The background is a blurred image of a library with bookshelves. In the foreground, an open book is stacked on top of several closed books. Floating in the air are various white and yellow icons, including mathematical symbols like plus (+), minus (-), multiplication (x), and division (÷), as well as technical symbols like the Greek letter gamma (γ), the letter X, and question marks. There are also icons of a magnifying glass, a pencil, and a hand pointing. The overall theme is academic and technical.

ACADEMIC
THOUGHTS ON DATA
NEEDS FOR
FIGHTING DNS
ABUSE

Raffaele Sommese - University of Twente

r.sommese@utwente.nl

OARC 44 – 6-7 February 2025 - Atlanta

WHY THIS PRESENTATION AT OARC?

- OARC is not about naming, but about operations!
- But several operators are active actors in the fight against DNS namespace abuse.
- I often discuss about the data needs for fighting DNS namespace abuse -- data that needs to come from operators.
- Only operators know how to design this data sharing (or why they cannot share!)
- Also, I will introduce you to some nice by-product feeds and a lot of fun measurement orchestration insights :)

WHY THIS PRESENTATION AT OARC?

- OARC is not about naming, but about operations!
- But several operators are active actors in the fight against DNS namespace abuse.
- I often discuss about the data needs for fighting DNS namespace abuse -- data that needs to come from operators.
- Only operators know how to design this data sharing (or why they cannot share!)
- Also, I will introduce you to some nice by-product feeds and a lot of fun measurement orchestration insights :)

WHY THIS PRESENTATION AT OARC?

- OARC is not about naming, but about operations!
- But several operators are active actors in the fight against DNS namespace abuse.
- I often discuss about the data needs for fighting DNS namespace abuse -- data that needs to come from operators.
- Only operators know how to design this data sharing (or why they cannot share!)
- Also, I will introduce you to some nice by-product feeds and a lot of fun measurement orchestration insights :)

WHY THIS PRESENTATION AT OARC?

- OARC is not about naming, but about operations!
- But several operators are active actors in the fight against DNS namespace abuse.
- I often discuss about the data needs for fighting DNS namespace abuse -- data that needs to come from operators.
- Only operators know how to design this data sharing (or why they cannot share!)
- Also, I will introduce you to some nice by-product feeds and a lot of fun measurement orchestration insights :)

IN THE PREVIOUS FALL

DarkDNS: Revisiting the Value of Rapid Zone Update

Raffaele Sommese
University of Twente
Enschede, The Netherlands
r.sommese@utwente.nl

Gautam Akiwate
Stanford University
Stanford, CA, USA
gakiwate@cs.stanford.edu

Antonia Affinito
University of Twente
Enschede, The Netherlands
a.affinito@utwente.nl

Moritz Müller
SIDN Labs
Arnhem, The Netherlands
University of Twente
Enschede, The Netherlands
moritz.muller@sidn.nl

Mattijs Jonker
University of Twente
Enschede, The Netherlands
m.jonker@utwente.nl

KC Claffy
CAIDA
San Diego, CA, USA
kc@caida.org

Abstract

Malicious actors exploit the DNS namespace to launch spam campaigns, phishing attacks, malware, and other harmful activities. Combating these threats requires visibility into domain existence, ownership and nameservice activity that the DNS protocol does not itself provide. To facilitate visibility and security-related study of the expanding gTLD namespace, ICANN introduced the Centralized Zone Data Service (CZDS) that shares daily zone file snapshots

1 Introduction

Malicious actors exploit (abuse) the DNS namespace to launch spam campaigns, phishing attacks, malware, and other harmful activities. In many dimensions the DNS ecosystem is more opaque than other aspects of Internet transport. Unlike BGP, DNS is a pull protocol, so learning internal dynamics requires an entry point *i.e.*, a domain name. Without knowing this entry point, any abuse that lies behind a domain remains opaque to everyone except the targets.

TRANSIENT DOMAINS?!

By leveraging CT Logs (stream) we demonstrated that:

- We detected 42% of newly registered domains before they appeared in the CZDS snapshot (Nov 2023 – Jan 2024).
- ~76K domains per day - - almost 1 domain per second.
- 1% of newly registered domains never appear in the next CZDS snapshot.
- Predominately malicious!
- With half of them died within their first 6 hours of life.

TRANSIENT DOMAINS?!

By leveraging CT Logs (stream) we demonstrated that:

- We detected 42% of newly registered domains before they appeared in the CZDS snapshot (Nov 2023 – Jan 2024).
- ~76K domains per day - - almost 1 domain per second.
- 1% of newly registered domains never appear in the next CZDS snapshot.
- Predominately malicious!
- With half of them died within their first 6 hours of life.

TRANSIENT DOMAINS?!

By leveraging CT Logs (stream) we demonstrated that:

- We detected 42% of newly registered domains before they appeared in the CZDS snapshot (Nov 2023 – Jan 2024).
- ~76K domains per day - - almost 1 domain per second.
- 1% of newly registered domains never appear in the next CZDS snapshot.
- Predominately malicious!
- With half of them died within their first 6 hours of life.

FROM RESEARCH TO OPERATION

- The paper focused on detecting transient domains in a post-mortem scenario.
- Being able to block them quicker can help against attackers who may leverage caching.
- How can we detect them as soon as they die?
 - Measuring newly registered domains* (every 10 min, for 48h. A/AAAA/NS (@TLDs)/MX).
 - Detect deletion: 3x NXDOMAIN at TLD Level.
 - Issuing 2 RDAP requests, one when the domain is first detected, one when the domain is marked as deleted.
 - Checking RDAP responses for status.

FROM RESEARCH TO OPERATION

- The paper focused on detecting transient domains in a post-mortem scenario.
- Being able to block them quicker can help against attackers who may leverage caching.
- How can we detect them as soon as they die?
 - Measuring newly registered domains* (every 10 min, for 48h. A/AAAA/NS (@TLDs)/MX).
 - Detect deletion: 3x NXDOMAIN at TLD Level.
 - Issuing 2 RDAP requests, one when the domain is first detected, one when the domain is marked as deleted.
 - Checking RDAP responses for status.

FROM RESEARCH TO OPERATION

- The paper focused on detecting transient domains in a post-mortem scenario.
- Being able to block them quicker can help against attackers who may leverage caching.
- How can we detect them as soon as they die?
 - Measuring newly registered domains* (every 10 min, for 48h. A/AAAA/NS (@TLDs)/MX).
 - Detect deletion: 3x NXDOMAIN at TLD Level.
 - Issuing 2 RDAP requests, one when the domain is first detected, one when the domain is marked as deleted.
 - Checking RDAP responses for status.

FROM RESEARCH TO OPERATION

- The paper focused on detecting transient domains in a post-mortem scenario.
- Being able to block them quicker can help against attackers who may leverage caching.
- How can we detect them as soon as they die?
 - Measuring newly registered domains* (every 10 min, for 48h. A/AAAA/NS (@TLDs)/MX).
 - Detect deletion: 3x NXDOMAIN at TLD Level.
 - Issuing 2 RDAP requests, one when the domain is first detected, one when the domain is marked as deleted.
 - Checking RDAP responses for status.

...AND

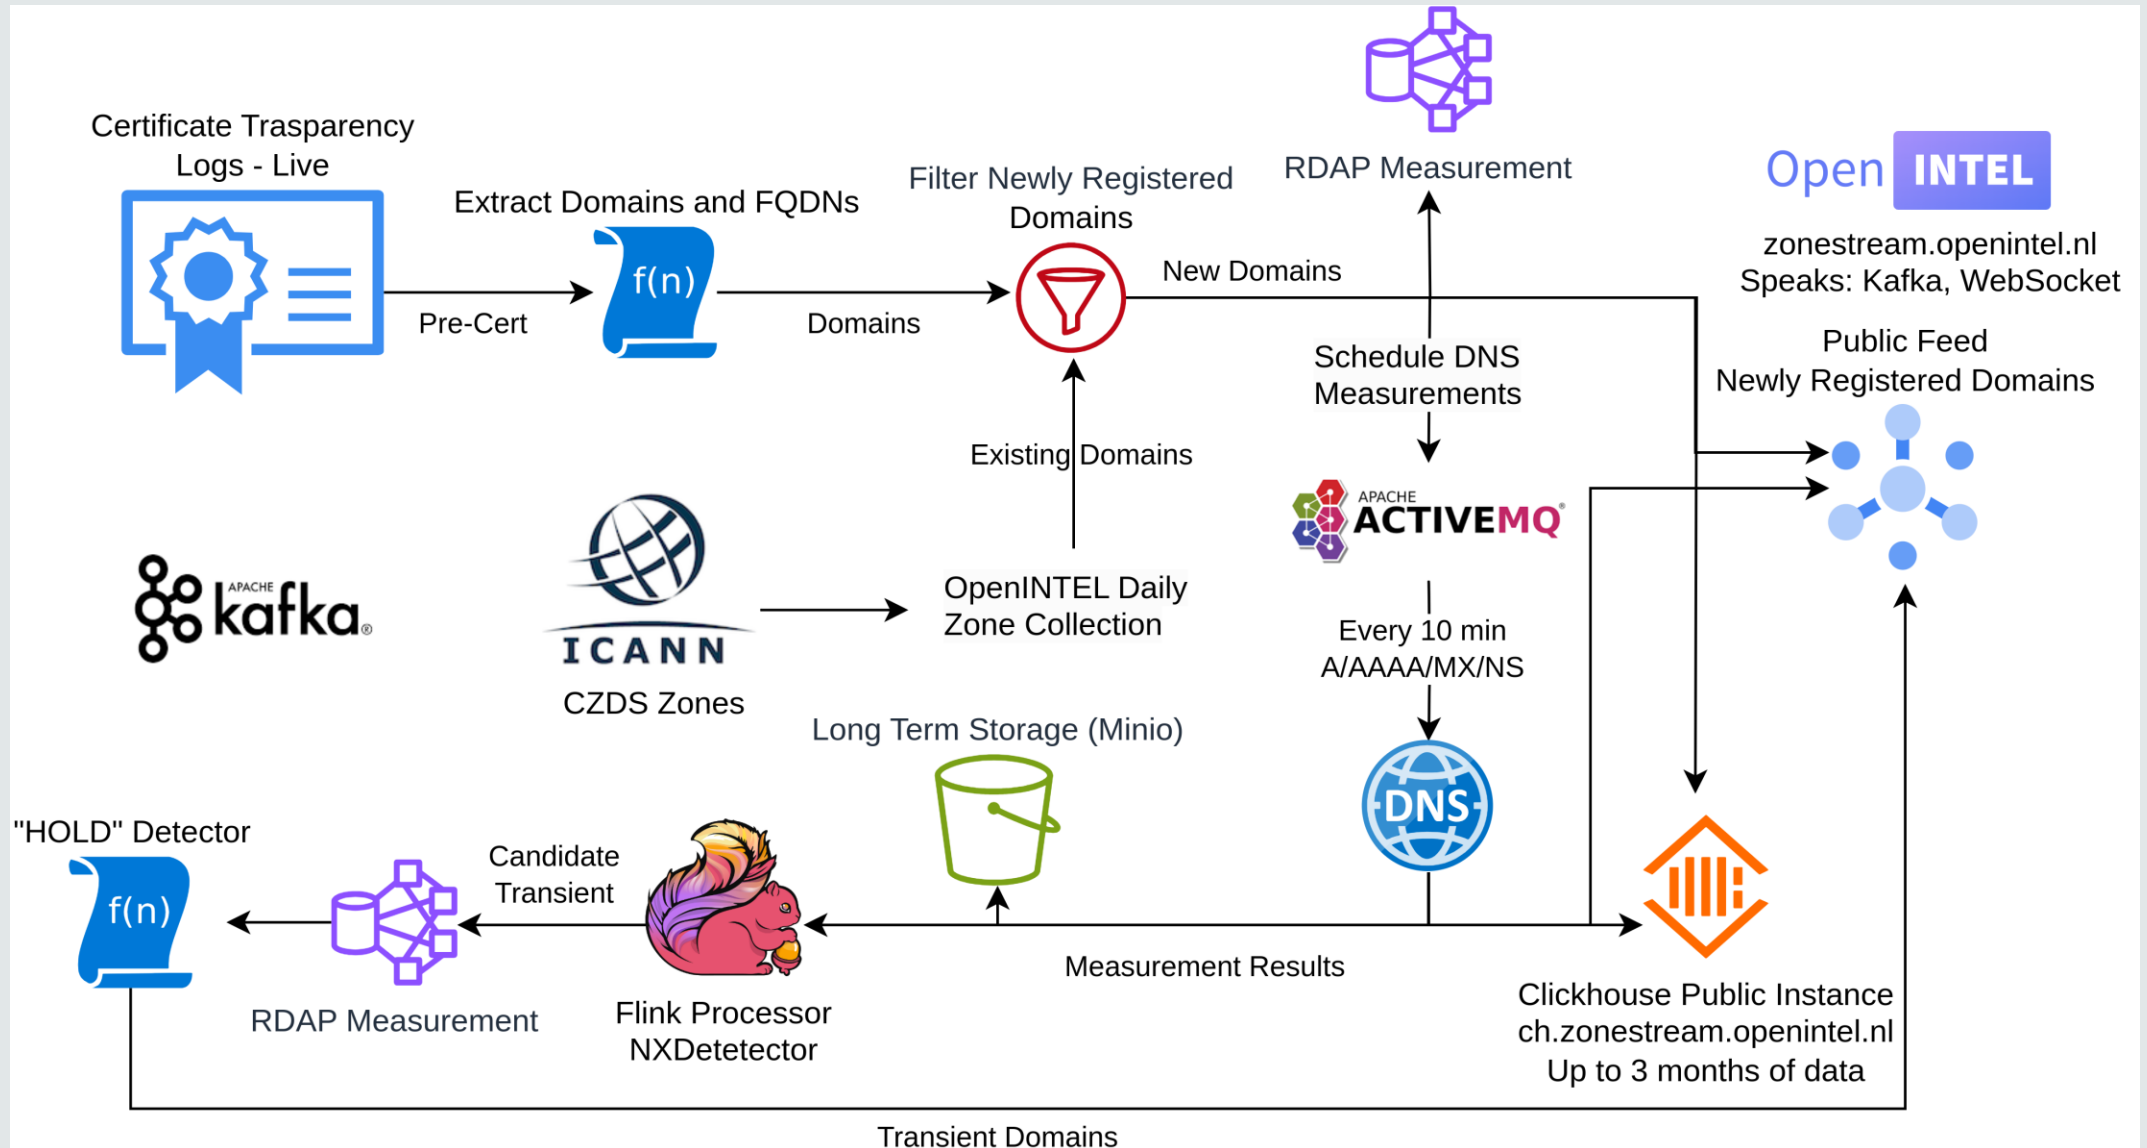
It seems someone (the registry) is holding those transient domains back.

Likely abuse!

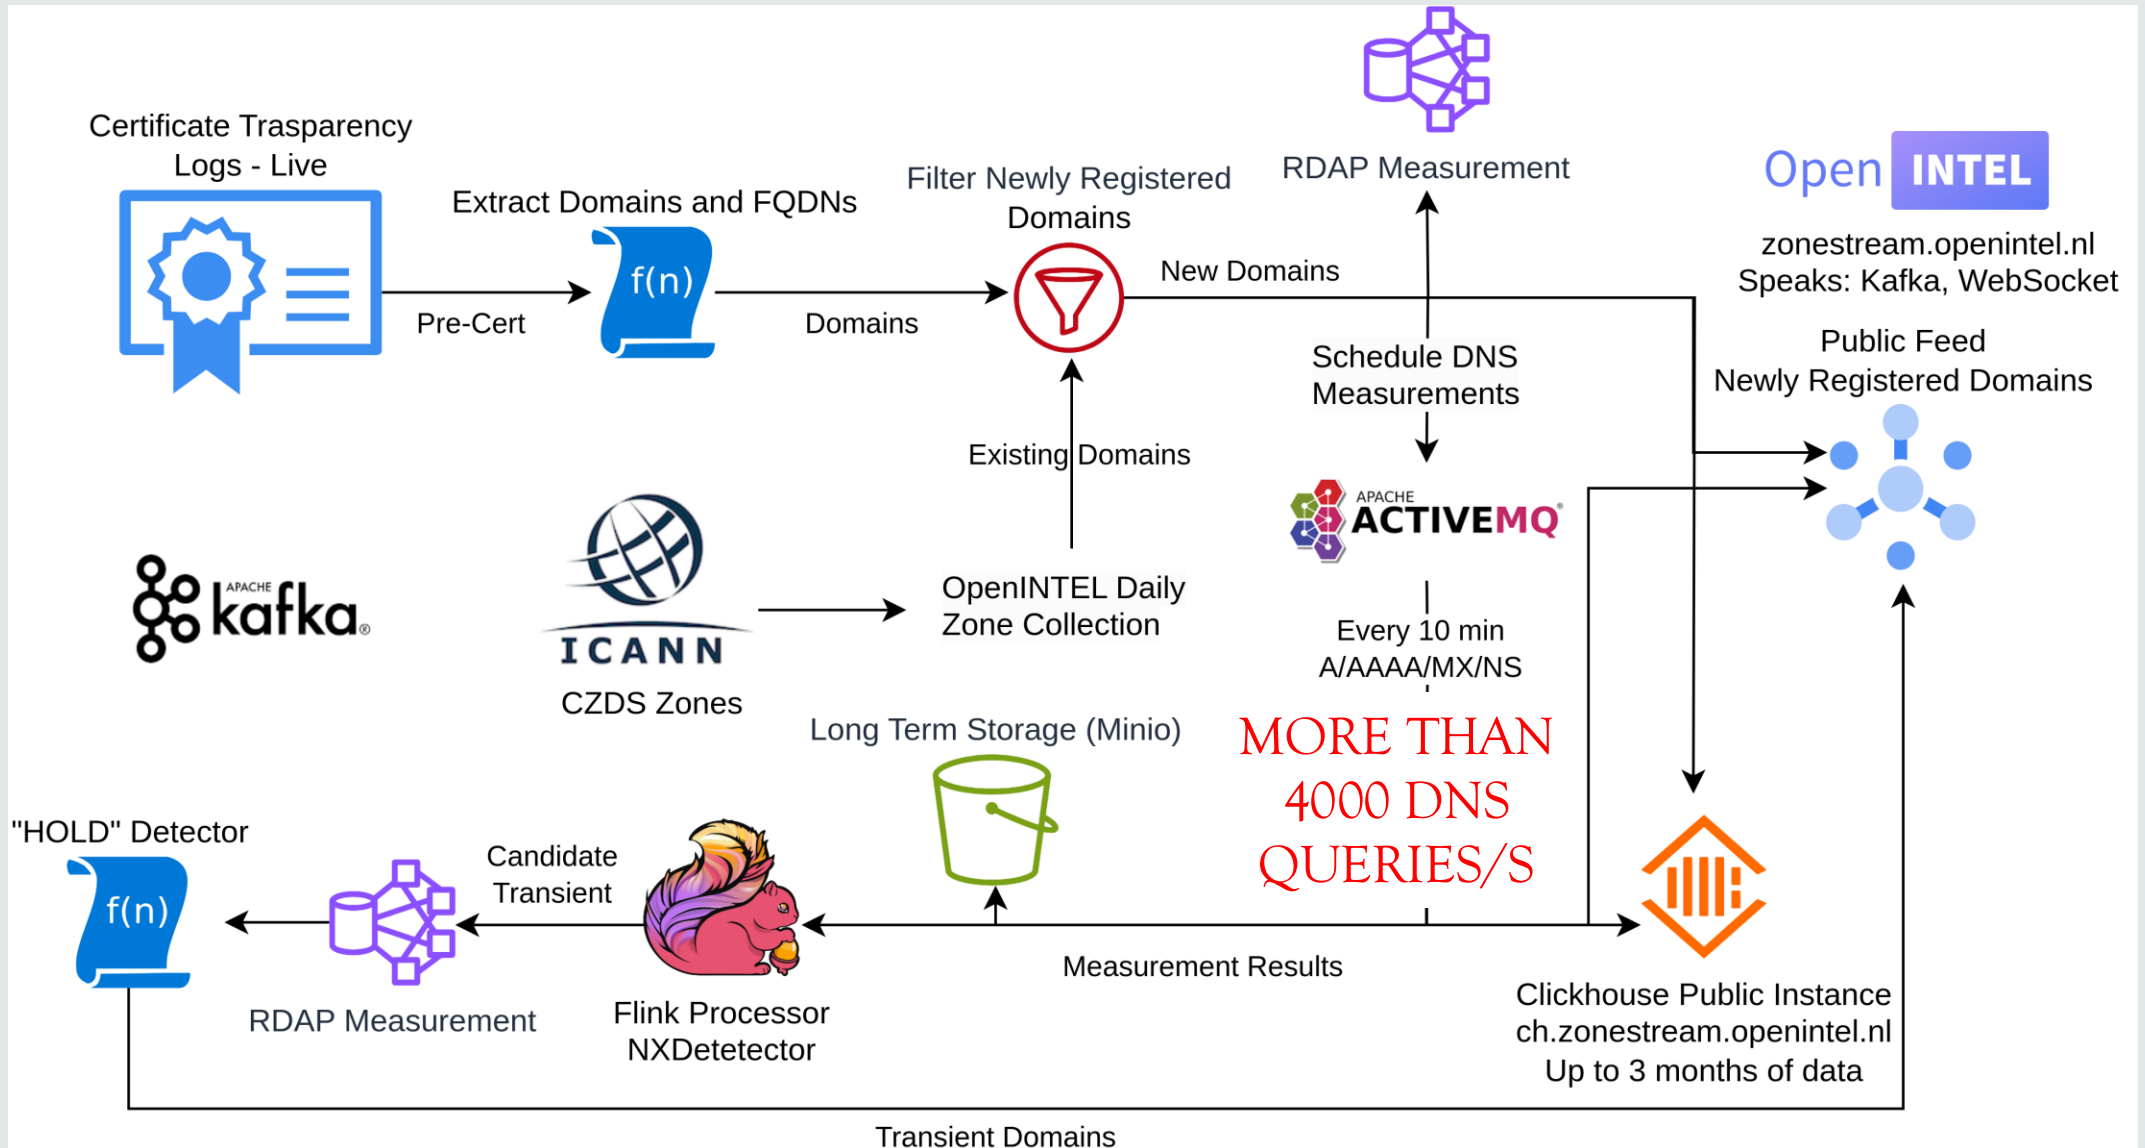
(Sometimes is client hold, registrar acting)

```
▼ rdap:
  copyright_notice: ""
  description: []
  dnssec: false
  ► entities: [{"abuse":{"email":"abuse@gnome.com","handle":"NOT APPLICABLE","name":"Gnome.com Pte. Ltd","type":"entity"},"registrant":{"address":{"Ltd","type":"entity","url":"https://rdap.centralnic.com/store/entity/1923"},"technical":{"handle":"C1559112222-CNIC","name":"","type":"entity"}}}]
  expiration_date: "2026-01-20 23:59:59+00:00"
  handle: "D520613940-CNIC"
  last_changed_date: "2025-01-20 12:48:56+00:00"
  name: "aocjjo223.store"
  ► nameservers: ["a.share-dns.com","b.share-dns.net"]
  parent_handle: ""
  registration_date: "2025-01-20 09:50:17+00:00"
  rir: ""
  ► status: ["server hold","server transfer prohibited","client transfer prohibited","add period"]
  terms_of_service_url: "https://www.centralnicregistry.com/"
  type: "domain"
  unicode_name: ""
  url: "https://rdap.centralnic.com/store/domain/aocjjo223.store"
  whois_server: "whois.nic.store"
  success: true
  ► tags: {"authoritative_level":"TLD","cert_index":"529532406","cert_timestamp":"1737376899","ctlog_name":"Let's Encrypt 'Oak2025h1","end_time":"2025
  timestamp: 1737384500
```

GETTING TO THIS....



GETTING TO THIS....



SOME NICE BY-PRODUCTS AS LIVE STREAM

- newly_registered_domain – List of newly registered domains detected from CT Logs.
- confirmed_newly_registered_domain – Same of above, RDAP-Confirmed.
- newly_registered_fqdn – List of FQDNs of newly registered domains detected from CT Logs.
- newly_registered_domains_measurements – A/AAAA/NS/MX (every 10 min for 48h).
- newly_registered_fqdn_measurements – A/AAAA (every 10 min for 48h).
- certstream – Certificate Transparency Logs stream (based on <https://certstream.calidog.io/>).
- certstream_domains – Stream of domain names learned from CT Logs.
- newly_issued_certificates_measurements - A/AAAA/NS for every domain name learned from CT Logs.

Publicly available as best-effort and for non-for-profit usage only.
How to access: <https://kafka.zonestream.openintel.nl/>

SOME NICE BY-PRODUCTS AS LIVE STREAM

- newly_registered_domain – List of newly registered domains detected from CT Logs.
- confirmed_newly_registered_domain – Same of above, RDAP-Confirmed.
- newly_registered_fqdn – List of FQDNs of newly registered domains detected from CT Logs.
- newly_registered_domains_measurements – A/AAAA/NS/MX (every 10 min for 48h).
- newly_registered_fqdn_measurements – A/AAAA (every 10 min for 48h).
- certstream – Certificate Transparency Logs stream (based on <https://certstream.calidog.io/>).
- certstream_domains – Stream of domain names learned from CT Logs.
- newly_issued_certificates_measurements - A/AAAA/NS for every domain name learned from CT Logs.

Publicly available as best-effort and for non-for-profit usage only.
How to access: <https://kafka.zonestream.openintel.nl/>

SOME NICE BY-PRODUCTS AS LIVE STREAM

- newly_registered_domain – List of newly registered domains detected from CT Logs.
- confirmed_newly_registered_domain – Same of above, RDAP-Confirmed.
- newly_registered_fqdn – List of FQDNs of newly registered domains detected from CT Logs.
- newly_registered_domains_measurements – A/AAAA/NS/MX (every 10 min for 48h).
- newly_registered_fqdn_measurements – A/AAAA (every 10 min for 48h).
- certstream – Certificate Transparency Logs stream (based on <https://certstream.calidog.io/>).
- certstream_domains – Stream of domain names learned from CT Logs.
- newly_issued_certificates_measurements - A/AAAA/NS for every domain name learned from CT Logs.

Publicly available as best-effort and for non-for-profit usage only.
How to access: <https://kafka.zonestream.openintel.nl/>

SOME NICE BY-PRODUCTS AS LIVE STREAM

- newly_registered_domain – List of newly registered domains detected from CT Logs.
- confirmed_newly_registered_domain – Same of above, RDAP-Confirmed.
- newly_registered_fqdn – List of FQDNs of newly registered domains detected from CT Logs.
- newly_registered_domains_measurements – A/AAAA/NS/MX (every 10 min for 48h).
- newly_registered_fqdn_measurements – A/AAAA (every 10 min for 48h).
- certstream – Certificate Transparency Logs stream (based on <https://certstream.calidog.io/>).
- certstream_domains – Stream of domain names learned from CT Logs.
- newly_issued_certificates_measurements - A/AAAA/NS for every domain name learned from CT Logs.

Publicly available as best-effort and for non-for-profit usage only.
How to access: <https://kafka.zonestream.openintel.nl/>

A close-up photograph of a person's hand holding a vintage-style compass. The hand is wearing a teal-colored long-sleeved shirt. The compass is round with a white face and a black bezel, showing cardinal and intercardinal directions (N, S, E, W, NE, SE, SW, NW) and degree markings. The background is a blurred view of a road intersection with a white stop line, set against a backdrop of dry, brownish hills under a clear sky. The lighting is bright, suggesting a sunny day.

ISN'T THERE A
BETTER WAY?

DNS: FROM PULL TO PUSH

- DNS is a *pull protocol*!?! Right?
- (Rapid) Zone Updates? -> EPP
- A signal for any update in the zone, enabling a CT Logs-like system to detect changes.
- DNS-Transparency -> <https://www.internetfire.org/projects/dns-transparency>
- Why?
 - Detecting early removal of domain names.
 - Evicting from cache of public and private resolvers.
 - Revoking certificates of expired domains.
 - Build better ML models to detect future threats.
 - Identifying hijacking
 - ...add your use case that requires you monitoring continuously the DNS

DNS: FROM PULL TO PUSH

- DNS is a *pull protocol*!?! Right?
- (Rapid) Zone Updates? -> EPP
- A signal for any update in the zone, enabling a CT Logs-like system to detect changes.
- DNS-Transparency -> <https://www.internetfire.org/projects/dns-transparency>
- Why?
 - Detecting early removal of domain names.
 - Evicting from cache of public and private resolvers.
 - Revoking certificates of expired domains.
 - Build better ML models to detect future threats.
 - Identifying hijacking
 - ...add your use case that requires you monitoring continuously the DNS

DNS: FROM PULL TO PUSH

- DNS is a *pull protocol*!? Right?
- (Rapid) Zone Updates? -> EPP
- A signal for any update in the zone, enabling a CT Logs-like system to detect changes.
- DNS-Transparency -> <https://www.internetfire.org/projects/dns-transparency>
- Why?
 - Detecting early removal of domain names.
 - Evicting from cache of public and private resolvers.
 - Revoking certificates of expired domains.
 - Build better ML models to detect future threats.
 - Identifying hijacking
 - ...add your use case that requires you monitoring continuously the DNS

DNS: FROM PULL TO PUSH

- DNS is a *pull protocol*!? Right?
- (Rapid) Zone Updates? -> EPP
- A signal for any update in the zone, enabling a CT Logs-like system to detect changes.
- DNS-Transparency -> <https://www.internetfire.org/projects/dns-transparency>
- Why?
 - Detecting early removal of domain names.
 - Evicting from cache of public and private resolvers.
 - Revoking certificates of expired domains.
 - Build better ML models to detect future threats.
 - Identifying hijacking
 - ...add your use case that requires you monitoring continuously the DNS

DNS: FROM PULL TO PUSH

- DNS is a *pull protocol*!? Right?
- (Rapid) Zone Updates? -> EPP
- A signal for any update in the zone, enabling a CT Logs-like system to detect changes.
- DNS-Transparency -> <https://www.internetfire.org/projects/dns-transparency>
- Why?
 - Detecting early removal of domain names.
 - Evicting from cache of public and private resolvers.
 - Revoking certificates of expired domains.
 - Build better ML models to detect future threats.
 - Identifying hijacking
 - ...add your use case that requires you monitoring continuously the DNS

DNS: FROM PULL TO PUSH

- DNS is a *pull protocol*!? Right?
- (Rapid) Zone Updates? -> EPP
- A signal for any update in the zone, enabling a CT Logs-like system to detect changes.
- DNS-Transparency -> <https://www.internetfire.org/projects/dns-transparency>
- Why?
 - Detecting early removal of domain names.
 - Evicting from cache of public and private resolvers.
 - Revoking certificates of expired domains.
 - Build better ML models to detect future threats.
 - Identifying hijacking
 - ...add your use case that requires you monitoring continuously the DNS



WILL THIS BE ENOUGH
TO MITIGATE ABUSE?



WILL THIS BE ENOUGH
TO MITIGATE ABUSE?

DEFINITELY NOT!

WE MAY NEED TO GO
BEYOND DNS

BEYOND DNS: WHERE? WHO? MONEY?

- Where?
 - Exposing the registrar information in a more public manner (e.g., in a DNS Transparency Log) will enable researchers and security firms to detect and identify bad registrars' behaviors (... someone said bad CAs identified by CT Logs?)
- Who?
 - GDPR complicate the matter veiling RDAP registrant information's behind privacy. But we need the ability to correlate behaviors of actors across different system, not of knowing their identity!
- How to follow the money?
 - Registrant may still lie about their identity, the (global, unique, anonymous) identifier should enable independent third-party, vetted security researcher, law enforcement to track the money flow (e.g., a hash linked to the payment method).

BEYOND DNS: WHERE? WHO? MONEY?

- Where?
 - Exposing the registrar information in a more public manner (e.g., in a DNS Transparency Log) will enable researchers and security firms to detect and identify bad registrars' behaviors (... someone said bad CAs identified by CT Logs?)
- Who?
 - GDPR complicate the matter veiling RDAP registrant information's behind privacy. But we need the ability to correlate behaviors of actors across different system, not of knowing their identity!
- How to follow the money?
 - Registrant may still lie about their identity, the (global, unique, anonymous) identifier should enable independent third-party, vetted security researcher, law enforcement to track the money flow (e.g., a hash linked to the payment method).

BEYOND DNS: WHERE? WHO? MONEY?

- Where?
 - Exposing the registrar information in a more public manner (e.g., in a DNS Transparency Log) will enable researchers and security firms to detect and identify bad registrars' behaviors (... someone said bad CAs identified by CT Logs?)
- Who?
 - GDPR complicate the matter veiling RDAP registrant information's behind privacy. But we need the ability to correlate behaviors of actors across different system, not of knowing their identity!
- How to follow the money?
 - Registrant may still lie about their identity, the (global, unique, anonymous) identifier should enable independent third-party, vetted security researcher, law enforcement to track the money flow (e.g., a hash linked to the payment method).

WILL THIS SOLVE THE PROBLEM OF ABUSE?

- Probably not!
- Malicious actor may still evade detection by using fake identities.
- But we will make their life harder (and ours easier)!
- We need a zone of trust where those information can be shared and design a mechanism to enable anonymization of those identifiers without violating the privacy of the users.

WILL THIS SOLVE THE PROBLEM OF ABUSE?

- Probably not!
- Malicious actor may still evade detection by using fake identities.
- But we will make their life harder (and ours easier)!
- We need a zone of trust where those information can be shared and design a mechanism to enable anonymization of those identifiers without violating the privacy of the users.

- Or we can keep playing the cat and the mouse game!



QUESTIONS?

Raffaele Sommese

r.sommese@utwente.nl

<https://zonestream.openintel.nl>

<https://kafka.zonestream.openintel.nl/>



**UNIVERSITY
OF TWENTE.**

