OCTOBER 2022 REPORT

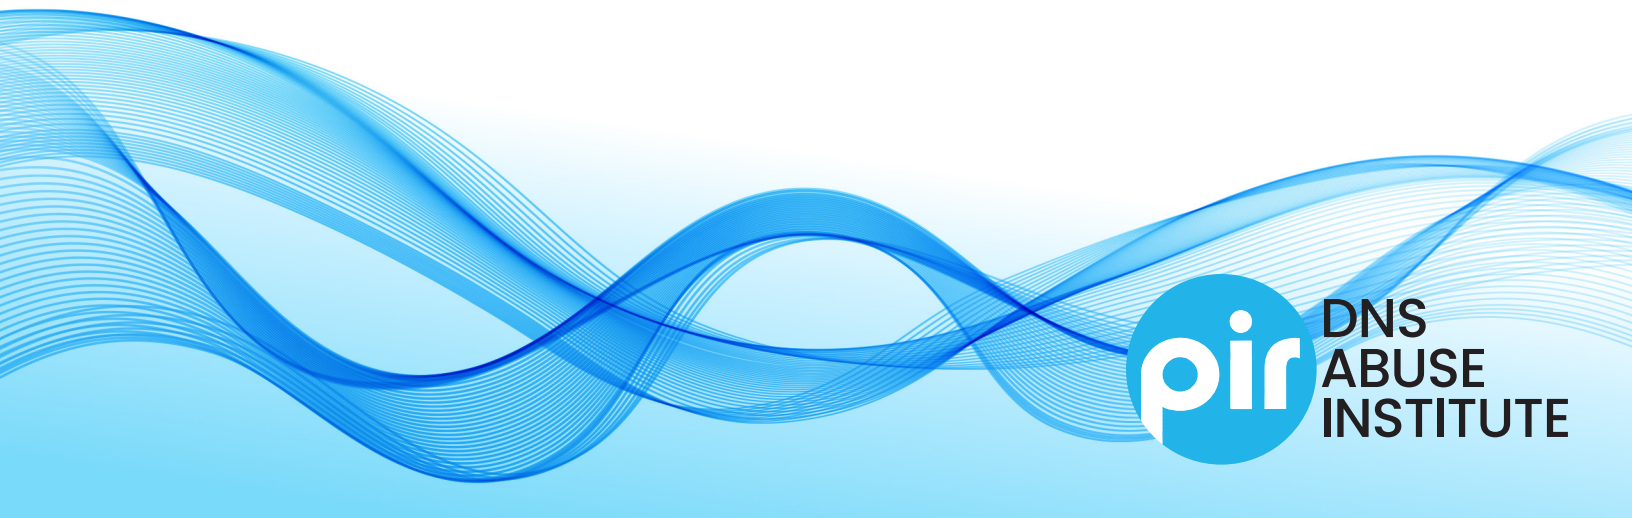

TABLE OF CONTENTS

EXECUTIVE SUMMARY

This report is the second publication from the DNS Abuse Institute's measurement initiative: **DNSAI Compass**, previously known as DNSAI Intelligence.

We've chosen the name DNSAI Compass ("Compass") in reference to the mathematical instrument that consistently measures the distance between two points. This reflects our intention for reliable, consistent, and accurate measurement. We hope that over time Compass also draws arcs between points of interest to improve our understanding of preventing and reducing DNS Abuse. In time, we hope it also guides us towards our own "North"—reducing DNS Abuse—as a navigational compass would.

After our initial report we conducted significant outreach across the DNS community. We focused on registrars and registries, offering them the opportunity to view their data. This process is still ongoing, we have already received considerable feedback from registries and registrars which has proven immensely helpful. We have gathered insight from every meeting and encourage interested registrars and registries to reach out to us to view their data and learn about this initiative. We also identified frequently asked questions which we endeavor to clarify in this report.

Compass is intended to reliably and consistently measure the prevalence and persistence of phishing and malware; it is not intended to capture all harm on the Internet, or to measure the total impact of this harm on end users.

We have prioritized measuring phishing and malware with a consistent, transparent, independent, and academically robust **methodology** as an indicator of wider trends in the DNS ecosystem. We have focused on phishing and malware because they generally provide sufficiently verifiable evidence of the security threat. In time, we may expand our methodology to cover more types of DNS abuse.

We intend to use this data to improve understanding of DNS Abuse, and the opportunities for mitigation and prevention. We are also working towards empowering registries and registrars by providing access to specific reports for their domains under management. Our initial outreach has indicated this would be welcomed.

At this point in our mission, we are particularly interested in the mitigation activity that falls within the control of a registrar or registry. There are a limited number of tools available for registrars and registries, all of which relate to a domain name. This is why our initiative focuses on unique domain names and does not attempt to count URLs; a registrar and registry can only act on the domain name.

This report focuses again on high-level aggregate data. We have added one more month to our previous report, and now cover May through August 2022. We continue to observe a drop in numbers of malware and continue to monitor this. Our detected mitigation activity, mitigation speed, and the composition of compromised and malicious domains is consistent with our previous month of reporting.

Our reports aim to celebrate and recognize good practice, as well as shine a spotlight on potential for areas of improvement in the industry. We are currently considering how to report publicly on individual TLD and registrar performance while recognising nuance and context. We look forward to slowly expanding the granularity of our data with future iterations of public reports.

The **methodology** is the same as our last report (v1.0) and we encourage readers to consider this detailed methodology and contact us with questions, ideas, or suggestions to help us improve this initiative. After all, we are here to support the DNS Community and make it better equipped to tackle DNS Abuse.

The DNS Abuse Institute will periodically publish reports on **DNSAI Compass**.

ABOUT

The **DNS Abuse Institute** (DNSAI or the "Institute") was created in 2021 by **Public Interest Registry** ("PIR") in pursuit of its non-profit mission. The Institute aims to reduce DNS Abuse and empower the DNS community.

The Institute created Compass as a reliable, independent, transparent, and sufficiently granular way of measuring DNS Abuse in order to ultimately reduce it at the DNS level.

Compass is a collaboration with **KOR Labs**, led by **Maciej Korczynski** from Grenoble INP-UGA. This data is then provided to the DNSAI. DNSAI then works with PIR's Data Analytics team to create the interactive charts and for the purposes of writing this report.

Our priorities for Compass are:

- **Transparency**: The methodology that collects, cleans, and aggregates the data must be as transparent as possible. To the extent that anyone should wish to, they could replicate the process.
- **Credibility and Independence**: We aim to have an academically robust and independent approach, separate from commercial interests.
- **Accuracy and Reliability**: The goal of these reports is to enable focused conversations, and to identify opportunities for abuse reduction. The data needs to be of high enough quality to serve as the foundation for meaningful changes to the ecosystem.

Our first report from **September 2022** [1] provides the methodology and further context on the background and development of this initiative.

Our approach is one of collaboration and engagement, and we endeavor to speak to interested parties and provide them with early access to data that concerns their organization. We are committed to refining this project as work continues and welcome insights from across the industry to help us iterate and improve. If you would like to review your data, please **contact us**.

For clarity, Compass exists completely independently of **NetBeacon**, the centralized abuse reporting service we created for the benefit of the DNS. Reports from NetBeacon do not go into our measurement work with Compass. This is a conscious choice to optimize and encourage usage of NetBeacon and prevent any abuse of NetBeacon as an attempt to influence Compass data. See the **methodology** for more information on how domains are included in Compass.

[1] Available at: https://dnsabuseinstitute.org/dnsai-compass/

Understanding this Report

This report is the second publication from the DNS Abuse Institute's measurement initiative: **DNSAI Compass**.

This report shows high level aggregate data from **May through August 2022.**

It focuses on the use of the DNS for phishing and malware:

- **Phishing** is an attempt to trick people into sharing important personal information— banking information, logins, passwords, credit card numbers.
- **Malware** is malicious software designed to compromise a device on which it is installed.

It includes the following charts:

- **Chart 1: Aggregate Trends**
- **Chart 2: Mitigation**
- **Chart 3: Registrar Median Mitigation Time**
- **Chart 4: Malicious vs. Compromised**

Our **methodology** provides important context and we recommend it is read in full.

Each chart is accompanied by:

- '**About this Chart**' to help the reader understand the data being displayed, and;
- '**Commentary**' where we have added any observations on the data.

Where we are showing data over time, the intent is to try and demonstrate trends, year over year, and we are therefore hoping to be able to display about two years of data depending on functionality and viewability.

GENERAL DNS ABUSE TRENDS

These charts are available in an interactive format on our **website**. They provide a broad overview of our findings on DNS Abuse trends.

Chart 1: Aggregate Trends

About this Chart

This chart provides a high level view on how much DNS Abuse has been identified by our methodology, and how it's changing over time. It shows the absolute volume of unique domains our methodology has identified are engaged in phishing and malware, broken out by category.

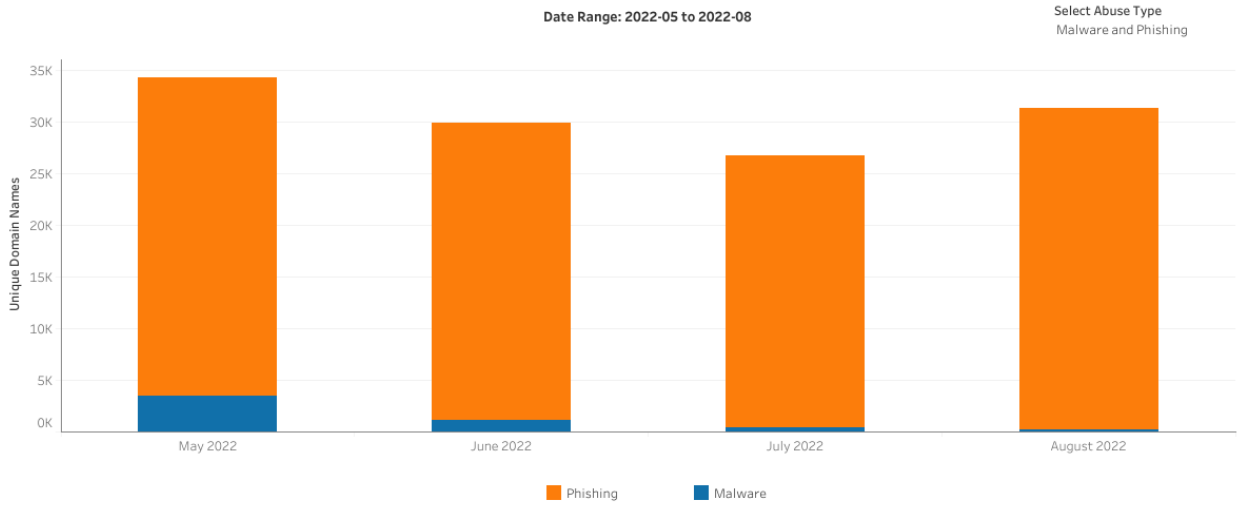

*Figure 1: Aggregate Trends - **Phishing and Malware***

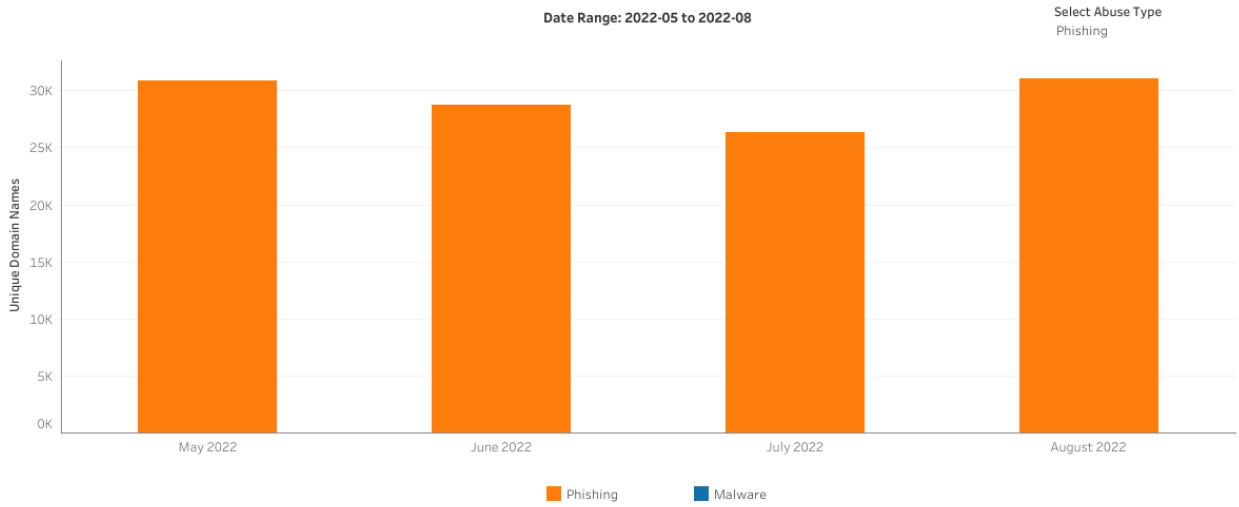

*Figure 2: Aggregate Trends - **Phishing***

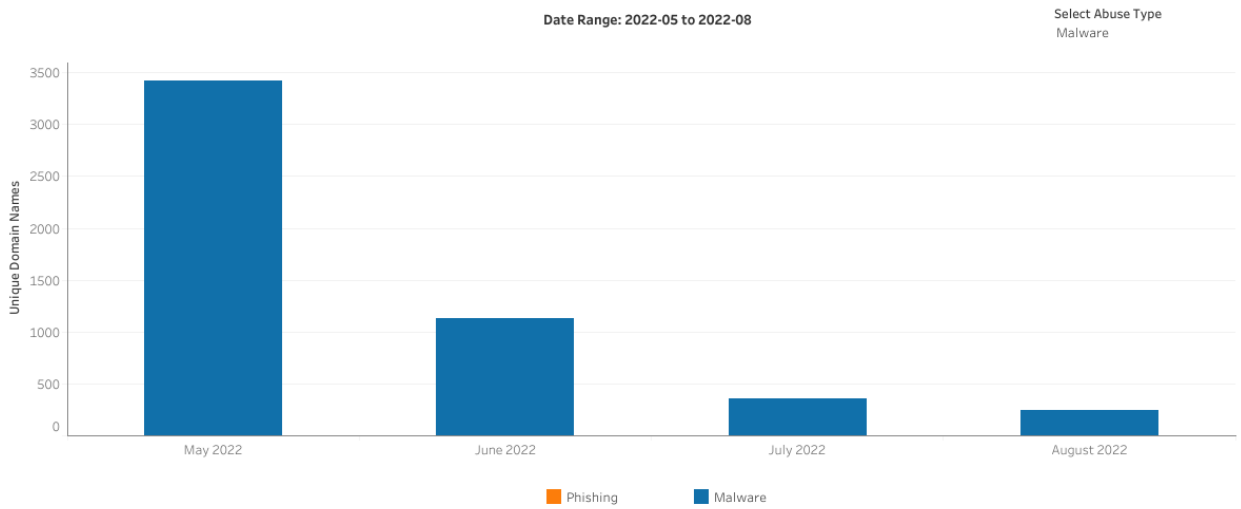

*Figure 3: Aggregate Trends - **Malware***

Commentary

Our data set currently only includes four months, which is insufficient to determine the existence of a sustained trend.

Our methodology continues to identify more occurrences of phishing than malware. This is inline with existing measurement reporting, such as ICANN's DAAR project [2], which also typically reports more phishing than malware. We've noted that this month there is once again a drop in the reported numbers of malware and we'll continue to monitor.

2 https://www.icann.org/octo-ssr/daar

Chart 2: Mitigation

About this Chart

This chart provides a high level view on how much DNS Abuse mitigation has been identified by our methodology, and how it's changing over time. The methodology includes a process to determine whether any mitigation has been observed. This involves taking an initial measurement of various factors related to the URL and repeating these measurements for one month. Further details are set out in the **methodology**.

This results in four labels:

- **Mitigated**: We believe a mitigating action has occurred. This action could be taken by a registrar, registry, a hosting provider, or another relevant actor.
- **Not Mitigated:** We did not detect any indication of mitigation.
- **Uncategorized**: We were unable to determine whether or not mitigation occurred.
- **Unprocessed**: The domains were not processed due to network connectivity, server problems, or other similar issues.

*Figure 4: Mitigation - **Phishing***

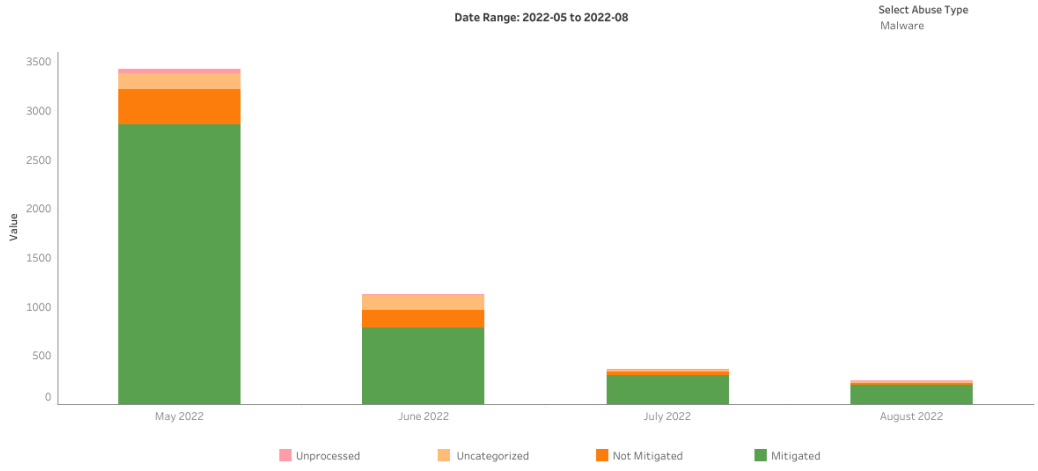

*Figure 5: Mitigation - **Malware***

Commentary

Our data set currently only includes four months, which is insufficient data to determine a sustained trend. As above, we've noted there is a considerable drop in the reported numbers of malware and we'll continue to investigate.

The proportion of domains for which we were unable to categorize is higher for phishing than malware. One possible reason for this is the evasion techniques outlined in the **methodology**.

Missing or limited data is challenging to manage in any data-driven project. In the pursuit of transparency, we have clearly identified the number of domains that we were unable to categorize or unable to process.

For future reports, we are working to balance principles of accuracy and reliability with the desire to compare trends over time. We want this to be a project of iterative improvement, with accuracy increasing over time. However, we also want the ability to compare trends over months and years. We are working on improving the breadth of coverage for categorization of mitigation activity, while avoiding significant changes to the existing categorization methodology for domains that could be categorized. See the **methodology** for further details.

Chart 3: Registrar Median Mitigation Time

About this Chart

This chart is intended to show the observed time taken to mitigate phishing and malware, and how it is changing over time. For the domains that our methodology determined were mitigated, this chart shows how many registrars had a median time to mitigation in each category.

After an initial measurement, KOR Labs repeats measurements for one month to determine if mitigation has occurred. The intervals used are (starting at the time of acquiring the URL from the blocklist): 5m, 15m, 30m, 1hr, 2hr, 3hr, 4hr, 5hr, 6hr, 12hr, 24hr, 36hr, 48hr, and then once every 12 hours for one month.

While we are describing this information as a "median registrar mitigation time", it should be noted that we do not know definitively that it was the registrar that took action. This data could include mitigation taken by the registry, the host, or any other relevant party. The reference to a registrar is indicative that the domain is under their management.

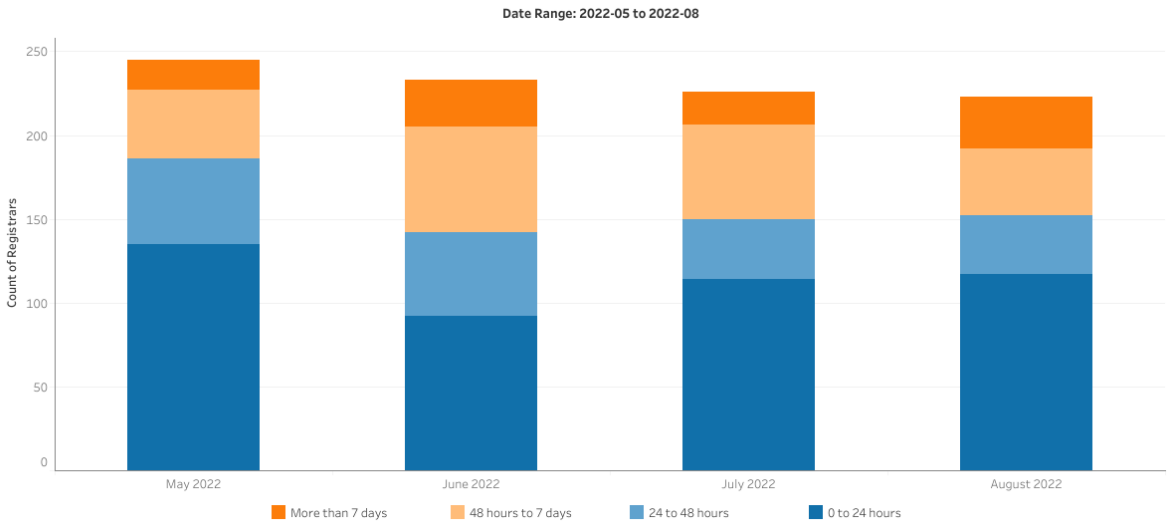

Figure 6: Registrar Median Mitigation Time - ***May, June, July , August 2022***

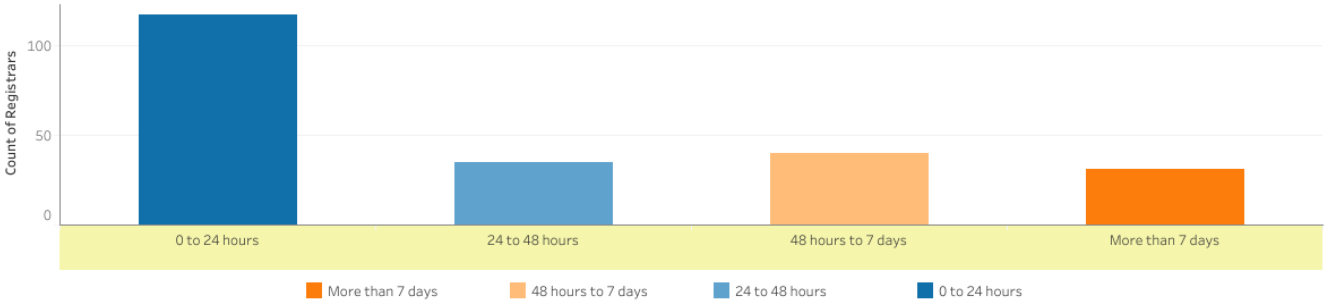

*Figure 7: Registrar Median Mitigation Time - **August 2022***

Commentary

Our data set currently only includes four months, which is insufficient data to determine a sustained trend.

There is no industry standard for how quickly mitigation should occur. This makes the presentation of mitigation time challenging. We believe there is a general industry view that mitigation within 24 hours is considered a quick response to evidence of phishing or malware. As phishing and malware are quite time-sensitive issues, with most harm happening at the start of the attack, we believe that mitigation after 7 days is not quick enough to prevent and disrupt harm, which is why we have included "More than 7 days" as a specific category.

Chart 4: Malicious vs. Compromised

About this Chart

This chart is intended to show the observed registration type (malicious vs. compromised) and how this is changing over time.

Our methodology includes three labels:

- **Malicious**: a domain registered for malicious purposes (i.e., to carry out DNS Abuse).
- **Compromised**: A benign domain name that has been compromised at the website, hosting, or DNS level.
- **Uncategorized**: A domain that our methodology was unable to categorize for a number of reasons, including problems in collecting the metadata necessary to categorize domain names accurately.

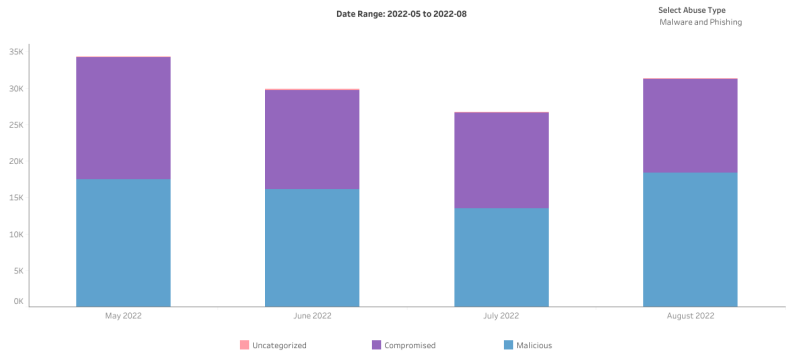

*Figure 8: Compromised vs Malicious - **Phishing and Malware***

*Figure 9: Compromised vs Malicious - **Phishing***

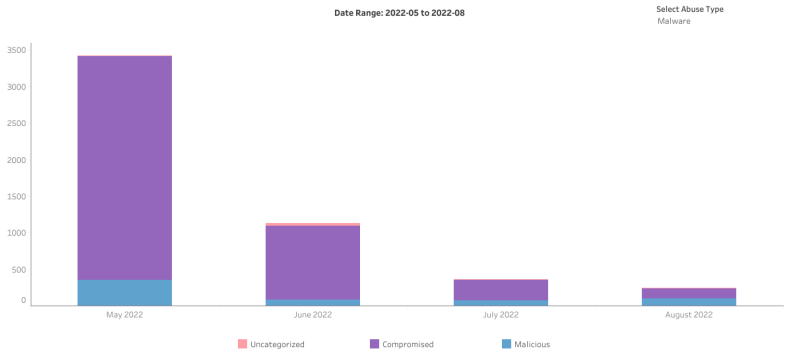

*Figure 10: Compromised vs Malicious - **Malware***

Commentary

Our data set currently only includes four months, which is insufficient data to determine a sustained trend.

The distribution between domains identified as malicious or compromised is different for phishing and malware over this initial period. The data shows more domains identified as maliciously registered for phishing. For malware, more domains were identified as compromised.

DNSAI COMPASS

DNS ABUSE INSTITUTE

www.dnsabuseinstitute.org