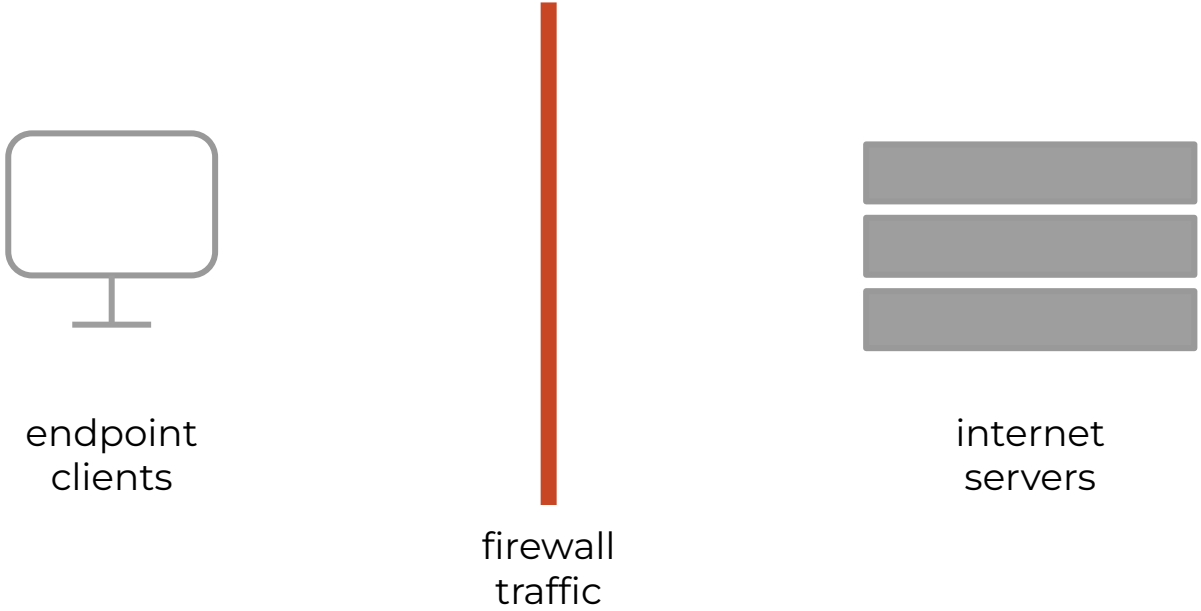

Our whole organization is solving cool, hard problems

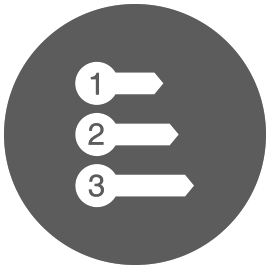

Prioritization

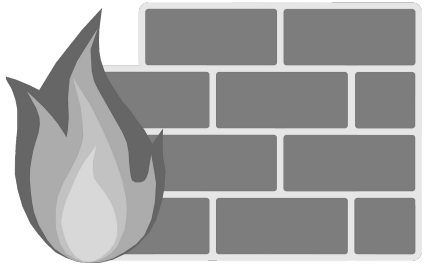

Zero-day blocking

Correct configuration

Data exfiltration

Malware identification

Thank you

Happy to answer questions!

Pamela Toman
Sr. Principal Machine Learning Engineer
Cortex Xpanse

ptoman@paloaltonetworks.com

paloaltonetworks.com

 |