Hewlett Packard
Enterprise

# Enterprise Data Exfiltration Detection and Prevention

Pratyusa K. Manadhata
Hewlett Packard Labs

# Unauthorized transfer of sensitive information from a victim enterprise to an attacker

https://www.google.com/imghp

# Examples..

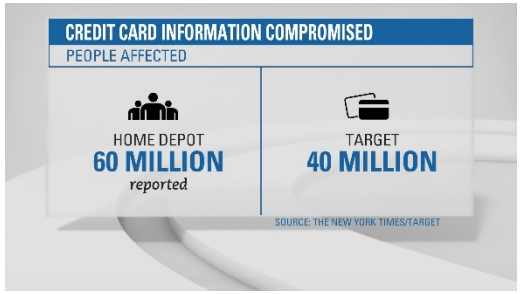










http://www.privacyrights.org/data-breach

# Cyber Espionage in 2015

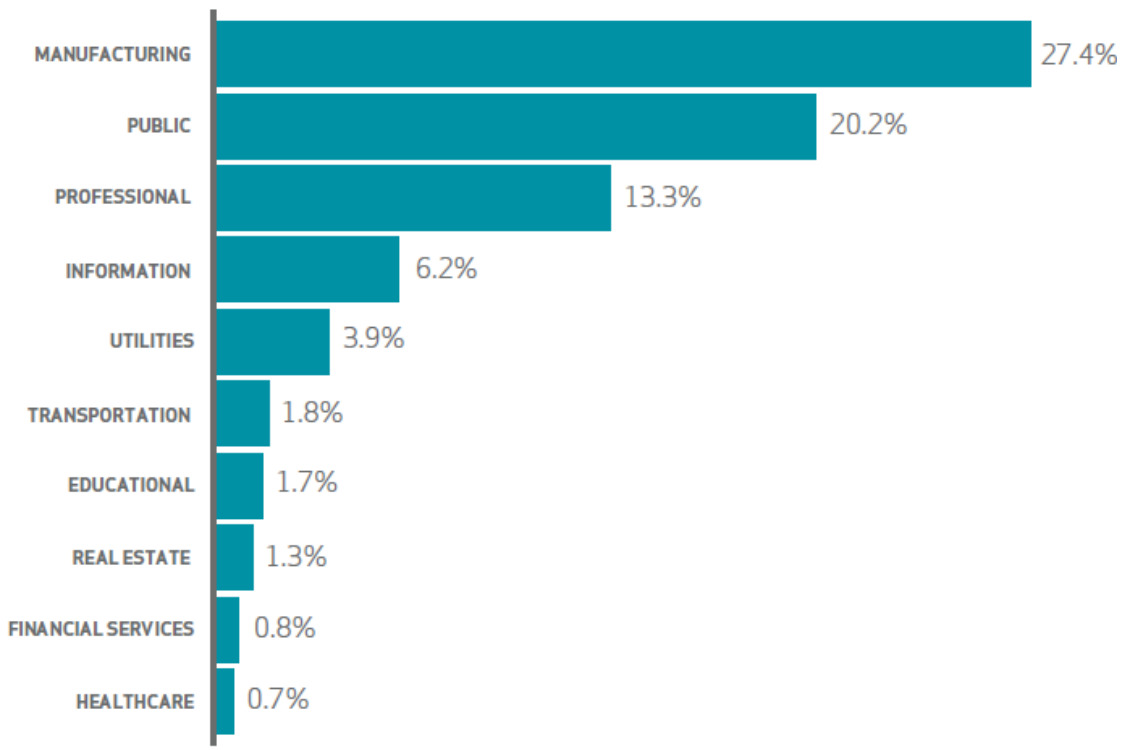


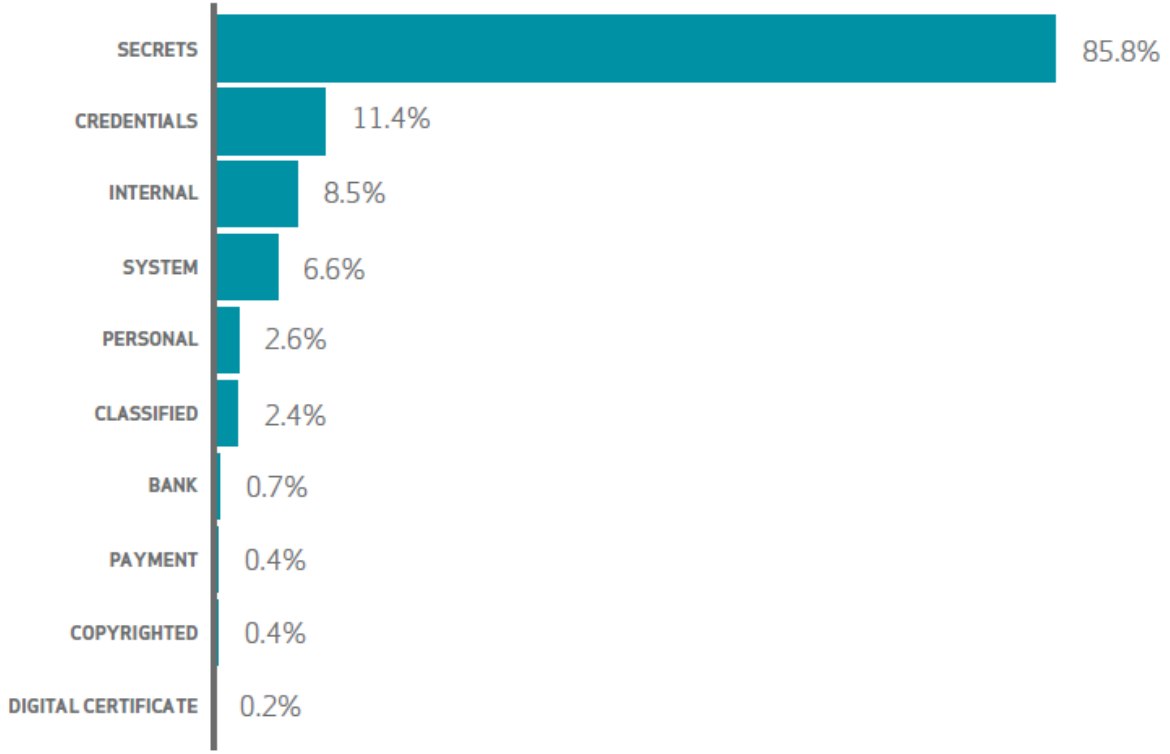


548 incidents reported in 2015.

2015 Data Breach Investigations Report, Verizon.

# Data Loss Prevention (DLP) Products (2006~)


Image: http://blog.skodaminotti.com/blog/data-loss-prevention-part-2-choosing-a-dlp-solution/

# Sensitive data identification

– Regular expression
  – social security numbers, telephone numbers, addresses, and other data that has a significant amount of structure.

– Keywords
  – small number of known keywords can identify private data, e.g., medical or financial records

– Fingerprints
  – Hashes of substrings of unstructured data

# Limitations

– Good at identifying “universally” confidential data, e.g., credit card number and SSN

  – Organization specific key word and fingerprint generation was challenging

– Prevents accidents and commodity attacks

  – Easy to circumvent

Can we learn organization specific sensitive data?

# Text classification for DLP

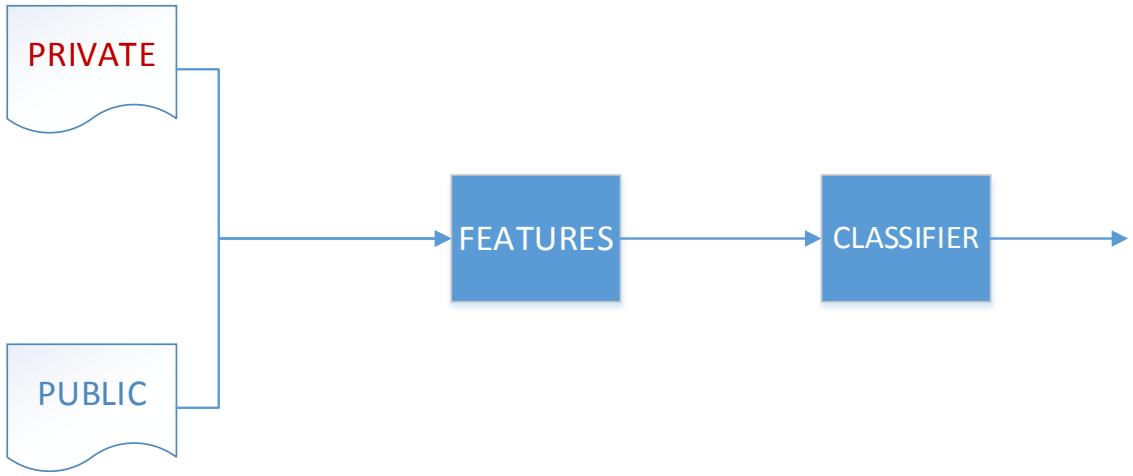


Text Classification for Data Loss Prevention, Michael Hart, Pratyusa Manadhata, and Rob Johnson
Privacy Enhancing Technology Symposium (PETS) 2011

# Real world is messy

– Enterprise Private (EPR), Enterprise Public (EPL), and Non Enterprise (NE)

– EPR and EPL likely to be relatively similar

– Many NE share almost no features with EPR and EPL

– Some NE may be quite similar to EPL

Contradicting goals: catch subtle differences between EPL and EPR,

but not overfit to be able identify NE as non-sensitive

# Our approach [PETS 2011]

– Supplement training data with NE

– Adjust SVM boundary toward EPL

– Two-step classifier to reduce FPs

| Dataset | Baseline FDR | Our classifier FDR |
|---|---|---|
| DynCorp | 4.49% | 0.00% |
| Enron | 47.05% | 0.92% |
| Google | 8.99% | 1.06% |
| Mormon | 0.88% | 0.36% |
| TM | 22.06% | 0.01% |

**Table 3.** The False Discovery Rate of the baseline approach far exceeds our classifier, implying that the baseline approach would fare poorly in real world networks whereas ours would not raise much fewer alarms.

# “Vector Machine Learning”

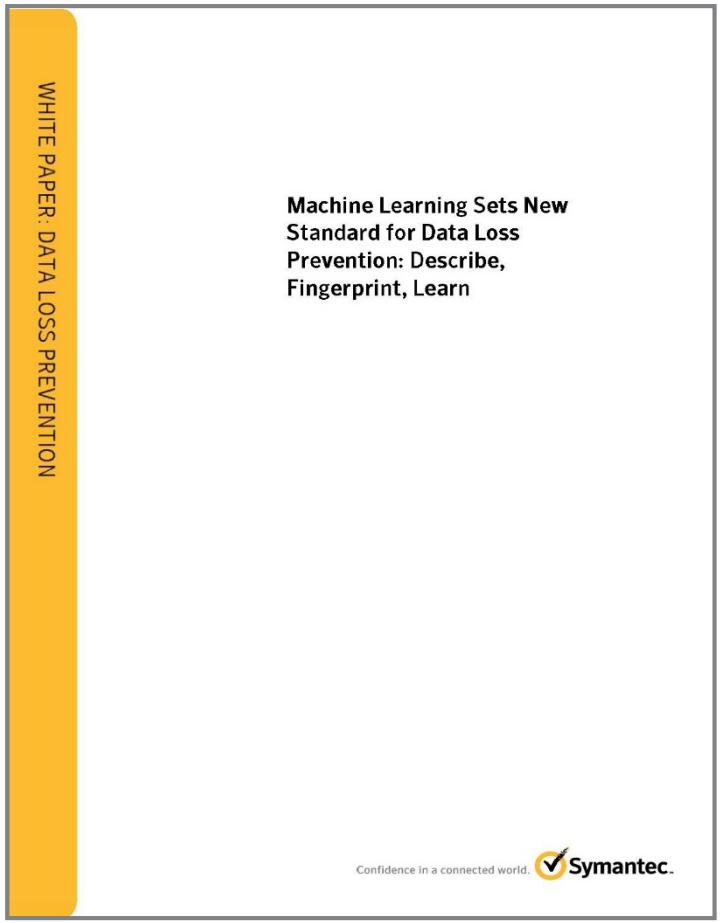


http://eval.symantec.com/mktginfo/enterprise/white_papers/b-dlp_machine_learning.WP_en-us.pdf

The threat landscape has evolved.

# Advanced threats: The new landscape

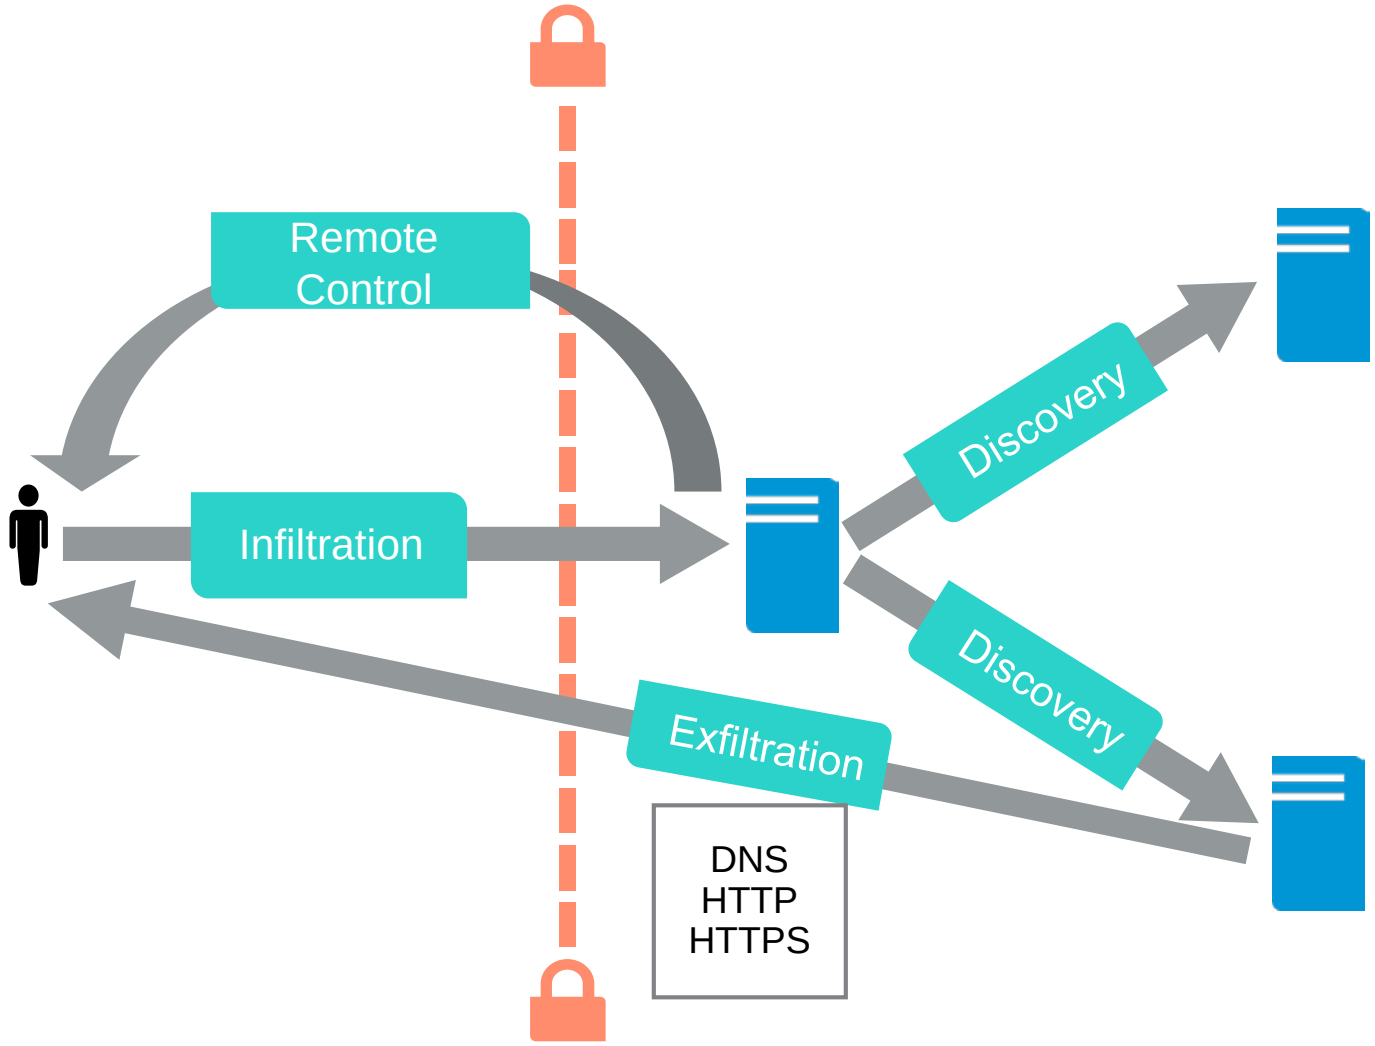

# Example: DNS Exfiltration

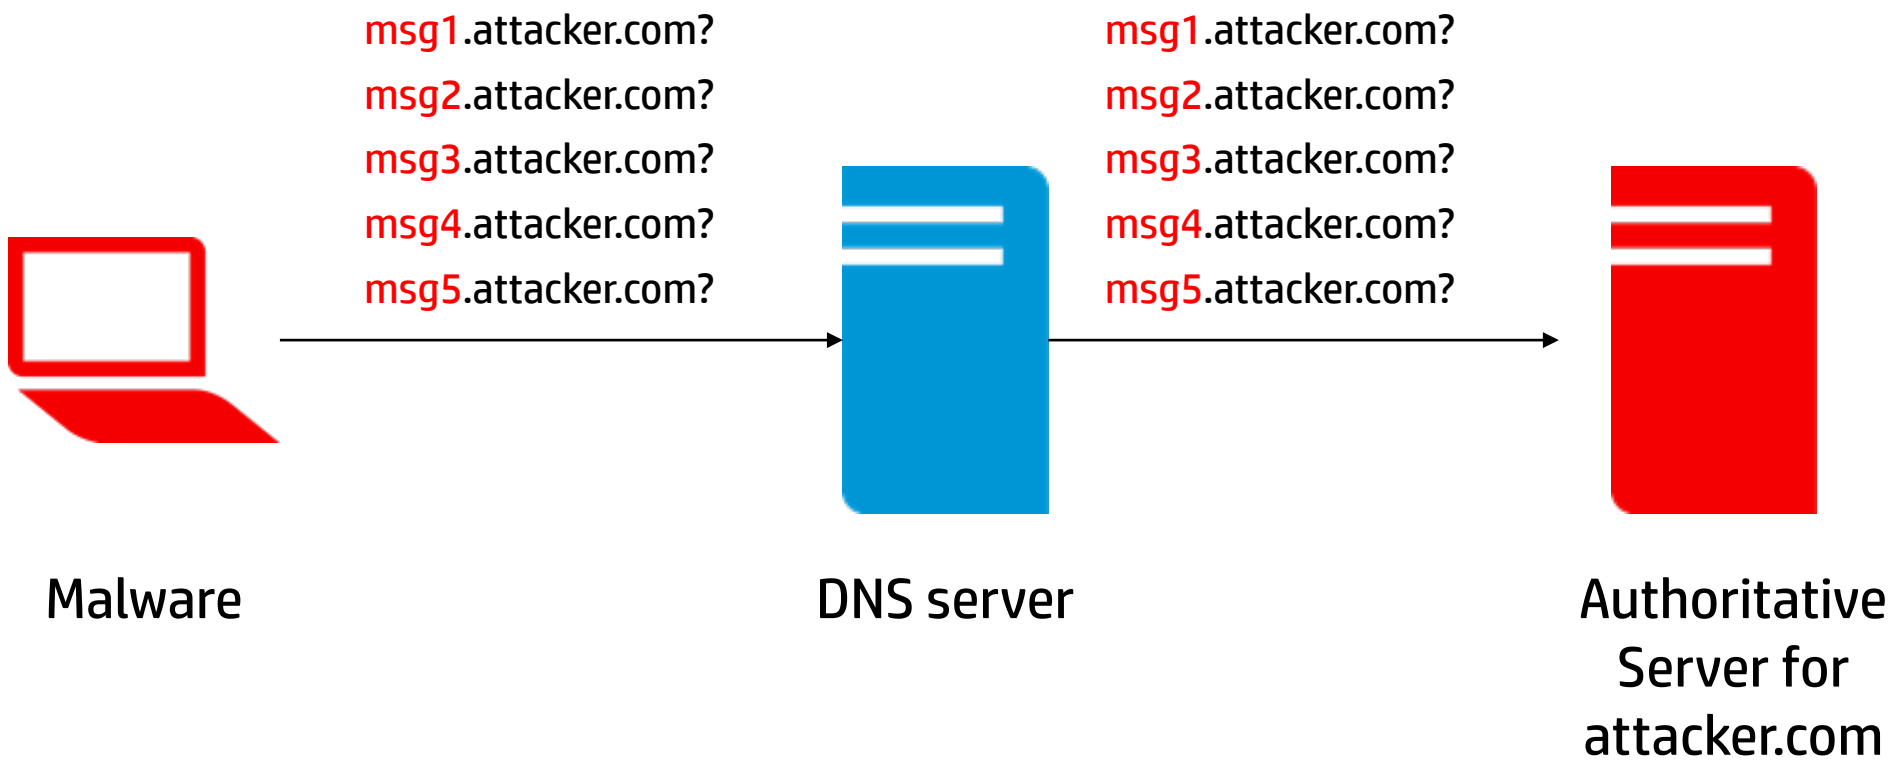

# A real world example

– Queries

BLGCOFDAGOOOESDULBOOBOOOOOOOOOOOOOOOOOLDOSESKGKHHF.detacsufbo.ru

EUJSFLDAGOOOESDUDBOOBOOOOOOOOOOOOOOOOOOSSJHGHFCLFOHCHLGHSAHAHU.CHLAAFHLSGHAFGFUOOEUGDKLCSHEKLJBOCOSECHFFUGBS
KGDJGGGHOJHJCGJG.KCDOELDUOEGUCUOUHJUAKEGGGFGEKHLGFDFESJOEL.detacsufbo.ru

SHUDHFDAGOOOESDUGBOOBOOOOOOOOOOOOOOOOOEDKDFBBHLEGGJLGUFABHCCU.DHDFFCHHKSHGHAOUBGEGEJLGFHUBDFGUGJDFFEAKFSBFFG
SDACGHCSKBHLSCGHH.EHSHHJFHUAAOOGKKSDDAHAUBBJDCCKGSHKLGJGAS.detacsufbo.ru

OHDOBHDAGOOESDUGBOOHOOOAOOOOOOOOOOOOOOO.detacsufbo.ru

HBSGGCDAGOOESDUUSOOBOOOOOOOOOOOOOOOOOO.detacsufbo.ru

– Responses (TXT records)

LLCDGHDABOOOSSUHOOOFOOOOOOOOOOOOOOOOOO

KJGDUDABOOOSBSUHOOOFOOOOOOOOOOOOOOOOOO

JJDHUDABOOOSBSUHOOOFOOOOOOOOOOOOOOOOOO

HBEAGDABOOOSBSUHOOOUOOOOOOOOOOOOOOOOOO

KALFCSDAOOOSBSUHOOOFOOOOOOOOOOOOOOOOOO

How do we deal with the new threat landscape?

# Detection mechanisms ??

– Many queries from a client to a domain -> misconfigured devices

– Many distinct queries from a client to a domain -> signaling needs few subdomains

– Many long queries from a client to a domain -> credit card numbers are not long

# Detection mechanism: Information content

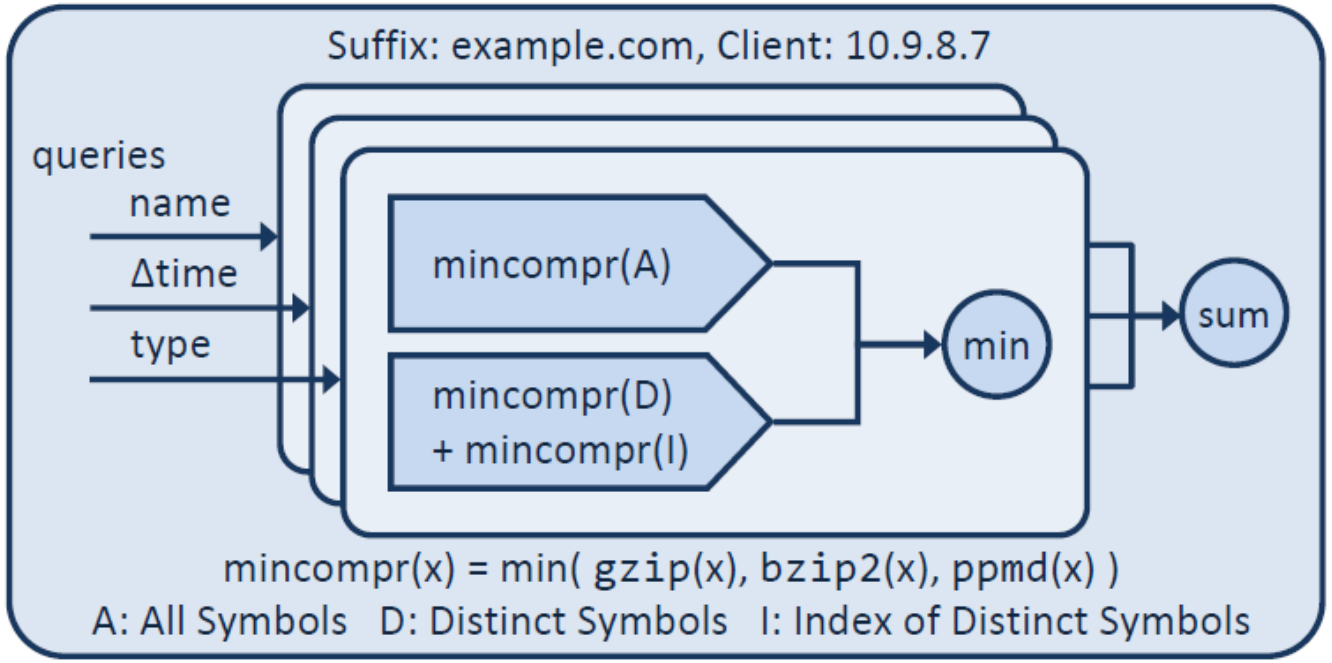


Vern Paxson, Mihai Christodorescu, Mobin Javed, Josyula Rao, Reiner Sailer, Douglas Schales, Marc Ph. Stoecklin, Kurt Thomas, Wietse Venema, and Nicholas Weaver.
Practical comprehensive bounds on surreptitious communication over DNS.
In *Proceedings of the 22nd USENIX conference on Security* (SEC'13). USENIX Association, Berkeley, CA, USA, 17-32.

# Industry heuristics for DNS exfiltration detection

– Lengths of DNS queries and responses

– Sizes of request and reply packets

– Entropy

– Total number/volume of DNS queries from a device

– Total number/volume of DNS queries to a domain

– …

# Characteristics of heuristics driven solutions

– Complex to build/maintain, but very fragile
  – False negatives and false positives

– Independent of other security products in the network

– Need a human being in the loop for response/remediation

Can we build a reliable detector on top of ‘noisy’ detectors?

# Problem setting

– Defender
  – observes a stream of noisy alerts (or absence of alerts)
  – has partial knowledge of the network

– Must identify domains involved in exfiltration and decide whether to block traffic (plan of action)

– While weighting the cost of
  – Deploying noisy detectors
  – Data loss due to false negatives
  – Disruption due to false positives

Decision making and planning under uncertainty

# Virtually Distributed Partially Observable Markov Decision Processes (VD-POMDP) [GameSec 16]

Network of enterprise devices, web domains,
and noisy detector nodes

Original POMDP: Impractical to solve for 3-4 nodes

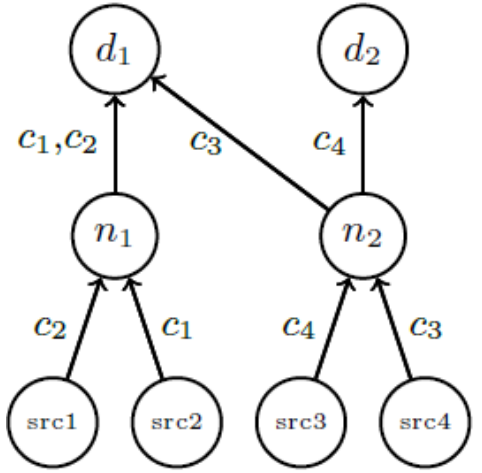


VD-POMDP:

1. Abstract action and observation space
2. Factor the original POMDP into one sub-POMDP per domain, solve them “offline”
3. Online policy aggregation using MILP to get final joint action

Sara Mc Carthy, Arunesh Sinha, Milind Tambe and Pratyusa K. Manadhata, Data Exfiltration Detection and Prevention: Virtually Distributed POMDPs for Practically Safer Networks, 7th Conference on Decision and Game Theory for Security (GameSec), New York, NY, Nov 2016.

# Experimental setting

– DETER testbed at USC

– Iodine to simulate DNS exfiltration

– Normal DNS query behavior simulation via scripts

– Paxson et al.’s approach as the noisy detector

– Synthetic workload for parameter sensitivity testing

# Results

– Runtime
  – POMDP doesn’t scale beyond 3 domains, whereas VD-POMDP scales linearly
  – For small networks (~1000 nodes), offline phase takes hours and online phase takes seconds

– Performance
  – VD-POMDP’s accuracy and time to detection is similar to POMDP

– Robustness
  – Very high detection rate (>0.95%) even with extremely noise detectors

# Summary

– Data exfiltration detection and prevention is an ongoing arms race

– Existing approaches don’t work in the new threat landscape

– Point detection approaches will always be noisy

– We need to build robust detectors on top of noisy detectors

– POMDP/Game theory may help us build robust detectors

Hewlett Packard
Enterprise

# Thank you

Pratyusa K. Manadhata
manadhata@hpe.com