

Detecting and preventing domain name abuse in .eu

Lieven Desmet, KU Leuven – Marc Van Wesemael, EURid



Registration of DNs with fraudulent/criminal intentions

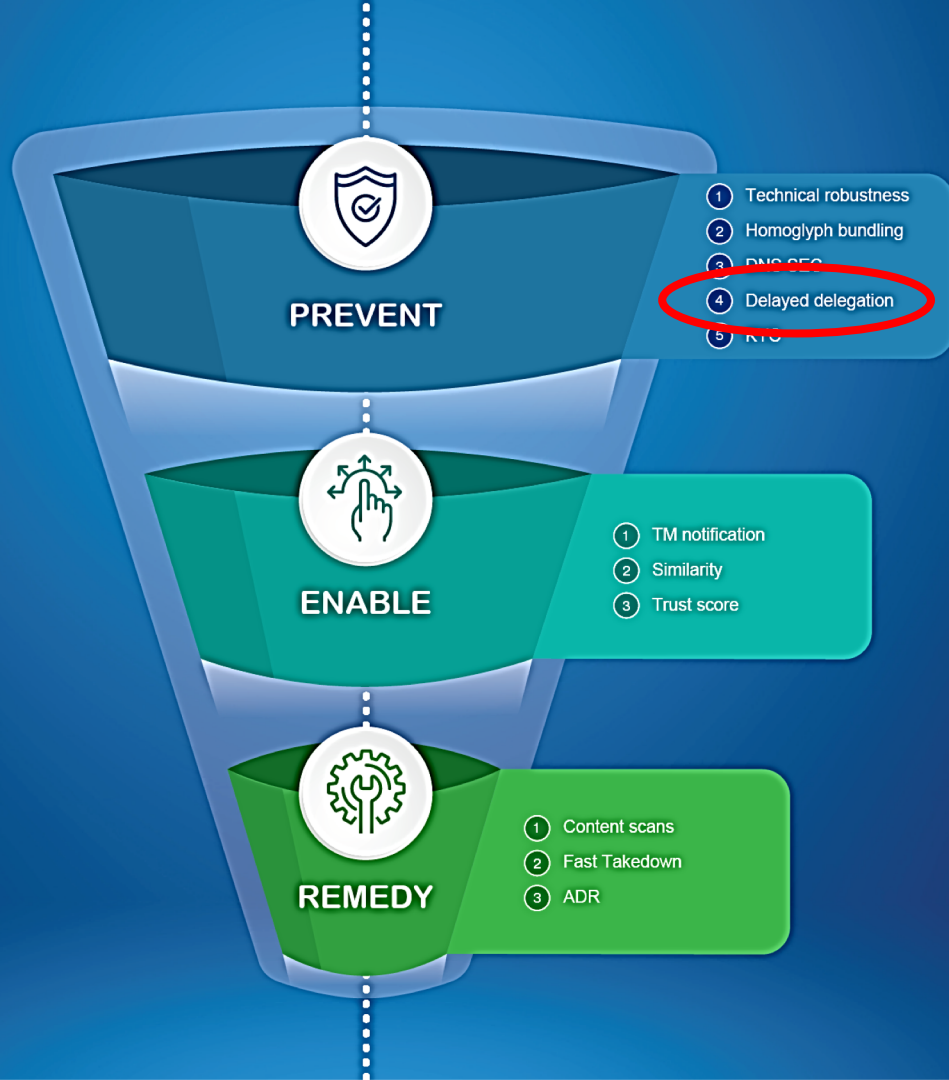
- › Types of abuse
 - › To attract traffic to websites (use of reputation of somebody else)
 - › To distribute malware
 - › To send SPAM
 - › To sell illegal products (drugs, counterfeited goods, fake medication, etc.)
 - › To sell products and not deliver
 - › Often (very) short term use (hours or days)

- › Key issue: fake identity of the registrant

The .eu trust strategy

› Delayed delegation

- ›› Predict at time of registration whether a domain name will be used abusively






Insights in malicious domain registrations

T. Vissers et al., Exploring the ecosystem of malicious domain registrations in the .eu TLD, Research in Attacks, Intrusions, and Defenses (RAID 2017), September 2017.



Activity of identified campaigns

Registrations per day ● 100 ● 200 ● 300 ● 400



Insight 1: Varying campaign characteristics



- › Simple campaign (c_14)
- › Single (fake) registrant used throughout the campaign

- **41 days active**
- **989 blacklisted registrations**
(= 95.37%)

Example campaign (c_11)

› Multiple fake registrant details

›› Combinations of

- 2 email accounts,
- 3 phone numbers,
- 4 street addresses

- **8 months active**
- **1,275 blacklisted registrations**
(= 53.96%)

Example of an advanced campaign (c_15)


- › Registrant details:
 - › 98 fake registrants
 - › Generated by Laravel Faker tool
- › Domain names:
 - › Consist out of 2-3 Dutch words
 - › Dutch words are reused across registrants
- › Batches of 8, 16, 24 or 32 registrations

- **8+ months active**
- **514 blacklisted registrations**
(= 26.95%)

Insight 2: Small set of malicious actors



Registrations per day ● 100 ● 200 ● 300 ● 400



At most 20 actors represent 80% of malicious registrations

Insight 3: Top facilitators for malicious registrations



	Nb of malicious	Contribution Malicious	Benign	Toxicity
1. registrar_5	10,353	49.61%	2.27%	36.25%
2. registrar_3	3,004	14.39%	2.64%	12.41%
3. registrar_7	2,327	11.15%	0.46%	38.67%
1. gmail.com	4,221	20.23%	24.79%	2.08%
2. yahoo.com	3,348	16.04%	1.49%	21.85%
3. aol.com	2,134	10.23%	0.31%	46.28%

Insight 3: Top facilitators for malicious registrations



	Nb of malicious	Contribution Malicious	Benign	Toxicity
1. registrar_5	10,353	49.61%	2.27%	36.25%
2. registrar_3	3,004	14.39%	2.64%	12.41%
3. registrar_7	2,327	11.15%	0.46%	38.67%
1. gmail.com	4,221	20.23%	24.79%	2.08%
2. yahoo.com	3,348	16.04%	1.49%	21.85%
3. aol.com	2,134	10.23%	0.31%	46.28%

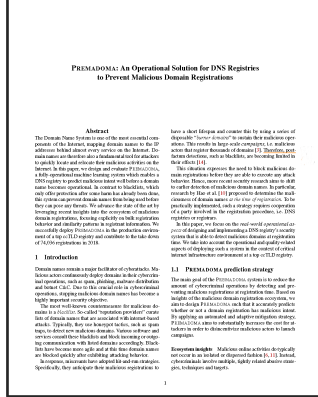
Quick overview of other insights



- › Majority of blacklisted domains is spam-based
- › Blacklisting happens shortly after registration
- › Malicious actors exhibit “human behavior”
 - ›› Some work 9-to-5, take holidays, ...
 - ›› They sometimes make typos
 - ›› They vary over time

Registration-time prediction of malicious intent

*J. Spooren et al., **PREMADOMA: An Operational Solution for DNS Registries to Prevent Malicious Domain Registrations**, Annual Computer Security Applications Conference (ACSAC 2019), December 2019.*

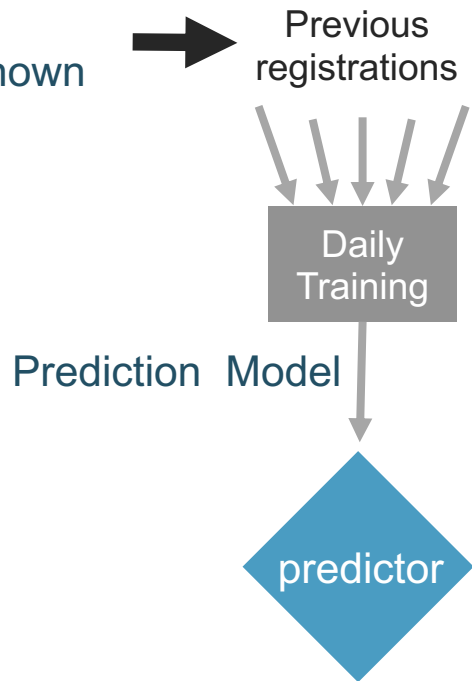


Pro-active detection and prevention

Previous registrations for which
the results (abuse/no abuse) is known

Pro-active detection and prevention

Previous registrations for which the results (abuse/no abuse) is known



Pro-active detection and prevention

Previous registrations for which the results (abuse/no abuse) is known



Previous registrations



Daily Training

Prediction Model



For each new registration, the system predicts if the domain will be used for malicious activity

New registration



Pro-active detection and prevention

Previous registrations for which the results (abuse/no abuse) is known

Previous registrations




Prediction Model

- Domains with malicious intent can be
- Detected early
 - Delayed
 - Prevented from being registered

For each new registration, the system predicts if the domain will be used for malicious activity

New registration



Underlying assumptions/rationales for our predictors

- › Similarity-based agglomerative clustering
 - ›› Domains belonging to the same campaign have very similar registration details

Underlying assumptions/rationales for our predictors


- › Similarity-based agglomerative clustering
 - ›› Domains belonging to the same campaign have very similar registration details
- › Reputation-based classification
 - ›› Domains using registration facilitators with a bad reputation (e.g. email providers or registrars), are likely to be malicious as well

Evaluation on historical data

- › Ground truth-based evaluation
 - › Recall: 66.23%
 - › Precision: 84.57
 - › False positive rate: 0.30%

- › Campaign-based evaluation
 - › 17 out of the 20 campaigns are well predicted

Detecting and preventing abuse in .eu: “1 picture ...”



Operational results

- › Period: July 2017 – December 2018 (18 months)
 - › Recall: 85.51%
 - › Precision: 72.04%
 - › False positive rate: 2.86%
- › Very big campaigns (October 2017 - March 2018)
- › Incomplete ground truth

Abstract. Domain blacklisting is a widely used technique to prevent access to malicious websites. This paper evaluates the effectiveness of domain blacklisting against malicious DNS registrations. We use a large dataset of malicious DNS registrations and compare them to a dataset of domain blacklists. We analyze the overlap between the two datasets and the impact of domain blacklisting on the visibility of malicious DNS registrations. Our results show that domain blacklisting is effective in reducing the visibility of malicious DNS registrations, but that it is not perfect. We discuss the reasons for this and provide recommendations for improving domain blacklisting.

1. INTRODUCTION

Domain blacklisting is a widely used technique to prevent access to malicious websites. This paper evaluates the effectiveness of domain blacklisting against malicious DNS registrations. We use a large dataset of malicious DNS registrations and compare them to a dataset of domain blacklists. We analyze the overlap between the two datasets and the impact of domain blacklisting on the visibility of malicious DNS registrations. Our results show that domain blacklisting is effective in reducing the visibility of malicious DNS registrations, but that it is not perfect. We discuss the reasons for this and provide recommendations for improving domain blacklisting.

2. BACKGROUND

Domain blacklisting is a widely used technique to prevent access to malicious websites. This paper evaluates the effectiveness of domain blacklisting against malicious DNS registrations. We use a large dataset of malicious DNS registrations and compare them to a dataset of domain blacklists. We analyze the overlap between the two datasets and the impact of domain blacklisting on the visibility of malicious DNS registrations. Our results show that domain blacklisting is effective in reducing the visibility of malicious DNS registrations, but that it is not perfect. We discuss the reasons for this and provide recommendations for improving domain blacklisting.

3. METHODOLOGY

Domain blacklisting is a widely used technique to prevent access to malicious websites. This paper evaluates the effectiveness of domain blacklisting against malicious DNS registrations. We use a large dataset of malicious DNS registrations and compare them to a dataset of domain blacklists. We analyze the overlap between the two datasets and the impact of domain blacklisting on the visibility of malicious DNS registrations. Our results show that domain blacklisting is effective in reducing the visibility of malicious DNS registrations, but that it is not perfect. We discuss the reasons for this and provide recommendations for improving domain blacklisting.

4. RESULTS

Domain blacklisting is a widely used technique to prevent access to malicious websites. This paper evaluates the effectiveness of domain blacklisting against malicious DNS registrations. We use a large dataset of malicious DNS registrations and compare them to a dataset of domain blacklists. We analyze the overlap between the two datasets and the impact of domain blacklisting on the visibility of malicious DNS registrations. Our results show that domain blacklisting is effective in reducing the visibility of malicious DNS registrations, but that it is not perfect. We discuss the reasons for this and provide recommendations for improving domain blacklisting.

5. CONCLUSIONS

Domain blacklisting is a widely used technique to prevent access to malicious websites. This paper evaluates the effectiveness of domain blacklisting against malicious DNS registrations. We use a large dataset of malicious DNS registrations and compare them to a dataset of domain blacklists. We analyze the overlap between the two datasets and the impact of domain blacklisting on the visibility of malicious DNS registrations. Our results show that domain blacklisting is effective in reducing the visibility of malicious DNS registrations, but that it is not perfect. We discuss the reasons for this and provide recommendations for improving domain blacklisting.


Ground truth analysis

T. Vissers et al., Assessing the Effectiveness of Domain Blacklisting Against Malicious DNS Registrations, IEEE Workshop on Traffic Measurements for Cybersecurity (WTMC 2019), May 2019.



Incompleteness of the blacklists

- › Failed to detect?
- › Never active/malicious?



	Active	Dormant
Blacklisted	Blocked	Pro-actively blocked
Non-blacklisted	Missed	Unused

Active vs Dormant – Blacklisted vs Non-blacklisted

- › 5 largest campaigns in .eu (Q1-Q2 2018)
- › Based on passively-logged DNS requests (.eu TLD server)

Active vs Dormant – Blacklisted vs Non-blacklisted

- › 5 largest campaigns in .eu (Q1-Q2 2018)
- › Based on passively-logged DNS requests (.eu TLD server)

	Active	Dormant
Blacklisted	Blocked 54.8%	Proactive 2.9%
Non-blacklisted	Missed 14.1%	Unused 14.0%

Active vs Dormant – Blacklisted vs Non-blacklisted

- › 5 largest campaigns in .eu (Q1-Q2 2018)
- › Based on passively-logged DNS requests (.eu TLD server)

	Active	Dormant
Blacklisted	Blocked 54.8%	Proactive 2.9%
Non-blacklisted	Missed 14.1%	Unused 14.0%



Key takeaways

Rather small set of bad actors

- › Up to 20 campaigns are responsible for 80% of malicious registrations

- › Top facilitators:
 - ›› About half of the malicious registrations via 1 registrar
 - ›› 1 public email provider are malicious with a high toxicity

Registration-time detection and prevention

- › Two prediction models predict at registration-time the malicious intent
- › Captures the majority of malicious domain registrations
- › Incompleteness of ground truth makes analysis hard
- › Interesting to see how this will further impact the security landscape

Detecting and preventing domain name abuse in .eu

Lieven Desmet, KU Leuven – Marc Van Wesemael, EURid