Chair of Network Architectures and Services
School of Computation, Information, and Technology
Technical University of Munich

Propagating Threat Scores With a TLS Ecosystem Graph Model Derived by Active Measurements

Markus Sosnowski, Patrick Sattler, Johannes Zirngibl, Tim Betzer, and Georg Carle

Thursday 23rd May, 2024

Network Traffic Measurement and Analysis Conference 2024

Chair of Network Architectures and Services
School of Computation, Information, and Technology
Technical University of Munich

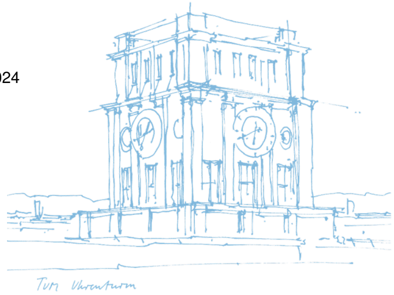

Motivation

Active Internet-wide DNS and TLS measurements can provide new information on known threats:

Motivation

Active Internet-wide DNS and TLS measurements can provide new information on known threats:

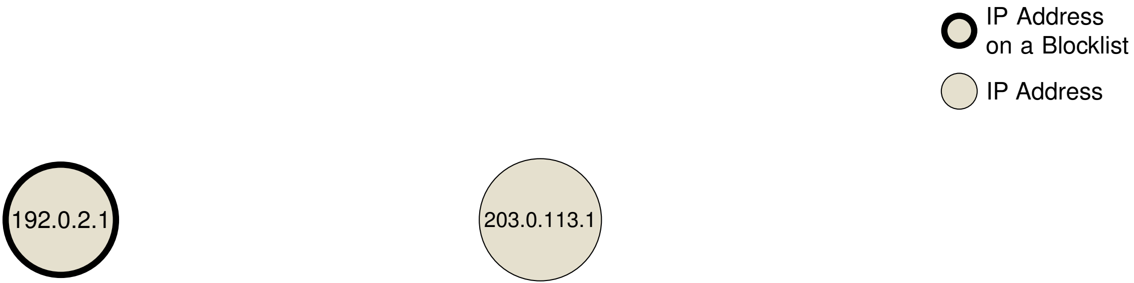

Should we block **`203.0.113.1`**?

Motivation

Active Internet-wide DNS and TLS measurements can provide new information on known threats:

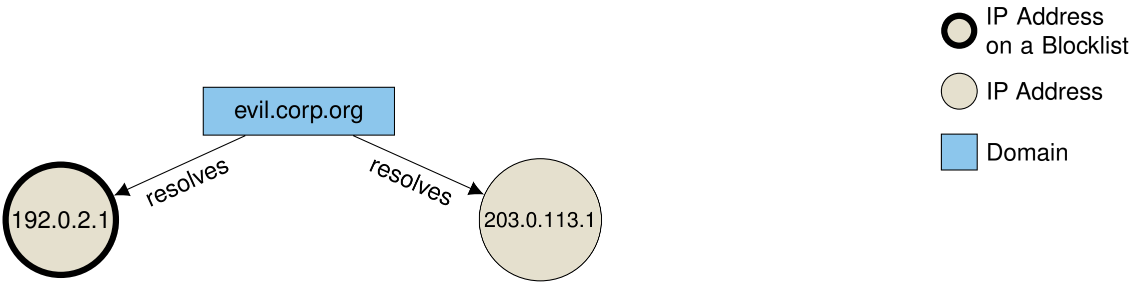

Should we block **`203.0.113.1`**?

Motivation

Active Internet-wide DNS and TLS measurements can provide new information on known threats:

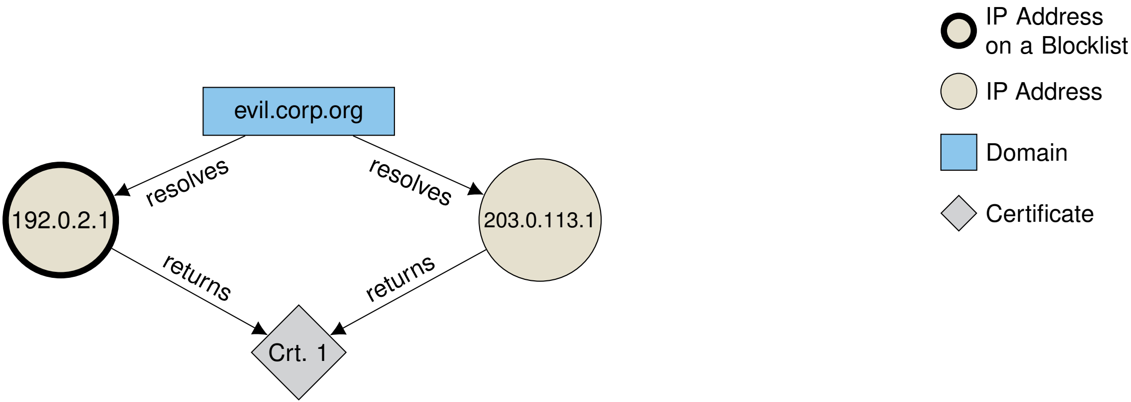

Should we block **`203.0.113.1`**?

Active Internet-wide DNS and TLS measurements can provide new information on known threats:

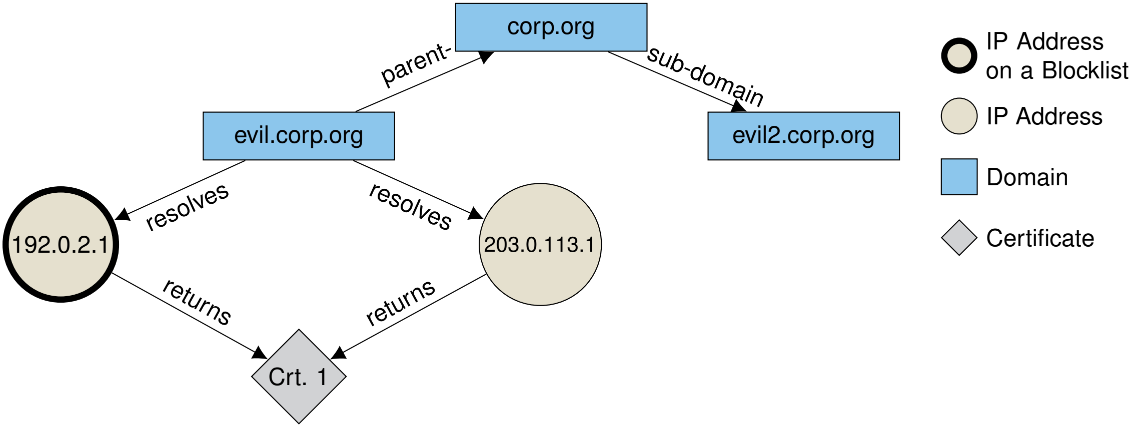

Should we block **`203.0.113.1`**? What about the domains?

Motivation

Active Internet-wide DNS and TLS measurements can provide new information on known threats:

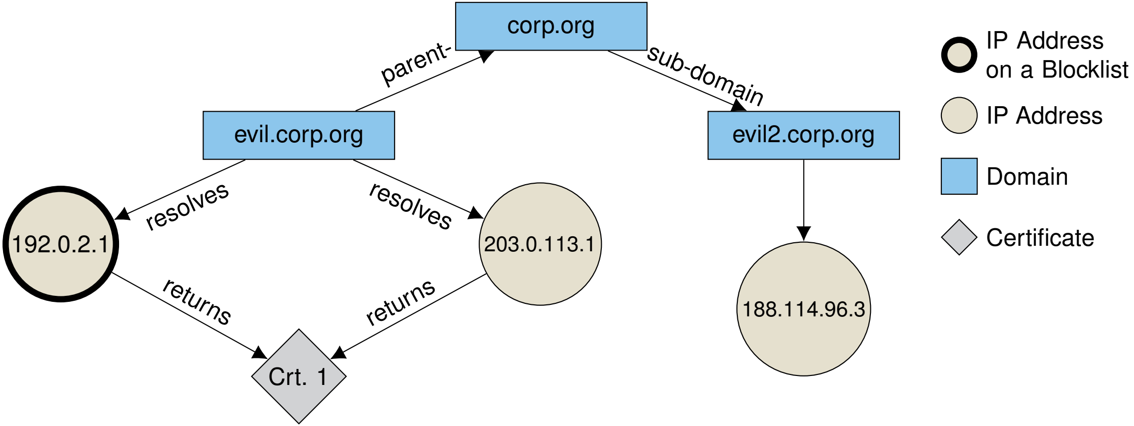

Should we block **203.0.113.1**? What about the domains? What about **188.113.96.3**?

Motivation

Active Internet-wide DNS and TLS measurements can provide new information on known threats:

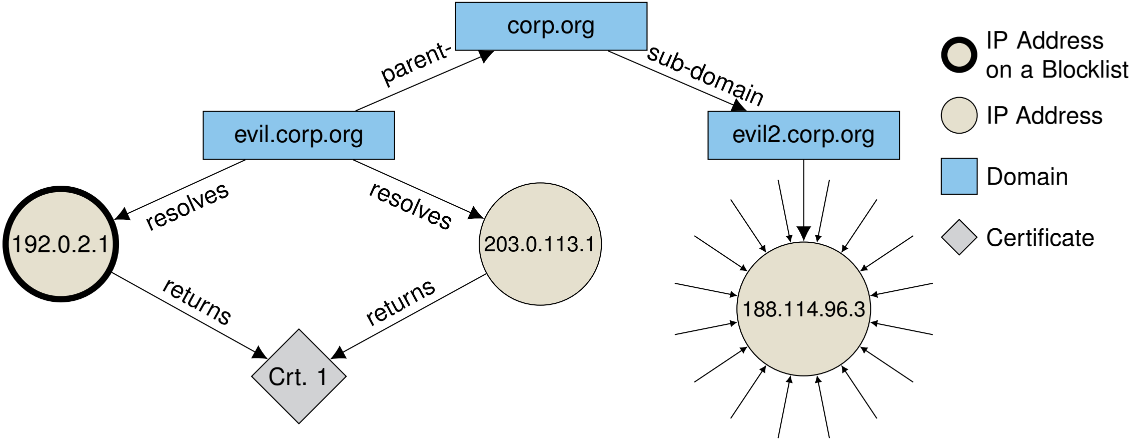

Should we block **`203.0.113.1`**? What about the domains? What about **`188.113.96.3`**?

Challenge: “Internet-wide” is quite large

Challenge: “Internet-wide” is quite large

An Internet-wide TLS scan from Jan. 2024

Type	Count
Domains	628 M
IPv4 TLS Handshakes	608 M
IPv6 TLS Handshakes	146 M

Motivation

Challenge: “Internet-wide” is quite large

An Internet-wide TLS scan from Jan. 2024

Type	Count
Domains	628 M
IPv4 TLS Handshakes	608 M
IPv6 TLS Handshakes	146 M

$\Longrightarrow$

- Any algorithm used on such large datasets has to scale!
- $O(n)$ or faster

Methodology

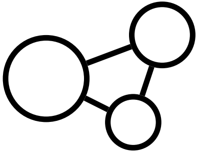

Modeling the TLS Ecosystem as Graph

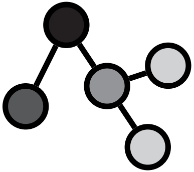

Propagating Threat Scores

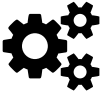

An Internet-wide TLS Scanning Pipeline

Methodology

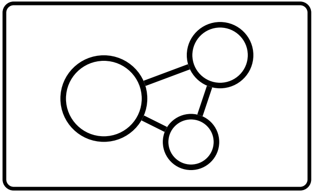

Modeling the TLS Ecosystem as Graph

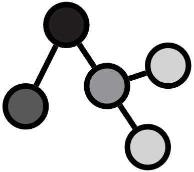

Propagating Threat Scores

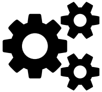

An Internet-wide TLS Scanning Pipeline

Modeling the TLS Ecosystem as Graph

- The Internet is a network, modeling collected data as a graph is intuitive

Modeling the TLS Ecosystem as Graph

- The Internet is a network, modeling collected data as a graph is intuitive
- The generalized structure allows applying standard graph algorithms

Modeling the TLS Ecosystem as Graph

- The Internet is a network, modeling collected data as a graph is intuitive
- The generalized structure allows applying standard graph algorithms
- Labeled property graph:

Modeling the TLS Ecosystem as Graph

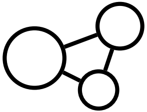

- The Internet is a network, modeling collected data as a graph is intuitive
- The generalized structure allows applying standard graph algorithms
- Labeled property graph:
  - Data is represented as nodes and edges

Modeling the TLS Ecosystem as Graph

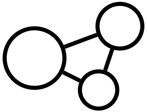

- The Internet is a network, modeling collected data as a graph is intuitive
- The generalized structure allows applying standard graph algorithms
- Labeled property graph:
  - Data is represented as nodes and edges
  - Nodes and edges are labeled and can have arbitrary properties

Modeling the TLS Ecosystem as Graph

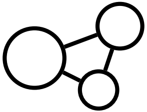

- The Internet is a network, modeling collected data as a graph is intuitive
- The generalized structure allows applying standard graph algorithms
- Labeled property graph:
  - Data is represented as nodes and edges
  - Nodes and edges are labeled and can have arbitrary properties
  - Edges are directed

Modeling the TLS Ecosystem as Graph

Designing the graph schema:

Modeling the TLS Ecosystem as Graph

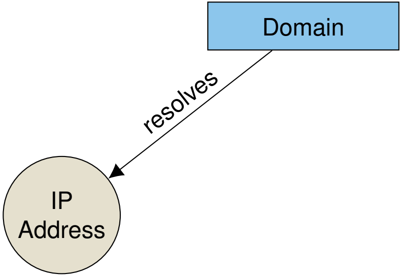

Designing the graph schema:

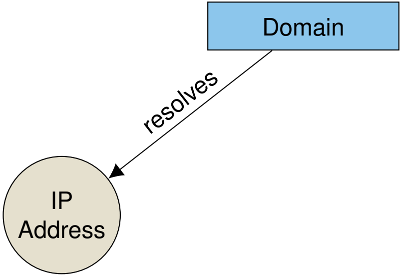

Designing the graph schema:

- directions in the graph should reflect deliberate actions of the actor controlling a node

Modeling the TLS Ecosystem as Graph

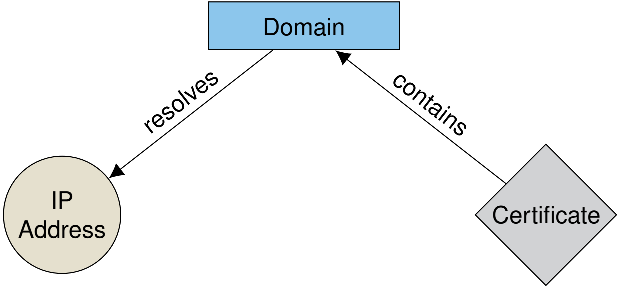

Designing the graph schema:

- directions in the graph should reflect deliberate actions of the actor controlling a node

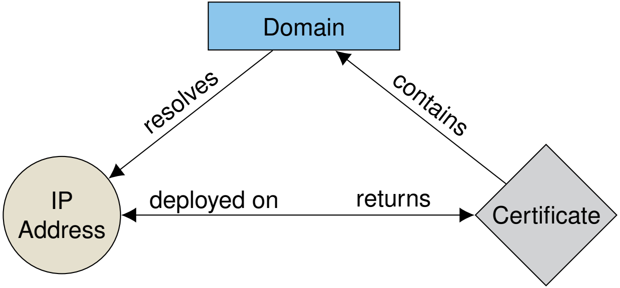

Designing the graph schema:

- directions in the graph should reflect deliberate actions of the actor controlling a node

Modeling the TLS Ecosystem as Graph

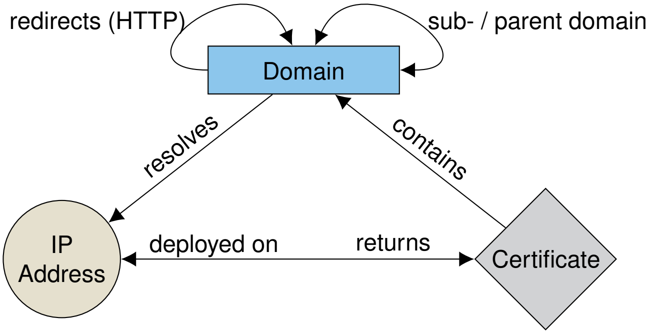

Designing the graph schema:

- directions in the graph should reflect deliberate actions of the actor controlling a node

Methodology

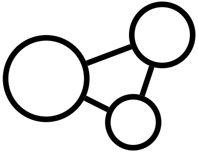

Modeling the TLS Ecosystem as Graph

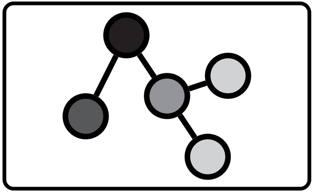

Propagating Threat Scores

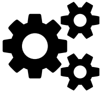

An Internet-wide TLS Scanning Pipeline

Propagating Threat Scores

The Probabilistic Threat Propagation (PTP) [1] algorithm:

- PTP meets our intuition how scores should propagate (considers locality and edge directions)

[1] K. M. Carter et al., "Probabilistic threat propagation for malicious activity detection," in Proc. IEEE Int. Conference on Acoustics, Speech and Signal Processing, 2013

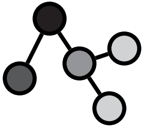

The Probabilistic Threat Propagation (PTP) [1] algorithm:

- PTP meets our intuition how scores should propagate (considers locality and edge directions)
- it's fast $O(n)$

[1] K. M. Carter et al., "Probabilistic threat propagation for malicious activity detection," in Proc. IEEE Int. Conference on Acoustics, Speech and Signal Processing, 2013

Propagating Threat Scores

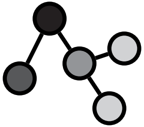

The Probabilistic Threat Propagation (PTP) [1] algorithm:

- PTP meets our intuition how scores should propagate (considers locality and edge directions)
- it's fast $O(n)$
- we can use existing blocklists as input

[1] K. M. Carter et al., "Probabilistic threat propagation for malicious activity detection," in Proc. IEEE Int. Conference on Acoustics, Speech and Signal Processing, 2013

Propagating Threat Scores

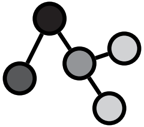

The Probabilistic Threat Propagation (PTP) [1] algorithm:

- PTP meets our intuition how scores should propagate (considers locality and edge directions)
- it's fast $O(n)$
- we can use existing blocklists as input
- highly connected nodes (e.g., from CDNs) will automatically get low scores

[1] K. M. Carter et al., "Probabilistic threat propagation for malicious activity detection," in Proc. IEEE Int. Conference on Acoustics, Speech and Signal Processing, 2013

Propagating Threat Scores

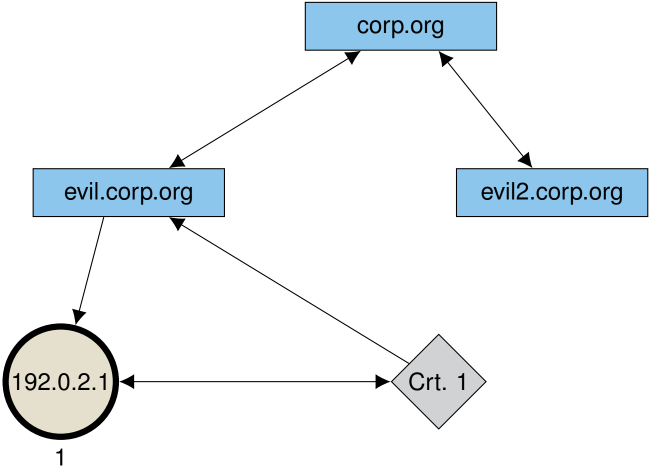

Message-based approximate PTP:

- the input has a fixed score of one (e.g., nodes on a blocklist)

Propagating Threat Scores

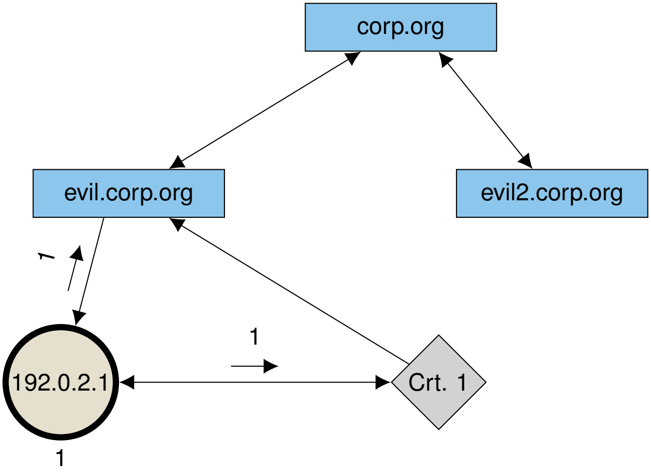

Message-based approximate PTP:

- the input has a fixed score of one (e.g., nodes on a blocklist)
- each node sends its score to neighbors in reversed graph direction

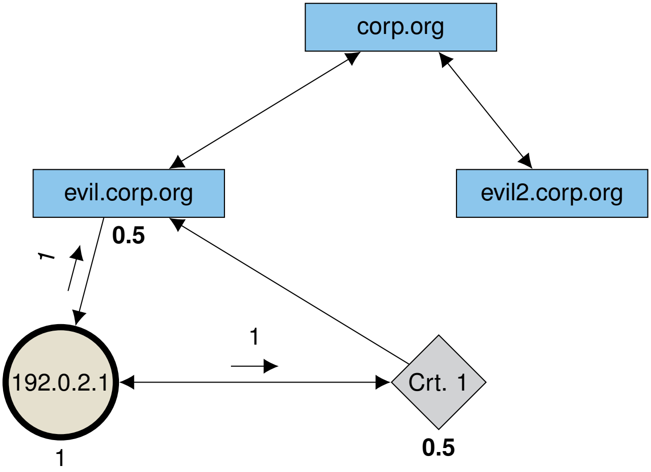

Message-based approximate PTP:

- the input has a fixed score of one (e.g., nodes on a blocklist)
- each node sends its score to neighbors in reversed graph direction
- a node will get the average of received scores as new score

Propagating Threat Scores

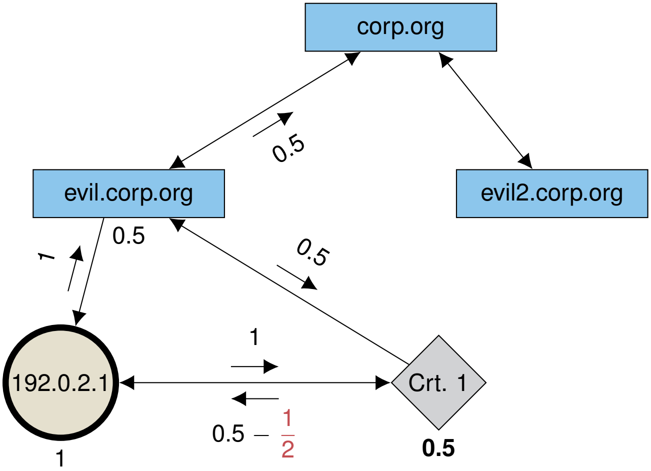

Message-based approximate PTP:

- the input has a fixed score of one (e.g., nodes on a blocklist)
- each node sends its score to neighbors in reversed graph direction
- a node will get the average of received scores as new score
- repeat until convergence

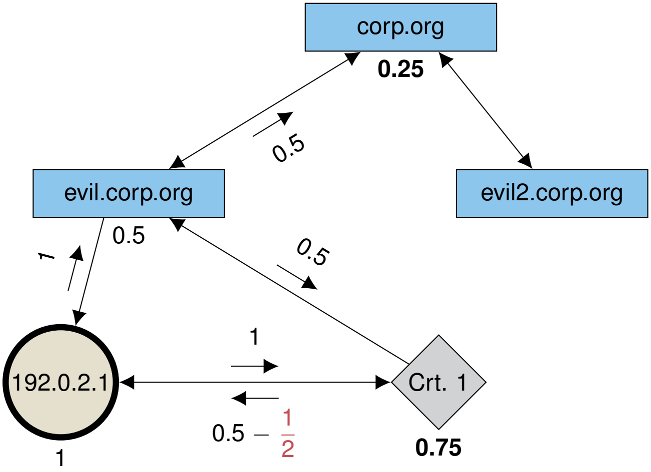

Message-based approximate PTP:

- the input has a fixed score of one (e.g., nodes on a blocklist)
- each node sends its score to neighbors in reversed graph direction
- a node will get the average of received scores as new score
- repeat until convergence

Propagating Threat Scores

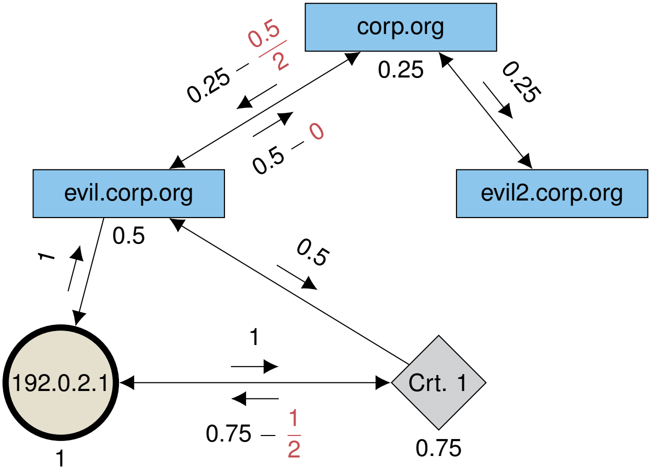

Message-based approximate PTP:

- the input has a fixed score of one (e.g., nodes on a blocklist)
- each node sends its score to neighbors in reversed graph direction
- a node will get the average of received scores as new score
- repeat until convergence
- core aspect of PTP is to minimize the **error** introduced by nodes falsely increasing their own score

Propagating Threat Scores

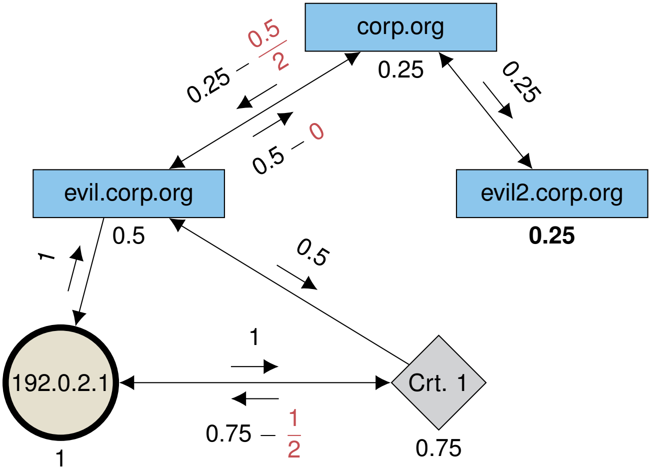

Message-based approximate PTP:

- the input has a fixed score of one (e.g., nodes on a blocklist)
- each node sends its score to neighbors in reversed graph direction
- a node will get the average of received scores as new score
- repeat until convergence
- core aspect of PTP is to minimize the **error** introduced by nodes falsely increasing their own score

Methodology

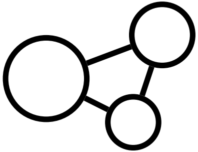

Modeling the TLS Ecosystem as Graph

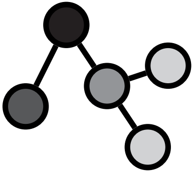

Propagating Threat Scores

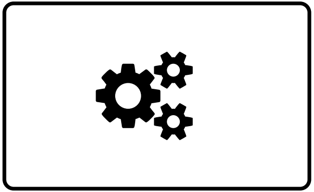

An Internet-wide TLS Scanning Pipeline

An Internet-wide TLS Scanning Pipeline

Internet-wide measurements at GINO[1]:

- Special interest group since 2016

[1] https://net.in.tum.de/projects/gino/

An Internet-wide TLS Scanning Pipeline

various domain sources

Internet-wide measurements at GINO[1]:

- Special interest group since 2016
- Among others: Internet-wide DNS, TLS, HTTPS scans on port 443

[1] https://net.in.tum.de/projects/gino/

An Internet-wide TLS Scanning Pipeline

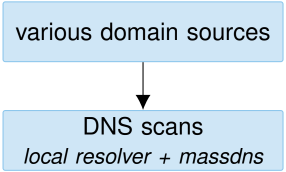

Internet-wide measurements at GINO[1]:

- Special interest group since 2016
- Among others: Internet-wide DNS, TLS, HTTPS scans on port 443

[1] https://net.in.tum.de/projects/gino/

An Internet-wide TLS Scanning Pipeline

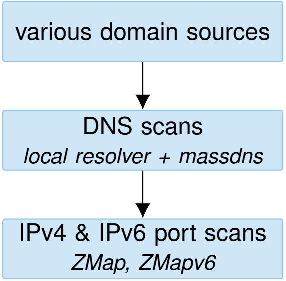

Internet-wide measurements at GINO[1]:

- Special interest group since 2016
- Among others: Internet-wide DNS, TLS, HTTPS scans on port 443

[1] `https://net.in.tum.de/projects/gino/`

An Internet-wide TLS Scanning Pipeline

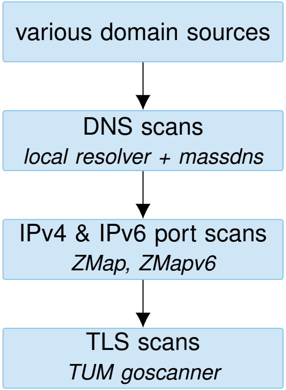

Internet-wide measurements at GINO[1]:

- Special interest group since 2016
- Among others: Internet-wide DNS, TLS, HTTPS scans on port 443

[1] https://net.in.tum.de/projects/gino/

An Internet-wide TLS Scanning Pipeline

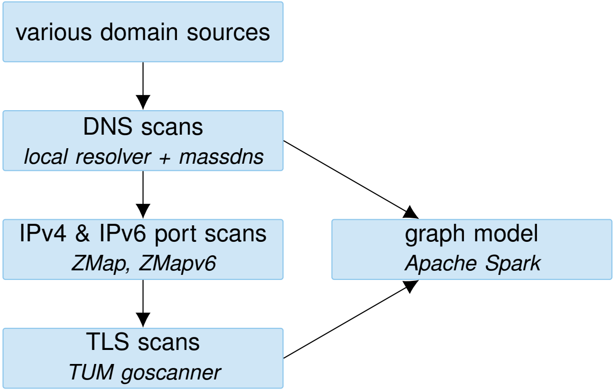

Internet-wide measurements at GINO[1]:

- Special interest group since 2016
- Among others: Internet-wide DNS, TLS, HTTPS scans on port 443
- New: Apache spark app to merge our scans and construct the graph model

[1] `https://net.in.tum.de/projects/gino/`

Results

We created 13 monthly Internet-wide TLS Ecosystem Graphs throughout the last year[2]

[2] starting Jan. 2023

We created 13 monthly Internet-wide TLS Ecosystem Graphs throughout the last year[2]

Overview of the latest graph from Jan. 2024

Node Type	Amount	
Domains	628 M	70.0%
Certificates	171 M	19.1%
IPv4 & IPv6 Addresses	98 M	10.9%

2 starting Jan. 2023

We created 13 monthly Internet-wide TLS Ecosystem Graphs throughout the last year[2]

Overview of the latest graph from Jan. 2024

Node Type	Amount	
Domains	628 M	70.0%
Certificates	171 M	19.1%
IPv4 & IPv6 Addresses	98 M	10.9%

- 90% of edges targeting IP addresses accumulated on only 2% of the nodes

[2] starting Jan. 2023

We created 13 monthly Internet-wide TLS Ecosystem Graphs throughout the last year[2]

Overview of the latest graph from Jan. 2024

Node Type	Amount	
Domains	628 M	70.0%
Certificates	171 M	19.1%
IPv4 & IPv6 Addresses	98 M	10.9%

- 90% of edges targeting IP addresses accumulated on only 2% of the nodes

⇒ we saw a high centralization of the TLS ecosystem, especially for IP addresses

[2] starting Jan. 2023

Results

For each graph and blocklist, we ran the PTP algorithm

Blocklist
abuse.ch Feodo
Blocklist.de Strongips
abuse.ch SSLBL
Openphish

Results

For each graph and blocklist, we ran the PTP algorithm

Blocklist	Type
abuse.ch Feodo	C&C IP addresses
Blocklist.de Strongips	abusive IP addresses
abuse.ch SSLBL	C&C certificates
Openphish	phishing domains

Results

For each graph and blocklist, we ran the PTP algorithm

Blocklist	Type	Observed
abuse.ch Feodo	C&C IP addresses	34
Blocklist.de Strongips	abusive IP addresses	161
abuse.ch SSLBL	C&C certificates	19
Openphish	phishing domains	3 461

Challenge: The Internet is like a black box and we never know if an entity is actually malicious!

Results

Challenge: The Internet is like a black box and we never know if an entity is actually malicious!

- we only have indicators

Challenge: The Internet is like a black box and we never know if an entity is actually malicious!

- we only have indicators
- we can show the value of our approach if the identified domains / IP addresses are largely suspicious

Results

How to evaluate whether we found something suspicious?

1. Manual Inspection
2. Comparison with External Threat Intelligence
3. Analysis Over Time

How to evaluate whether we found something suspicious?

1. **Manual Inspection**
2. Comparison with External Threat Intelligence
3. Analysis Over Time

Manual Inspection

We quickly noticed several clusters of outliers due to their uniform score and large size

Manual Inspection

We quickly noticed several clusters of outliers due to their uniform score and large size

1. 155k domains resolving to a single IP address with a blocked certificate

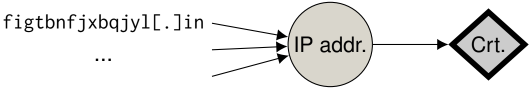

We quickly noticed several clusters of outliers due to their uniform score and large size

1. 155k domains resolving to a single IP address with a blocked certificate

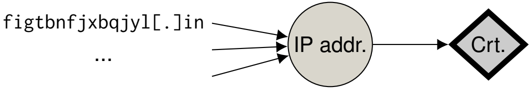

2. 38k unbouncepages subdomains

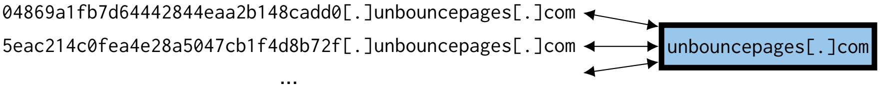

We quickly noticed several clusters of outliers due to their uniform score and large size

1. 155k domains resolving to a single IP address with a blocked certificate

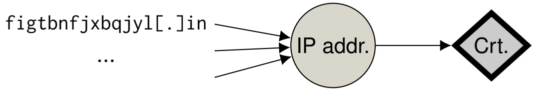

2. 38k unbouncepages subdomains

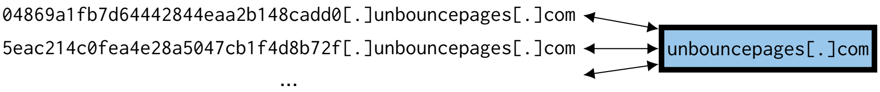

3. 27k sole IP address returning a blocked certificate

Manual Inspection

We quickly noticed several clusters of outliers due to their uniform score and large size

1. 155k domains resolving to a single IP address with a blocked certificate

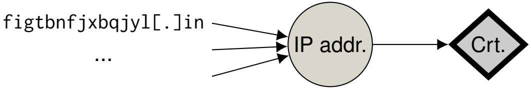

2. 38k unbouncepages subdomains

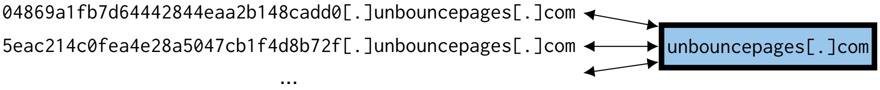

3. 27k sole IP address returning a blocked certificate
4. 3k (seemingly) random domains redirecting to a known phishing domain

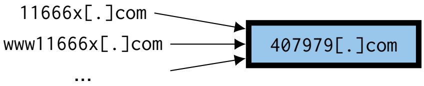

Results

How to evaluate whether we found something suspicious?

1. Manual Inspection
2. **Comparison with External Threat Intelligence**
3. Analysis Over Time

Comparison with External Threat Intelligence

Threat intelligence services:

- Provide API to check a domain or IP address

Comparison with External Threat Intelligence

Threat intelligence services:

- Provide API to check a domain or IP address
- VirusTotal (VT)[3]
  - aggregates a large amount of threat intelligence feeds (e.g., blocklists)

[3] `https://www.virustotal.com`

Threat intelligence services:

- Provide API to check a domain or IP address
- VirusTotal (VT)[3]
  - aggregates a large amount of threat intelligence feeds (e.g., blocklists)
- Google Safe Browsing (GSB)[4]
  - threats information detected by Google

[3] https://www.virustotal.com
[4] https://safebrowsing.google.com

Comparison with External Threat Intelligence

Threat intelligence services:

- Provide API to check a domain or IP address
- VirusTotal (VT)[3]
  - aggregates a large amount of threat intelligence feeds (e.g., blocklists)
- Google Safe Browsing (GSB)[4]
  - threats information detected by Google
- However, both have a very rate-limited API

[3] https://www.virustotal.com

[4] https://safebrowsing.google.com

Domains with a PTP score above the threshold[5] (without the first three manually identified clusters):

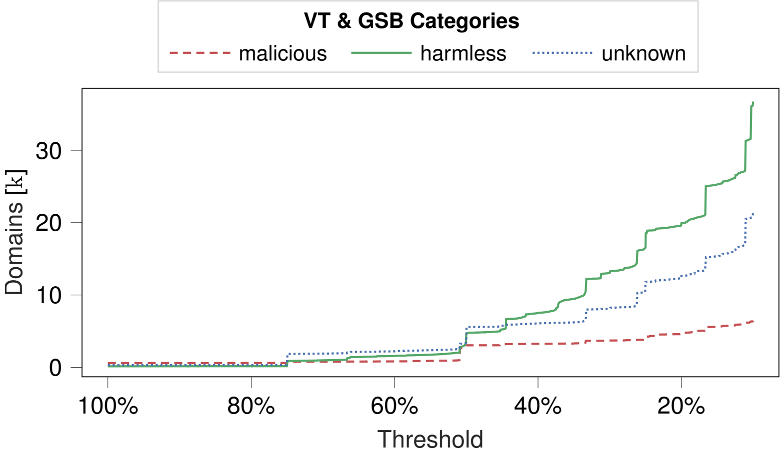

[5] only the latest graph from Jan. 2024

IP Addresses with a PTP score above the threshold[6] (without the first three manually identified clusters):

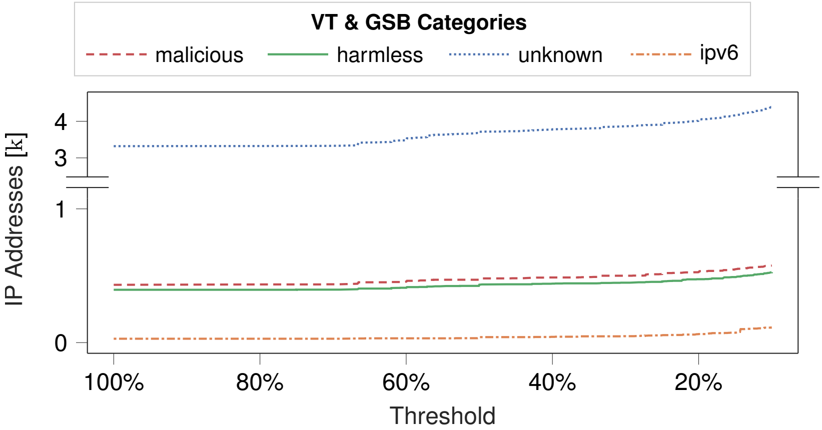

[6] only the latest graph from Jan. 2024

How to evaluate whether we found something suspicious?

1. Manual Inspection
2. Comparison with External Threat Intelligence
3. **Analysis Over Time**

Analysis Over Time

- **Reminder:** We created monthly graphs over the last year

Analysis Over Time

- **Reminder:** We created monthly graphs over the last year
- For each graph, we can evaluate whether nodes with a high score appeared later on the same blocklist

- **Reminder:** We created monthly graphs over the last year
- For each graph, we can evaluate whether nodes with a high score appeared later on the same blocklist
- Calculating the Appearance Rate:

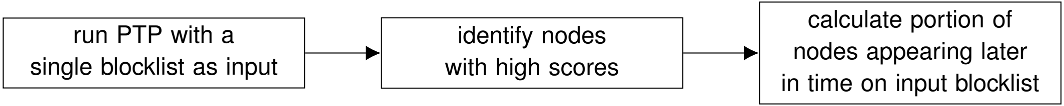

Nodes with a score above an optimized threshold and the portion appearing later on the same blocklist

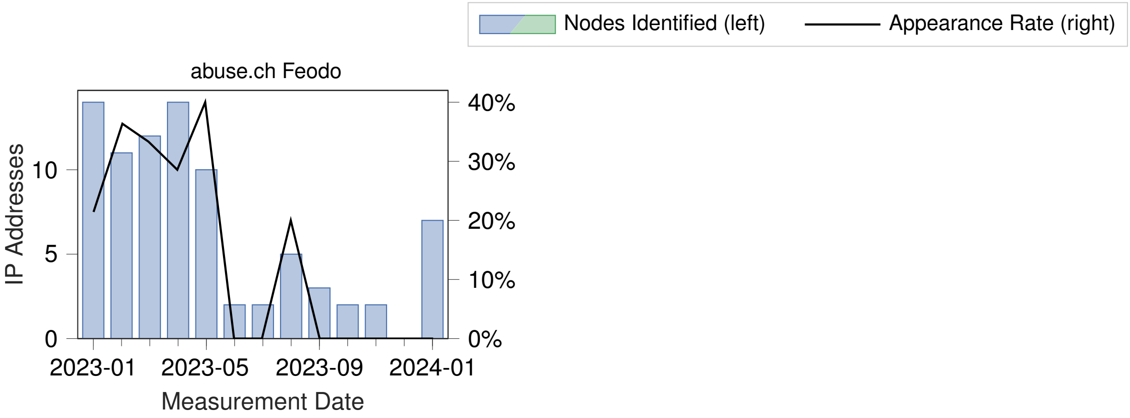

Nodes with a score above an optimized threshold and the portion appearing later on the same blocklist

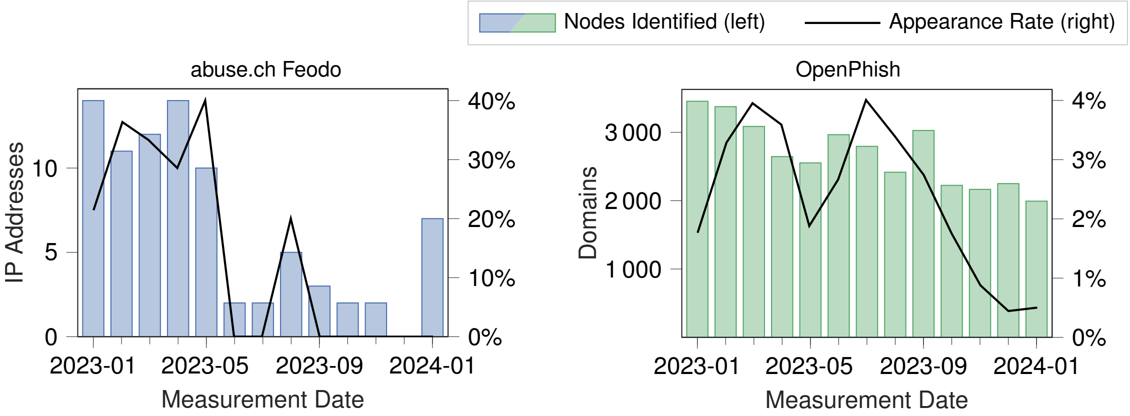

Conclusion

We offer an approach than can navigate the millions of possible domains and IP addresses, to help security researchers focus on suspicious subsets of the Internet when searching for unknown threats.

Read our paper! We provide:

- a versatile TLS ecosystem graph model build around deliberate actions
- a PTP algorithm to propagate threat scores
- three analyses that highlight how our approach focuses on malicious activity
- published results, interactive plots, scripts, and code

`https://tumi8.github.io/iteg/`

Appendix

Example - tma.ifip.org

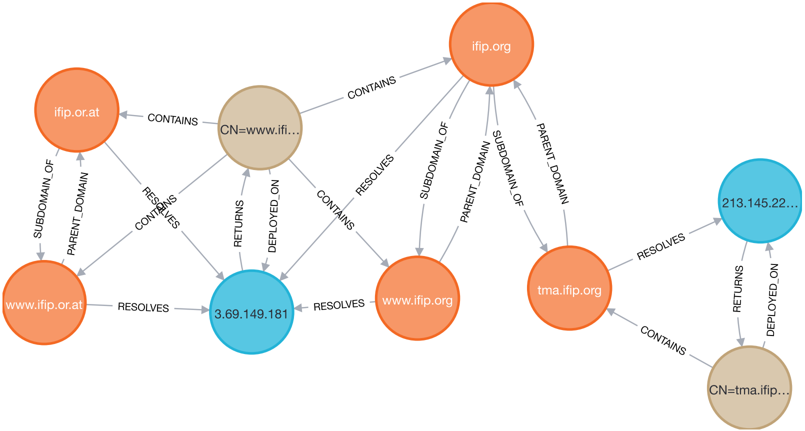

- loading the graph model in Neo4J allows to quickly explore server infrastructure
- did you knew `ifip.org` is also hosted under `ifip.or.at`, although TMA only under `tma.ifip.org`?
- loading the neighbors of `ifip.org` would reveal many more IFIP conferences

Appendix

Example - Early Detection of a Domain

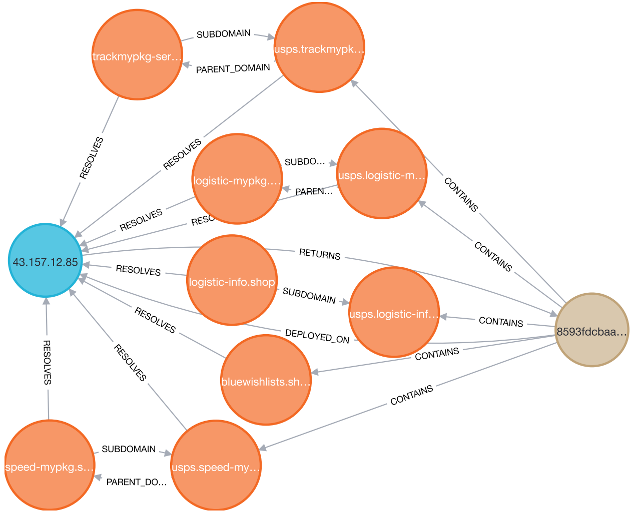

- our graph loaded into Neo4J for easy manual navigation
- only `usps[.]trackmypkg-servi[.]shop`, `usps[.]logistic-mypkg[.]shop`, and `usps[.]speed-mypkg[.]shop` were blocked by OpenPhish
- `bluewishlists[.]shop` appeared later on the blocklist (threat score 67%)
- `usps[.]logistic-info[.]shop` never appeared on the list

Appendix

Optimizing the Detection Threshold

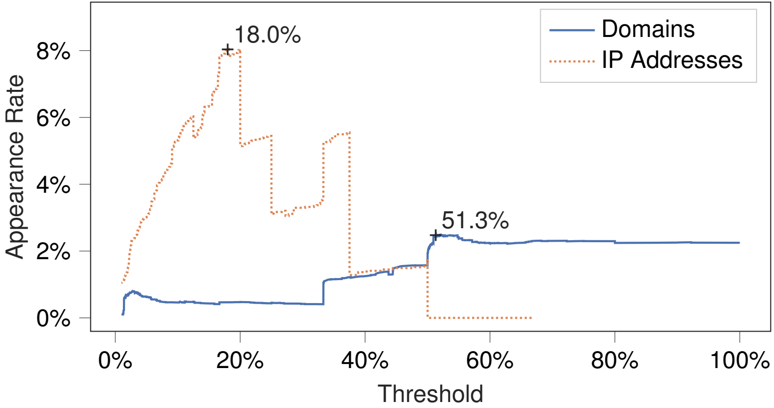

Best performing thresholds:

- Domains: 51%
- IP addresses: 18%

Appendix

Centralization of the TLS Ecosystem