Domijn: The Security of Domain Registrars and the Risk of a Domain Name Takeover

Koen van Hove
NLnet Labs & University of Twente
koen@nlnetlabs.nl

Jeroen van der Ham-de Vos
University of Twente
j.vanderham@utwente.nl

Roland van Rijswijk-Deij
University of Twente
r.m.vanrijswijk@utwente.nl

Abstract

Domain names are key assets for organisation. They anchor an organisation's online presence and reputation, and serve as linking pin for web services and, e.g., email. Consequently, a malicious takeover of a domain can lead to significant damages. Organisations register domain names through so-called registrars, a type of business that plays a key role in the domain name industry. This implies that registrars play an important part in safeguarding against malicious takeovers of domains. In this paper we empirically study how registrars implement security controls to prevent against such takeovers. We focus on the top 10 most popular registrars for the .nl ccTLD. We present the results of this study in light of a model for the impact of domain takeovers, that analyses the possible consequence of a takeover. We contrast this against the impact of two other well-known threats: ransomware and DDoS attacks. We find that all registrars in our study implement relatively effective security measures, but that they fall short in more advanced security controls, such as the proper implementation of two-factor authentication. We also find that a domain takeover can have significant impact, potentially equalling that of a ransomware attack.

1 Introduction

A domain name is an important asset for organisations. It is where their website is, what is used to receive email, where many of their internal services live, and more. Generally an organisation registers their domain name a so-called *registrar*. A registrar manages the registration at the registry and may also provide additional services, such as managed DNS hosting to also provide the technical services needed to operate an Internet domain. While registrars are the primary actor responsible for the administrative handling of domain name registrations, the actual business model is sometimes more complicated. Larger registrars often allow their services to be offered by resellers. This allows, e.g., hosting companies that themselves do not have the proper accreditation to also offer domain name registrations to their customers.

Most registrars and resellers offer their services through a web portal where information can be managed, such as: 1. Ownership information; 2. Data that ends up in public databases containing information about the domain name and the holder; 3. Domain Name System (DNS) nameservers[1], which tell those resolving the domain name where to find the records for that domain; 4. Transfer codes, a code allowing someone to authorise moving a domain name to another registrar; 5. DNSSEC keys[1], allowing someone to set the keys used to cryptographically sign the records in the DNS for this domain.

This setup means that a registrar's web portal is a crucial element in protecting the security of a domain name registration. If a malicious party manages to log in to the web portal on behalf of someone else, they can likely take over the domain name.

Whilst we did notice several large registrars explicitly focus on brand protection and the perceived risk of a domain takeover, we noticed a lack of scientific models describing the impact of such a takeover, as well as a lack of scientific work on the technical likelihood of such a takeover. In this paper we aim to fill that gap by modelling the risk of a domain takeover for an organisation, and assessing it on using the NIST Risk Assessment Scale [17] as shown in Figure 1. We do so by first analysing the likelihood of a malicious party being able to take over a domain name by systemically investigating the technical security of registrars and resellers. Then we analyse the impact of a domain takeover by taking inspiration from existing models for two other, well-studied attacks for ransomware and DDoS. We then combine the likelihood and impact to assess the general risk.

We focus on .nl as The Netherlands is considered a role-model according to the International Telecommunication Union's (ITU) Global Cybersecurity Index 2024 report [18], scoring 20 out of 20 in nearly every category including technical measures. The most popular registrars and resellers used for .nl are Dutch. This means our results likely paint a 'best

[1] For an in-depth discussion of the technical aspects of the Domain Name System, we refer to a tutorial by Van der Toorn et al. [37].

Likelihood of Threat Event Initiation or Occurrence	Likelihood Threat Events Result in Adverse Impacts				
	Very Low	Low	Moderate	High	Very High
Very High	Low	Moderate	High	Very High	Very High
High	Low	Moderate	Moderate	High	Very High
Moderate	Low	Low	Moderate	Moderate	High
Low	Very Low	Low	Low	Moderate	Moderate
Very Low	Very Low	Very Low	Low	Low	Low

Figure 1: The National Institute of Standards and Technology (NIST) risk assessment table

case' scenario and paint a too optimistic picture rather than a too pessimistic picture.

The contributions of this paper are:

1. An empirical study into the security measures of domain registrars and resellers, and a characterisation of the shortcomings in the security of domain registrars and resellers, with as key finding the improper implementation of two-factor authentication;

2. An analysis for the impact of a domain name takeover for an organisation;

3. A risk assessment of a domain name takeover for organisations based on this empirical study and impact model.

Outline – the structure of this paper is as follows: First we describe the background and technical terminology in Section 2. After that, in Section 3 we describe the related work. In Section 4 we describe the methodology for which registrars and resellers we test the security of, and how we test them. In Section 5 we show the results of our tests and the vulnerabilities we found. In Section 6 we analyse the impact of a domain takeover for those organisations and make the risk assessment. In Section 7 we analyse how applicable these results and impact are outside the Netherlands. In Section 8 we describe how we disclosed the vulnerabilities we found to the registrars and resellers. In Section 9 we discuss what these results mean in broader context and the limitations. In Section 10 we make recommendations on how to make hostile domain name takeovers less likely. Lastly, in Section 11 we draw our conclusions.

Our work does not exist in a vacuum outside of society. For that reason, in Section 4.6 we describe the ethical considerations and steps we have taken to minimise adverse effects.

2 Background

The Domain Name System (DNS) is a system to make the Internet accessible to human beings [16]. Rather than having to remember a string of numbers (such as 203.0.113.5) called an Internet Protocol (IP) address to access something on the Internet, the DNS uses human readable domain names (such as example.com), and allows for quick lookup of the associated IP address, much like a phone book can be used to look up phone numbers for an organisation or name. Apart from IP addresses, the DNS is also used to relay where email for that domain name should be delivered, among other things.

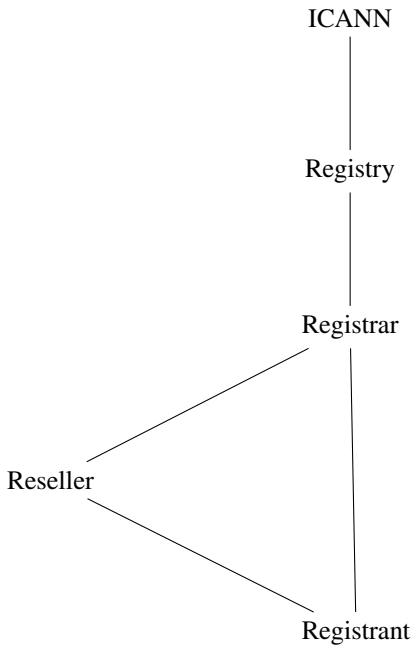

Figure 2: The Registry-Registrar-Reseller-Registrant relationship where ICANN is the root

The DNS is hierarchical, as shown in Figure 2. At the top is the Internet Corporation for Assigned Names and Numbers (ICANN), they coordinate the DNS. ICANN delegate the management of the top-level domains (TLDs) to the registries – examples of TLDs are '.nl', '.com', and '.edu'. Registries manage the list of all domain names under their TLD. Registries delegate the task of selling domain names to one or more registrars. They are the organisations where one would go to register, .e.g., 'example.org'. The person or organisation who registers such a domain name is called the registrant – they are the one who hold 'example.org'. Sometimes there is another layer between the registrar and registrant called a

reseller. From the perspective of the registrant, there is no difference between a registrar and reseller – both enable you to register your domain name. The only difference is that they do not have a direct contract with the registry, but with a registrar.

In order to find the IP address for a domain name, e.g. my.example.org, your computer does the following:

1. Your computer asks the root (ICANN) where to find .org, which tells you where to find .org;
2. Your computer asks .org where to find example.org, which tells where to find example.org;
3. Lastly your computer asks example.org where to find my.example.org, which then tells you where to find my.example.org.

In the original version of the DNS, none of the answers were cryptographically signed. This meant that, much like the telephone game, any adversary-in-the-middle could change any of the answers. DNSSEC was added later to make answers cryptographically verifiable. The registrar plays a role in this, as it has to send the key material for a domain name to the registry [5], so that when one asks the .org registry where to find 'example.org', the registry can answer both where to find example.org, but also which key answers from example.org (e.g. for my.example.org) will be signed with.

2.1 Terminology

We believe it is worthwhile to provide a brief description of the technical terms used throughout this paper:

TLD A Top-Level Domain (TLD) is a domain in the DNS at the highest hierarchical level. Examples are .com or .nl.

Domain name The part of a network address that defines a realm of control under a TLD, such as example.com or voorbeeld.nl.

DNS The Domain Name System (DNS) as defined in RFC 1035 [22] and subsequent standards are the set of protocols defining the technical workings of domain names.

Resolver A resolver retrieves information (records) associated with the domain name, such as the IP address it can find a server, or which server email should be sent to.

Nameserver A nameserver contains the information about the records for a domain name, such as the IP address it can find a server, or which server email should be sent to, and provides that information when a resolver asks for it.

DNSSEC DNS Security Extensions (DNSSEC) [13] are a set of protocols that provide origin authentication of DNS data, allowing the resolver to check DNS data has not been tampered with.

Registry A registry is the operator of a TLD, who manages all the database of all domain registrations within that TLD. [14]

Registrar A registrar is an organisation accredited by a registry to sell domain registration services to the (general) public. They may, but need not, provide more services.

Reseller A reseller is an organisation that sells domain registration services to the (general) public via a registrar.

Registrant A registrant is an organisation or individual that buys and holds the domain name.

WHOIS WHOIS (not an acronym, pronounced as 'Who is') is a protocol with which information about the registrant can be retrieved about an internet resource such as a domain name [8]. Information can include administrative and technical names, addresses, phone numbers and email addresses.

RDAP The Remote Data Access Protocol (RDAP) is the follow-up protocol from WHOIS [26]. It is used for the same purpose, but retrieves it in a standardised format.

DDoS attack A Distributed Denial of Service (DDoS) attack is a type of cyber attack where many systems try to flood a service, making it unavailable. [30]

Ransomware Ransomware is malware that encrypts a users' files, effectively holding data and systems hostage until an amount of money (the ransom) is paid to the parties.

3 Related Work

Due to the nature of this paper, we split our related work up in three parts:

1. The first part is about the related academic work regarding the technical aspect of domain names and possible ways to abuse aspects of its operation;
2. The second part is about the related academic work on impact models for cyber attacks;
3. The third part is about instances of domain name takeovers we found in news articles.

3.1 DNS and Registrar Attacks

DNS and DNS vulnerabilities have been studied in great detail, for a good overview we refer to the survey by Schmid [32].

DNSSEC, its working, security, and deployment, have been studied greatly. We refer to the work by Chung et al. [5] and Lian et al. [20] for an overview. Registrars hold an important position in DNSSEC, as they are the ones who can add, change, and remove the DNSSEC keys used for a domain name. Hence

we focus on the role of the registrar. The UK National Cyber Security Centre published guidance for good security practice for domain registrars [25].

When it comes to the registrar's role, Chung et al. [5] studied the role registrar's play in DNSSEC. They find a worrying lack in security posture, e.g., demonstrating that some providers accept changes to key material via email without checking if these mails come from the legitimate domain name owner. Akiwate et al. [2] explored a renaming scheme some registrars use to handle expired domain names where those domain names are still referred to by another domain name, making those other domains susceptible to abuse.

Vissers et al. [38] do something similar by looking at what they refer to as 'nameserver typosquatting'. The DNS allows for domain names to refer to other domain names, similar to how a dictionary can refer to another word for the definition of a word. Unlike a dictionary however, the DNS allows referring to more than one other domain name to allow for redundancy. Vissers et al. analyse where typos have been made in one of these names, such that, if one were to register that domain, they could reply with malicious answers.

Zhang et al. [41] do something similar with stale 'glue' records. Glue records are necessary when the hierarchy is self-referential. For example, .org might answer that the data for example.org can be found at ns1.example.org. We cannot ask example.org for the data for ns1.example.org as we do not know where example.org can be found yet. In these cases, a 'glue' record allows .org to also inform us where we can find ns1.example.org directly.

Schlamp et al. [31] look at domain names that have expired that are listed in internet resource databases, which shows the holder information for those resources. By registering those domain names, one can claim ownership and take over those resources.

We could find no research that systematically looks into the security measures taken by domain registrars, although Akiwate et al. [2] do show that certain registrars are disproportionately associated with malicious domains, and Munny et al. [24] show that certain providers are used more commonly for toll scams.

3.2 Impact Models

In order to understand the risk of a domain name takeover, we need both the likelihood and the impact of such an event.

There are impact models and quantification for other kinds of cyber attacks. Khan et al. [19] design a generalised model for quantifying cyber security. The National Institute of Standards and Technology (NIST) have also created a framework for conducting cyber risk assessments [17]. Wolthuis et al. [39] show a model in action by using a DDoS attack as an example.

Risk has also been quantified from a financial perspective. Razavi et al. [30] analyse the financial loss for DDoS attacks at banks, Woods et al. [40] look at the impact of a company's stock value, and Gomez et al. [12] who look at how much money ransomware groups earned by analysing Bitcoin transactions.

We could not find a threat model if a registrar or reseller is breached and a domain name is taken over, although it is named by Schmid [32] and in a report by ICANN [15].

3.3 In the News

There are known instances of domain name takeovers. Davis [10] describes a breach at a registrar using social engineering, which allowed the attacker access to the organisation's entire cloud infrastructure. Cointelegraph [7] describes another breach where domain hijacking was used to steal cryptocurrency. The Register describes a breach that impacted Twitter.co.uk and the New York Times [4]. The Cybersecurity and Infrastructure Security Agency (CISA) mention that there has been an increase in malicious activity in DNS infrastructure [6], and does mention verifying DNS registrar accounts, which seems to be related to a report from Cisco [29] from around the same time, where email traffic for the Lebanese and United Arabic Emirates (UAE) was redirected.

In this paper we thus aim to fill the gap by investigating the state of registrar and reseller security, and analyse the impact by creating an impact model inspired by already quantified and modelled threats. This way, we can assess the risk of a domain name takeover.

4 Methodology

As described in Section 2, for the registrant there is little difference between a registrar and reseller. In most cases the registrant will log into the web portal provided by where they bought their domain name, unaware of whether they are a registrar or reseller. A change made by a registrant at the reseller will still end up at the registry in the same way as it would when making that change at a registrar. For that reason, **we group these registrars and resellers and refer to them as 'agents'**.

4.1 Determining Agents

We use the top 1,000,000 domains from Cloudflare Radar between 22-29 July 2024, which they base on data from their public DNS resolver *1.1.1.1*. We filter out the .nl domains – this results in just over 9000 .nl domains. We elaborate why we focus on .nl in Section 7. Based on this we aggregated the most common agents by querying for their registrar and reseller data over RDAP. This can be seen in figure 3.

Some agents allow for sign-ups without human interaction, mainly through a web portal. *A* has two distinct registrar/reseller names they use in the WHOIS data, but both use

the same portal. *KPN Zakelijk* has a sign-up process with specific requirements we did not meet.

Not all registrars that only act as portals for resellers (i.e. do not do registrations from customers directly) always accurately include the reseller information in the WHOIS data (such as Registrar.eu, Key-Systems GmbH, and The Registrar Company B.V.), and these have been excluded.

The agent landscape in The Netherlands (where most .nl domains are used) is varied. Over 300 agents are required to reach over 90% market share. In total we noticed over 1000 different agents within these 9000 domain names. We believe this selection provides an accurate view of agents that are commonly being used for .nl domain names. Based on the names from the RDAP data, most agents appear to be Dutch.

We look at the top 10 agents from this set that allow for 'automated' sign-up, meaning by entering our details on a website and spending a small fee. The others, such as *CSC Global*, require an onboarding process through a sales representative, which would require us to have a business at scale, or are closed for signing up entirely, such as *SURF*, which only provides services to its members.

Even though the national government ('Rijksoverheid') is its own registrar, some organisations within the government use a different registrar. When checking the WHOIS data of the list of national government domains, all agents listed in Figure 7 appeared at least once (though often several times).

Agent	Market share	Automated	Country
*A**	20.31%	✓	NL
B	3.05%	✓	NL
C	2.47%	✓	FR
D	2.37%	✓	NL
*A**	1.98%	✓	NL
E	1.93%	✓	NL
F	1.92%	✓	NL
CSC Global	1.70%	✗	US
Rijksoverheid	1.53%	✗	NL
KPN Zakelijk	1.23%	✓**	NL
SURF	1.14%	✗	NL
G	1.14%	✓	NL
H	1.07%	✓	DE
I	1.07%	✓	NL
MarkMonitor	0.91%	✗	US
J	0.87%	✓	NL

Figure 3: The top .nl registrars and resellers based on the Cloudflare Radar data from 22-29 July 2024 based on data from their public DNS resolver 1.1.1.1.

4.2 Registering Domains

We determined which agents to look at. As we also highlight in Section 4.6, we only use our own domains for testing. For that reason, we have registered the domains as a natural person at each agent. We did so using a real, albeit not colloquially used, name, and an address on the University of Twente. We used a newly created AOL email address, and a new phone number specifically for this purpose.

Figure 4: A screenshot of the website we created for the fictional writer. The domain name has been redacted

4.3 Authentication by Agent

For each agent we registered a domain with, we kept track of the information they asked of us during registration, as this can later be used to authenticate us. This can be seen in Figure 7. In all but one case this was only the name, email address, telephone number, and address. We also kept track of which two-factor authentication options were provided, and whether it was mandatory to have one. Timed One-Time Passwords (TOTP) seem to be far the most popular method of multi-factor authentication.

4.4 Information Available

We created a website (as seen in Figure 4, name changes per domain) for a fictional book writer that lists that it is still under construction. The website contains some text, and the author's email address specific to that domain, telephone number, and physical address.

Out of the 10 agents we registered a domain with, 8 show our email address we registered with as administrative and/or technical contact in the public WHOIS data. Only *C* and *H* do not show the address, but instead use their own address. However, this is likely to change in the future due to new policy from Stichting Internet Domeinregistratie Nederland (SIDN) since October 2023 [34]. Most other data is redacted for privacy reasons. *D*, *J*, *I*, *G*, and *F* use the email address that was used for sign-up as login name and email address in the WHOIS data by default.

The data the registrar has about the registrant is listed in Figure 7. We assume one is able to find out the following information available about the holder of the domain (e.g. through open source intelligence):

1. The domain name;
2. The name of the holder;
3. The email address associated with the account;

4.5 Experiment Setup

We have determined which agents to look at, what data they possess about the registrant, and what data an outsider might reasonably have about the registrant. Based on this, we systematically test, for each of the agents:

1. How likely it is that we are able to find a working password for a specific user;
2. What technical measures they use to protect against password / TOTP brute forcing, by testing 100 TOTP codes in quick succession;
3. What verification questions are asked when trying to get the email address associated with an account changed by claiming the password reset email does not arrive via a phone call;
4. Whether based on the information an outsider has available, they are able to make unauthorized modifications to a domain name they do not hold.

4.6 Ethical Considerations

From the start of this research, we involved the ethics commission of the University of Twente. We always made sure to only do testing using our own resources and own domains. The registrars and resellers where issues were found were informed before publication as described in Section 8. In all our testing we made sure to minimise the impact on existing customers of the registrars and resellers. The registry, Stichting Internet Domeinregistratie Nederland (SIDN) was informed beforehand.

We made very sure to not put the blame or responsibility on a specific person or to guilt-trip or pressure them to do something for us. Their names and identifying information were not recorded.

We made the conscious decision to not publish the names of the registrars and resellers that we interacted with. We are aware that their names can be deduced from the data, and of the risk it poses to these organisations and their customers, but we believe it is infeasible to do this research in a way that cannot be traced back to the registrars and resellers. We believe that there is a greater public benefit by publishing this paper than by withholding it.

5 Results

5.1 Password Availability

An estimated 43% to 51% of users reuse their passwords [9]. According to Ablon et al., 43% received a breach notification in their lifetime [1]. Research at Google discovered that out of password stolen or leaked online, between 7% and 25% matched the Google account's login data [35]. Using these numbers we can obtain a rough estimation for the number of passwords being available.

Out of the 9000 .nl domains, we identify just over 7500 unique email addresses in the WHOIS data. We checked for these email addresses whether they had any data in Have I Been Pwned [36]. We found breach notifications for roughly 54% of those email addresses. Of those accounts with breach notifications, 56% had a breach that involved the leaking of passwords. These numbers also include 'privacy' email addresses as provided by, for example *C*, which are unique and unlikely to be used anywhere else. It is noticeable that a lot of these domains use an administrative and/or technical contact that uses a private email address. For example, around 5% uses an @gmail.com address.

Combining these statistics, we believe that between 2% and 8% of domains currently have a working password available online from a leak. We have not tried to confirm this number.

5.2 TOTP Bypass

We can approach the time it takes to brute force the TOTP token using the binomial distribution. We assume the current TOTP code, the previous code, and the next code are valid at any point in time. We fire off around 100 requests per second. By running this for an hour, our probability to correctly guess the TOTP code becomes:

$$1 - B(100 \cdot 3600, \frac{3}{1,000,000}) \approx 55\%$$

We therefore check whether we can burst 100 requests in one second – if the agent does not prevent our attempts, we could likely brute force the TOTP. We do this by logging in as usual, sending off 100 TOTP requests, and then trying to log in using the correct TOTP code on the original page. If the

Agent	BFP	BFP method	Reset
A	■	Lifetime expires	Text message
B	□		Contact support
C	□		Recovery codes
D	◩	IP rate limiting	Recovery code
E	◩	IP rate limiting	Contact support
F	□		Contact support
G	□		Contact support
H	□		Contact support
I	■	Session expires	Contact support
J	◩	IP rate limiting	Contact support

Figure 5: The agents, their TOTP brute force protections (BFP) – if applicable – and reset mechanism in case the TOTP token was lost. ■ means the agent implemented rate limiting on a per-account basis, ◩ means rate limiting was implemented, but not per account but per IP address, and □ means no rate limiting was implemented.

correct TOTP code still works, it gives a strong indication no rate-limiting is being applied. If the correct TOTP code does not work, we try changing our IP address and entering the code again. If that does work, it gives us a strong indication that rate limiting is only applied per IP address, and not per account. We only burst 100 requests to limit the possible damage on the agent's servers – also see Section 4.6. Our findings as tested in the end of 2024 and early 2025 are shown in Figure 5.

RFC 4226 section 7.3 recommends several throttling techniques for one-time passwords [23], but it states they have to be applied across login sessions. Only *A* and *I* did this on a per-account basis. *D* did lock out our IP address for half an hour but not our account. *E* returned a "429 Too Many Requests" page. In both cases changing our IP address and reloading the page allowed us to sign in. *J* did not tell us that we were rate limited, but refused our correct TOTP token with a 429 error. However, after changing our IP address we were also let in.

To our knowledge organisations such as the National Institute of Standards and Technology (NIST) and the National Cyber Security Centre (NCSC) provide no guidelines for this. The Open Worldwide Application Security Project (OWASP) does not mention rate limiting, but does mention alerting the user at a failed two-factor authentication login attempt [27]. None of the agents we tested did this.

5.3 Helpdesk Calls

Phone calls involve another person directly. Hence for that reason we have taken the utmost care to not harm any of the operators in the process. See Section 4.6 for more details.

We created a phone script to mimic someone calling to reset their password, and not receiving the reset email. The information available to the person calling is:

1. the domain name of the registrant;
2. the name of the registrant;
3. the email address of the registrant;

This is in line with the data we considered to be available described in Section 4.4. The caller then tries to either receive the (new) password via phone or get the email send to their own email address, thereby gaining access to the account. We created a rough call flow diagram as shown in Appendix A – we consider it infeasible to account for every possible answer.

Our goal is not to come up with the most persuasive social engineering scenario, but rather determine what information is required to gain access to an account. Our hypothesis is that with fairly limited data about a domain name, access can be acquired, similar to what Chung et al. [5] found. The information available to the operator on the other end about us is limited. The information we provided, as can also be seen in Figure 7, tends to be not particularly secret. Additional information they have consists of things such as receipts and invoices, or asking for ID.

We do want to explicitly point out that simple measures exist to prevent this kind of impersonation. Calling the phone number on file, sending an email to the known address, or sending new login details per letter to the registered address, are all first-line defence options to prevent account hijacking. In all of these cases, the legitimate account holder would notice the attempt.

A does not have a phone number. Their system is outlined in their knowledge base, and involves a copy of an identity document and copy of payment, as well as a business register extract. *C*, *J*, and *D* also state they ask for proof of ID via email. We have thus excluded these from this test. We hid the phone number we called from.

B and *F* share a customer service team – we only called *B*. *E* does not do phone support nor does their website list how to do it, the only option is to open a support ticket.

We list the information that was asked of us in Figure 6. None of our attempts resulted in gaining access to the account. We were not notified of any attempt by any of the registrars either.

- ***B*** asked for the name, domain, customer number (which we claimed to not have), and then asked us to send a copy of our ID card in a response to their email, as well as the new email address;
- ***G*** asked us to send a redacted copy of our ID card (with only the name visible) to their general support email address;
- ***H*** could not do anything without a customer number;
- ***I*** asked us to call back from the phone number known in their systems.

These results are in contrast with the results from Chung et al. in 2017 [5] where a lot could be arranged via a phone call.

Agent	Phone	Information requested				
		Name	*Domain*	*Cust. No.*	*Ident.*	*Tel. no.*
A	✗	Website asks to send identification				
B	✓	✓	✓	✓	✓	✗
C	✗	Website asks to send identification				
D	✓	Website asks to send identification				
E	✗	Unknown				
F	✓	Same as *B*				
G	✓	✓	✓	✓	✓	✗
H	✓	✓	✗	✓	✗	✗
I	✓	✓	✓	✓	✗	✓
J	✗	Website asks to send identification				

Figure 6: The information requested over the phone when we called the customer service of our ten agents.

5.4 Analysis of the Most Popular Domains

In Section 4 we describe that we look at the top 1,000,000 domains from Cloudflare Radar. We also specifically look at the top 10% of those domain names (i.e. the top 100,000) – of which there are 480 .nl domain names.

As can be seen in Figure 3, not all registrars make it possible to register a domain name through a web interface without human interaction. The most prominent ones are *CSC Global*, *Rijksoverheid*, *KPN Zakelijk*, *SURF*, and *MarkMonitor*. These five can again be split into two:

Private registrars – *Rijksoverheid* is the registrar for the Dutch national government and *SURF* is the registrar for Dutch education and research institutions. Access is only for their 'members' (i.e. institutions and departments within their scope) – they are thus fully closed off;

Corporate registrars – *CSC Global*, *KPN Zakelijk* and *MarkMonitor* are companies whose main service is "protecting your brand". They promote their services by claiming that they protect your digital assets (of which domain names are part), and that there is personalised around the clock support. They allow anyone to sign up (given they bring enough money), but only through human interaction (e.g. a sales representative).

When we analyse the 480 .nl domains names in the top 100,000 domain names from Cloudflare Radar between 22-29 July 2024 we see a shift towards these private and corporate registrars compared to Figure 3.

We considered every domain of these 480 by hand, and find that many of the domain names used as 'brand' by a major institution (i.e. the main domain an organisation is known for) are registered at either a private or corporate registrar. Out of those that are not, most are registered at *A*. The domains from main institutions that are registered at neither *A* nor a private/corporate registrar are the exception.

6 Impact of a Domain Takeover

In the previous sections we looked at the likelihood of being able to gain access to an agent's web portal. The question remains what one can actually do once access has been gained into the web portal of an agent, and what the impact of that is.

Stichting Internet Domeinregistratie Nederland (SIDN) introduced .nl Control [33]. With .nl Control, SIDN will call to verify the assigned person at SIDN to verify the changes, and later also require a signature before any changes are processed. This greatly limits the changes anyone can make. However, we are not aware of how wide-spread this is, but based on the list of supported providers listed on SIDN's website, none of the providers from Figure 7 support it yet. We thus believe that most .nl domains do not use this yet.

Without .nl Control in place, the registrant will receive an email on updates regarding the domain name, but the update will still go through nearly instantly. Once logged in, an attacker can generally do four main things:

1. The attacker can request a transfer code. This enables an attacker to transfer the domain, thereby fully taking control of a domain name (and taking control away from the original holder);
2. The attacker can change the DNS name servers to their own. This enables an attacker to (selectively) change responses to DNS requests;
3. The attacker can view and change the holder information, such as phone numbers and addresses;
4. The attacker can disable or change the DNSSEC keys.

With control over the DNS responses, an attacker can do many things. This list is not exhaustive:

1. Control what users see on the website;
2. Capture the cookies of users, thereby being able to impersonate that user on the website;
3. Obtain data victims enter on the attacker's website served on the original domain;
4. Receive all the email for that domain, including password reset emails;
5. Log in to any external systems that can be accessed after requesting a password reset, such as payroll and other administrative systems, as well as external cloud infrastructure;
6. Request Transport Layer Security (TLS) certificates for the domain, making the connecting appear secure in e.g. browsers;

Agent	Registration details						Two-factor authentication				
	Name	*Email*	*Tel. no.*	*Address*	*D.o.B.*	*Sec. Q.*	*Email*	*Phone*	*TOTP*	*FIDO*	*mandatory*
A	✓	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗
B	✓	✓	✓	✓	✓	✗	✗	✗	✓	✗	✗
C	✓	✓	✓	✓	✗	✗	✓	✗	✓	✓	✓
D	✓	✓	✓	✓	✗	✗	✗	✗	✓	✗	✗
E	✓	✓	✓	✓	✗	✗	✗	✗	✓	✗	✗
F	✓	✓	✓	✓	✗	✗	✗	✗	✓	✗	✗
G	✓	✓	✓	✓	✗	✗	✗	✗	✓	✗	✗
H	✓	✓	✓	✓	✗	✗	✗	✗	✓	✗	✗
I	✓	✓	✓	✓	✗	✗	✓	✗	✓	✗	✓
J	✓	✓	✓	✓	✗	✗	✗	✗	✓	✗	✗

Figure 7: We registered domain names at each organisation from Figure 3, and kept track of the information they require during the registration. From left to right: name, email address, telephone number, address, date of birth, and security questions, as well as its two-factor authentication options. We exclude protections only applied to high-risk IP addresses (such as public Virtual Private Networks (VPNs)), like typing in a code sent to an email or selecting all fire hydrants, as these can be circumvented by, e.g., using a residential IP address

7. Provide false responses to Application Programming Interface (API) calls, manipulating the data retrieved by e.g. third party integrations;

8. In the case of Internet of Things devices, publish and push updates to end-user devices;

When an attacker can control the DNS responses, they effectively have 'the keys to the castle'. This is comparable to other digital threats an organisation might face. We pick two we consider close when it comes to impact:

6.1 Comparison with Ransomware

Another threat for an organisation is ransomware. Ransomware causes losses of billions of dollars [28]. Ransomware is a subset of malware that blocks the user or organisation out of their own systems and/or files. In many cases it encrypts the files on the system, requiring a ransom to be paid to retrieve the key. This amount can be millions of dollars [42].

We argue that there are a lot of parallels to be drawn between the ransomware threat and the threat of a domain takeover, although there are differences as well:

1. Both grind operations in a lot of organisations to a halt. Systems will become unavailable (e.g. due to Single Sign-On (SSO), email, etc. being unavailable);

2. Both have the attacker as party who can solve the issue they created, though in the case of a domain takeover the registry has authority as well;

3. Ransomware is destructive by encrypting files, domain takeovers only get rid of the reference (e.g. a local DNS resolver that overrides the takeover would still allow internal services to keep operating);

4. Both allow for access to e.g. email (in the case of a domain takeover only to email that is received after the takeover), allowing for things like password resets in (remote) systems and access to potentially sensitive data;

5. Both allow an attacker to perform destructive actions on third-party platforms, such as Software-as-a-Service (SaaS);

6. Both allow an attacker to exfiltrate data from third-party platforms, such as Software-as-a-Service (SaaS), and extort the victim to prevent publication of secret or sensitive data;

6.2 Comparison with DDoS Attacks

DDoS attacks can make an organisation's digital infrastructure (e.g. website, email) unavailable for the duration of the DDoS attack. Like ransomware, we believe parallels can be drawn between the threat of a DDoS attack and the threat of a domain takeover.

1. Both have the potential to grind operations within an organisation to a halt. Things like email, Single Sign-On (SSO), become unavailable for the duration of the DDoS attack, or until the domain takeover is sorted;

2. In both cases the attacker is the party who can solve the issue, either by giving back control of the domain name, or stopping the DDoS attack;

3. DDoS attacks require the attacker to keep attacking the organisation, once they stop the services become available again. A domain takeover does not require this;

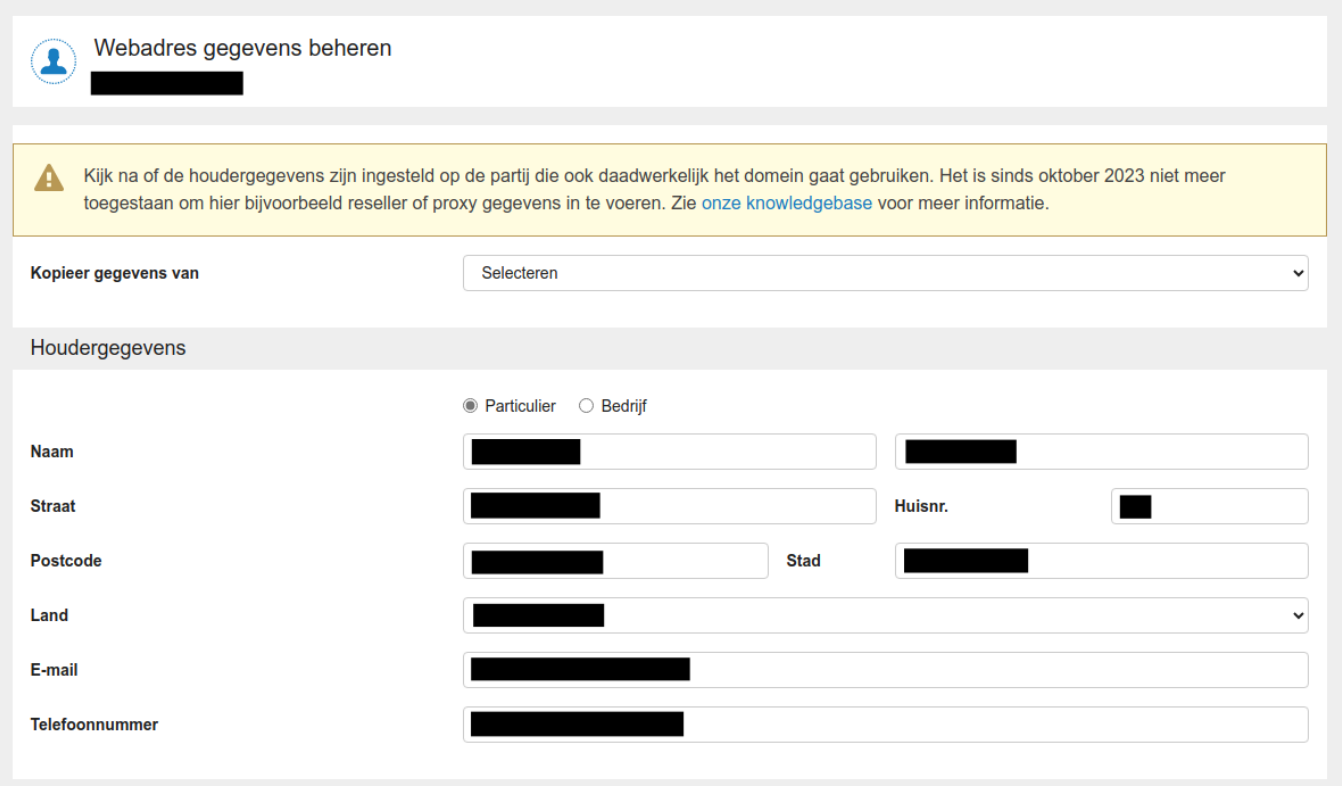

Figure 8: The information *A* shows as account holder data. From top to bottom: 1) Name, 2) Street / House no., 3) Postal code / City, 4) Country, 5) E-mail address, and 6) Phone number.

4. A DDoS attack does not allow the attacker access to sensitive data, whereas a domain takeover can result in receiving email for the organisation, including possibly password reset emails, allowing them access to sensitive systems.

6.3 Relative Impact

When we look at ransomware, we see that the impact of a domain takeover is likely lower. When we look at a DDoS attack, we see that the impact is likely higher. We thus believe that these two examples provide appropriate upper and lower bounds. This allows us to estimate the impact of a domain takeover.

We believe that this can be modelled in similar ways to the DDoS attacks described by Wolthuis et al. [39] and ransomware attacks from Gomez et al. [12]. We believe that there is reason to take the risk of a domain takeover into account in a similar vein to DDoS and ransomware attacks.

Combined with our findings from Section 5, we assess the likelihood of a threat occurring to be on the lower side of the spectrum, but the likelihood of adverse impacts on the higher end of the spectrum, based on the NIST assessment from Figure 1.

The precise impact depends too much on the organisation to make a clear assessment. For example, for an organisation providing publicly available weather data, the risk of data manipulation or data not being available might be far greater than data being leaked. Similarly, for a psychiatrist, confidentiality might be far more important than the availability of their systems. Due to the varying nature of the various kinds of attacks for various organisations, we believe we cannot make general statements about where in the risk table DDoS attacks, ransomware, and domain takeovers would land.

7 Comparison with Other Countries

We aim to analyse the best-case scenario. .nl has a diverse market of agents, and a lot of them are based in the Netherlands and used for things related to the Netherlands. This allows us to reasonably use .nl to approximate the state in the Netherlands. The moderate risk we find as described in the previous section is based on analysis from .nl, and the Netherlands is a role model country according to the ITU cybersecurity index [18]. The Netherlands scores 20 out of 20 in legal measures, technical measures, organization measures, and cooperation measures. Only capacity development is 19.22 out of 20. We hence believe that the risk assessment is skewed towards painting a ‘too positive’ picture rather than

a ‘too negative’ picture.

We analyse the state of agents, not a few agents specifically. The Netherlands has a competitive registrar market, and this applies to most countries in Europe. CENTR, the association of European country code top-level domain (ccTLD) registries, writes in a report that “The distribution networks of registrars selling European ccTLDs are on average considered competitive based on a widely used measure of market concentration”, based on the calculation of the Herfindahl-Hirschman Index (HHI) over 11 CENTR member registries in 2023 [3]. According to CENTR, in 2024 the median Herfindahl-Hirschman Index (HHI) for registrars under CENTR member registries was 1324. This suggests that registrar channels for most member registries are considered competitive (i.e. have low concentration). 17 members have competitive (HHI under 1500) registrar markets and a further 9 are considered moderately concentrated. There are no member registry’s with highly concentrated registrar channels.

We also used the same Cloudflare Radar data used in Figure 3 for .fr, .no, and .uk. We show their agent market share in Figure 9. We picked these three because these ccTLDs are mainly used by things related to France, Norway, and the United Kingdom respectively, all three countries are also considered role models by the ITU, and they also all have RDAP availability that includes agent information (as WHOIS data can be tricky to parse [21]) so we could analyse their data. For all three we see a large local presence in registrars, although American-based ones are popular in the United Kingdom as well. We also noticed some overlap, where *C*, *CSC Global*, and *MarkMonitor* tend to be popular at these three ccTLDs too.

The impact of a domain takeover we describe in Section 6 is not specific to any country or (cc)TLD. Only the likelihood of occurrence changes. We believe the risk in other role model countries is similar, and might be higher in other countries, for example in countries where TOTP is not as widespread.

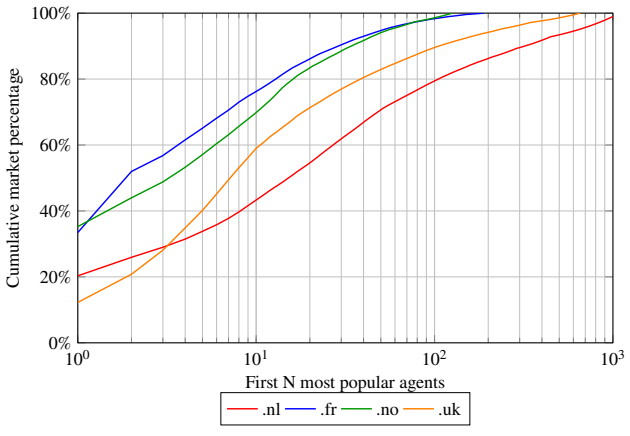

Figure 9: The market share for the TLDs .nl, .fr, .no, and .uk, shown as a cumulative percentage for the top N registrars.

8 Disclosure Process

We decided to contact the agents from Figure 5 that did not apply per-account rate limiting. *A*, *C*, *E*, *I*, and *J* have a *security.txt* [11] set up. *B*, *F*, *G* do not have a security.txt page, but do have a responsible disclosure page set up. Only *D* and *H* have neither.

We contacted *C*, *E*, and *J* using their contact address in security.txt, *B*, *F*, and *G* via their responsible disclosure pages, and *D* and *H* and via their general contact form. We sent them a message along the following lines (slightly adjusted for those who did do IP rate limiting):

> *Subject: TOTP rate limiting*
>
> Dear <organisation>,
>
> At the University of Twente , we have been conducting research into the security of domain names in the Netherlands and how registrars and resellers prevent unauthorised access to, for example, the web portal. Part of this involves analysing which security methods are used to prevent brute force attacks on two-factor authentication.
>
> We noticed that it appears that no rate limiting is being applied, which is why we are sending this email.
>
> Our test setup is as follows: we log in normally with a username and password, then we try 100 incorrect TOTP codes in 1 second, after which we enter the correct TOTP code. We analyse whether and how rate limiting is applied based on the responses we receive from the 100 incorrect TOTP codes.
>
> RFC 4226 section 7.3 stipulates that rate limiting must be implemented per user and not per session, in order to counter parallel attacks. We have limited ourselves to 100 attempts so as not to put unnecessary pressure on your servers. Without rate limiting, with 100 requests per second, a correct TOTP code can be guessed in just under an hour on average.
>
> Is our assumption that no rate limiting is applied correct? And if so, are there plans to implement rate limiting in the next three months? Our goal is not to ‘name and shame’; we therefore do not publish any names in the study. Nevertheless, we would like to mention that parties have implemented rate limiting, so we would appreciate a response.
>
> We look forward to hearing from you. Please feel free to contact us with any questions or comments.
>
> Yours faithfully,
> <name>
>
> <email>
> <phone number>
>
> <address>

We received a confirmation within a day from two out of eight notifications, namely from *D* and *H*, and also from *E* after a week. We have not received a message from the others till this date. None have stated that the issue was resolved.

9 Discussion

Signing up for the agents listed in Figure 7 was more difficult than expected. Two required manual intervention by customer support after we contacted them that the registration had seemingly not worked. One kept the old DS-record (a record used by DNSSEC) when we moved to our own DNS name servers. When we asked them how to remove or change that, they disabled DNSSEC for us without requiring any verification.

However, as we show in Section 5.3, when trying to gain access to an account, all agents seem to have procedures in place to prevent unauthorised takeovers. We did not expect this result. We also expected the registrar to ask verification question such as email, telephone number, address, and date of birth, but none of them requested that information.

These events, the detected protection methods against TOTP brute forcing, as well as the general impression of the online portals, gave us the feeling that many of these systems are rather brittle, but the limitations known to the organisations. Sadly we have no way to fully verify that without access to the source code.

We were surprised to find that none of the agents sent us an email after a failed TOTP login attempt. Whilst the result from Figure 5 do show that some of them prevent against brute forcing, it seems like TOTP support and the security implications were not fully thought through.

We were also surprised to find that there are simultaneously several popular agents such as *CSC Global* and *MarkMonitor* that offer brand protection services, indicating that there is a realisation at organisations that their domain name is important, yet simultaneously a lack of formal impact models of what would happen if a domain name were to be taken over.

9.1 Limitations

There are some limitations to our approach, which we want to specifically highlight separately.

As alluded to in Section 4.1, we could not test corporate registrars, which are the most prominent in the 480 most popular domain names as we describe in Section 5.4. Additionally, even though we believe our results are applicable outside the Netherlands as well as we describe in Section 7, care and attention is still required when applying these results outside the Netherlands. Furthermore, it should be noted that country-code TLDs such as .nl can also be used outside the context of the Netherlands and vice versa.

The password availability metrics mentioned in Section 5.1 assumes the same password reuse numbers as found in other research. These numbers may also vary wildly, e.g., in case there is a recent large-scale password breach. Ethical and legal limitations prevent us from verifying these claims. Similar ethical and legal limitations prevent us from sending a fake ID card scan as was requested in some of the helpdesk calls from Section 5.3.

10 Recommendations

Whilst we believe that it is worthwhile to analyse the market as a whole, we also want to share a concrete list of points to consider for organisations and individuals:

1. Make an impact assessment of what would happen if your domain name is taken over. In Section 6 we describe some of the potential impacts, but as every organisation is different, this list will be different as well;

2. Monitor your domain name, including its DNS records, DNSSEC configuration, and WHOIS data. Changes here can be an indication that someone is tampering with your domain;

3. Pick a registrar wisely. There are still registrars out there that do not support things like second factor authentication. Having a trustworthy and secure registrar is thus crucial;

4. Use an email alias for your contact address in the WHOIS, and do not use this email to sign up at other places or for logging in to the administrative portal.

11 Conclusion

We analysed the risk of a domain name takeover by examining the technical likelihood of a domain takeover and the impact if it were to occur.

Firstly, we looked at the security of domain names in the Netherlands. When it comes to the security at popular domain registrars in the Netherlands. The positive part is that multi-factor authentication is available in all cases we examined and processes are in place to verify the identity of the domain holder, even over the phone. Additionally, larger organisations seem to use their own registrar or a registrar specifically aimed at protecting brands. However, things as TOTP are still often easy to brute force in most cases, and notifications and confirmation of account access and changes are often lacking, making it potentially difficult to notice when an account has been breached. Email addresses for domain management are reused for other purposes, and are frequently personal email addresses that appear in public breaches and data leaks.

Secondly, we find that gaining illicit access to a domain name can have great impact. We analysed the impact of a domain name takeover, and find that its potential consequences are significant, especially with the uptake in use of cloud

services, and likely between that of a DDoS attack and a ransomware attack.

These two factors combined lead us to believe the risk of a domain name is potentially great. We have no reason to believe these results are unique to The Netherlands, and are likely relevant throughout Europe and even globally. We therefore urge organisations to take the impact of a domain takeover for their organisation into consideration, and provide recommendations and points for organisations to consider.

Future Work

We believe it is worthwhile to repeat this experiment in the future, especially after the disclosure we describe in Section 8. Additionally, we believe it is worth investigating this with different markets in and outside Europe to confirm our assumptions from Section 7.

References

[1] Lillian Ablon, Paul Heaton, Diana Lavery, and Sasha Romanosky. Consumer attitudes toward data breach notifications and loss of personal information. 2016. doi:10.7249/rr1187.

[2] Gautam Akiwate, Stefan Savage, Geoffrey M Voelker, and Kimberly C Claffy. Risky BIZness: risks derived from registrar name management. In *Proceedings of the 21st ACM Internet Measurement Conference*, pages 673–686, 2021.

[3] CENTR. CENTRstats Global TLD Report, 1 2024. URL: https://centr.org/images/global_tld_report_2024_1.pdf.

[4] Richard Chirgwin. New York Times, twitter domain hijackers "came in through front door", Nov 2013. URL: https://www.theregister.com/2013/08/27/twitter_ny_times_in_domain_hijack/.

[5] Taejoong Chung, Roland van Rijswijk-Deij, David Choffnes, Dave Levin, Bruce M Maggs, Alan Mislove, and Christo Wilson. Understanding the role of registrars in DNSSEC deployment. In *Proceedings of the 2017 Internet Measurement Conference*, pages 369–383, 2017.

[6] CISA. CISA insights – CYBER: Mitigate DNS Infrastructure Tampering. URL: https://www.cisa.gov/sites/default/files/publications/CISAInsights-Cyber-MitigateDNSInfrastructureTampering_S508C.pdf.

[7] Cointelegraph. What is DNS hijacking? how it took down Curve Finance's website, May 2025. URL: https://www.tradingview.com/news/cointelegraph:9a15fa371094b:0-what-is-dns-hijacking-how-it-took-down-curve-finance-s-website/.

[8] Leslie Daigle. WHOIS Protocol Specification. RFC 3912, September 2004. URL: https://www.rfc-editor.org/info/rfc3912, doi:10.17487/RFC3912.

[9] Anupam Das, Joseph Bonneau, Matthew Caesar, Nikita Borisov, and XiaoFeng Wang. The tangled web of password reuse. In *NDSS*, volume 14, pages 23–26, 2014.

[10] Lee Davis. Keys to the (SAAS) kingdom, May 2025. URL: https://cybercx.com/blog/keys-to-the-saas-kingdom/.

[11] Edwin Foudil and Yakov Shafranovich. A File Format to Aid in Security Vulnerability Disclosure. RFC 9116, April 2022. URL: https://www.rfc-editor.org/info/rfc9116, doi:10.17487/RFC9116.

[12] Gibran Gomez, Kevin van Liebergen, and Juan Caballero. Cybercrime bitcoin revenue estimations: Quantifying the impact of methodology and coverage. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, CCS '23, page 3183–3197, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3576915.3623094.

[13] Paul E. Hoffman. DNS Security Extensions (DNSSEC). RFC 9364, February 2023. URL: https://www.rfc-editor.org/info/rfc9364, doi:10.17487/RFC9364.

[14] ICANN. Registering Domain Names. URL: https://www.icann.org/resources/pages/register-domain-name-2017-06-20-en.

[15] ICANN. SAC 074 | SSAC Advisory on Registrant Protection: Best Practices for Preserving Security and Stability in the Credential Management Lifecycle. URL: https://www.icann.org/resources/files/1194801-2015-11-03-en.

[16] ICANN. What Does ICANN Do? URL: https://www.icann.org/resources/pages/what-2012-02-25-en.

[17] Joint Task Force Transformation Initiative. Guide for conducting risk assessments, Sep 2012. URL: https://csrc.nist.gov/pubs/sp/800/30/r1/final.

[18] International Telecommunication Union (ITU). Global Cybersecurity Index 2024, May 2024. URL: https://www.itu.int/epublications/publication/global-cybersecurity-index-2024.

[19] M. Asif Khan and Mureed Hussain. Cyber security quantification model. In *Proceedings of the 3rd International Conference on Security of Information and Networks*, SIN '10, page 142–148, New York, NY, USA, 2010. Association for Computing Machinery. doi:10.1145/1854099.1854130.

[20] Wilson Lian, Eric Rescorla, Hovav Shacham, and Stefan Savage. Measuring the practical impact of {DNSSEC} deployment. In *22nd USENIX Security Symposium (USENIX Security 13)*, pages 573–588, 2013.

[21] Suqi Liu, Ian Foster, Stefan Savage, Geoffrey M. Voelker, and Lawrence K. Saul. Who is .com? learning to parse whois records. In *Proceedings of the 2015 Internet Measurement Conference*, IMC '15, page 369–380, New York, NY, USA, 2015. Association for Computing Machinery. doi:10.1145/2815675.2815693.

[22] P. Mockapetris. Domain names - implementation and specification. RFC 1035, November 1987. URL: https://www.rfc-editor.org/info/rfc1035, doi:10.17487/RFC1035.

[23] David M'Raihi, Frank Hoornaert, David Naccache, Mihir Bellare, and Ohad Ranen. HOTP: An HMAC-Based One-Time Password Algorithm. RFC 4226, December 2005. URL: https://www.rfc-editor.org/info/rfc4226, doi:10.17487/RFC4226.

[24] Morium Akter Munny, Mahbub Alam, Sonjoy Kumar Paul, Daniel Timko, Muhammad Lutfor Rahman, and Nitesh Saxena. Infrastructure patterns in toll scam domains: A comprehensive analysis of cybercriminal registration and hosting strategies. In *2025 APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–13, 2025. doi:10.1109/eCrime66972.2025.11327851.

[25] NCSC-UK. Good security practice for domain registrars, Mar 2025. URL: https://www.ncsc.gov.uk/collection/security-practice-domain-registrars.

[26] Andy Newton and Scott Hollenbeck. JSON Responses for the Registration Data Access Protocol (RDAP). RFC 7483, March 2015. URL: https://www.rfc-editor.org/info/rfc7483, doi:10.17487/RFC7483.

[27] OWASP. Multifactor authentication cheat sheet. URL: https://cheatsheetseries.owasp.org/cheatsheets/Multifactor_Authentication_Cheat_Sheet.html.

[28] Harun Oz, Ahmet Aris, Albert Levi, and A. Selcuk Uluagac. A survey on ransomware: Evolution, taxonomy, and defense solutions. *ACM Comput. Surv.*, 54(11s), September 2022. doi:10.1145/3514229.

[29] Paul Rascagneres and Warren Mercer. Dnspionage campaign targets middle east, Sep 2018. URL: https://blog.talosintelligence.com/dnspionage-campaign-targets-middle-east/.

[30] Hooman Razavi, Mohammad Reza Jamali, Morvaridsadat Emsaki, Ali Ahmadi, and Mostafa Hajiaghei-Keshteli. Quantifying the Financial Impact of Cyber Security Attacks on Banks: A Big Data Analytics Approach. In *2023 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, pages 533–538, 2023. doi:10.1109/CCECE58730.2023.10288963.

[31] Johann Schlamp, Josef Gustafsson, Matthias Wählisch, Thomas C. Schmidt, and Georg Carle. The abandoned side of the internet: Hijacking internet resources when domain names expire. In Moritz Steiner, Pere Barlet-Ros, and Olivier Bonaventure, editors, *Traffic Monitoring and Analysis*, pages 188–201, Cham, 2015. Springer International Publishing.

[32] Giovanni Schmid. Thirty Years of DNS Insecurity: Current Issues and Perspectives. *IEEE Communications Surveys & Tutorials*, 23(4):2429–2459, 2021. doi:10.1109/COMST.2021.3105741.

[33] Stichting Internet Domeinregistratie Nederland. .NL control: No domain name changes without permission. URL: https://www.sidn.nl/en/product/nl-control.

[34] Stichting Internet Domeinregistratie Nederland. Vanaf 1 oktober geldt een verbod op privacy- en proxyservices onder .nl. URL: https://www.sidn.nl/nieuws-en-blogs/vanaf-1-oktober-geldt-een-verbod-op-privacy-en-proxyservices-onder-nl.

[35] Kurt Thomas, Frank Li, Ali Zand, Jacob Barrett, Juri Ranieri, Luca Invernizzi, Yarik Markov, Oxana Comanescu, Vijay Eranti, Angelika Moscicki, Daniel Margolis, Vern Paxson, and Elie Bursztein. Data breaches, phishing, or malware? understanding the risks of stolen credentials. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, page 1421–1434, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3133956.3134067.

[36] Troy Hunt. Have I Been Pwned. URL: https://haveibeenpwned.com/About.

[37] Olivier van der Toorn, Moritz Müller, Sara Dickinson, Cristian Hesselman, Anna Sperotto, and Roland van Rijswijk-Deij. Addressing the challenges of modern DNS a comprehensive tutorial. *Computer Science Review*, 45:100469, 2022. URL: https://www.scienc

edirect.com/science/article/pii/S1574013722000132, doi:10.1016/j.cosrev.2022.100469.

[38] Thomas Vissers, Timothy Barron, Tom Van Goethem, Wouter Joosen, and Nick Nikiforakis. The wolf of name street: Hijacking domains through their nameservers. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, page 957–970, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3133956.3133988.

[39] Reinder Wolthuis, Frank Phillipson, Hidde-Jan Jongsma, and Peter Langenkamp. A framework for quantifying cyber security risks. *Cyber Security: A Peer-Reviewed Journal*, 4(4):302, Jun 2021. doi:10.69554/cykn3231.

[40] Daniel W Woods and Rainer Böhme. SoK: Quantifying cyber risk. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 211–228. IEEE, 2021.

[41] Yunyi Zhang, Baojun Liu, Haixin Duan, Min Zhang, Xiang Li, Fan Shi, Chengxi Xu, and Eihal Alowaisheq. Rethinking the security threats of stale DNS glue records. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 1261–1277, Philadelphia, PA, August 2024. USENIX Association. URL: https://www.usenix.org/conference/usenixsecurity24/presentation/zhang-yunyi-rethinking.

[42] Aaron Zimba and Mumbi Chishimba. On the economic impact of crypto-ransomware attacks: The state of the art on enterprise systems. *European Journal for Security Research*, 4(1):3–31, 2019.

A Phone flow chart

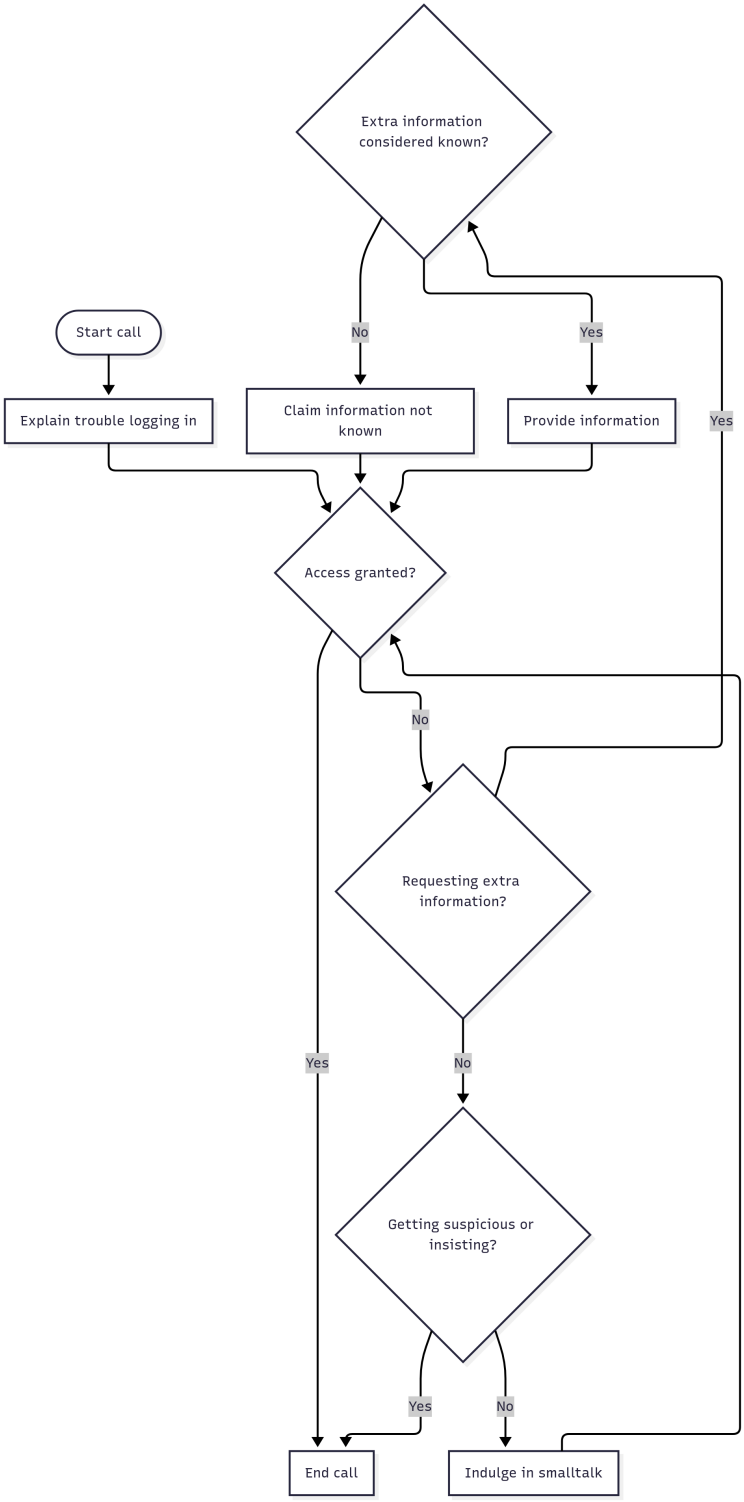

Figure 10: The flowchart we use for trying to gain access to an account by calling customer service. We do not press or guilt-trip the employee – we only provide answers to questions that are considered "available data", and claim not to know the information otherwise.