Governmental Domain Name Management: Policy versus Practice

W.J. Kobes

MSc Thesis
February 2023

Supervisors:
dr. J. van der Ham
dr. A. Abhishta
prof. dr. R. Torenvlied
dr. L.A.N. Long
dr. M.R. Koot

Programs:
MSc Computer Science
MSc Public Administration

Wouter Kobes: *Governmental Domain Name Management: Policy versus Practice,*
© February 18, 2023

CHAIRS:
Master Computer Science, 4TU Cyber Security
Design and Analysis of Communication Systems
Faculty of Electrical Engineering, Mathematics and Computer Science

Master Public Administration
Department of Public Administration
Faculty of Behavioural, Management and Social Sciences

LOCATION:
University of Twente
P.O. Box 217
7500 AE Enschede
The Netherlands

Secura B.V.
Karspeldreef 8
1101 CJ Amsterdam
The Netherlands

ABSTRACT

Domain names and the domain name system (DNS) are two core technologies that provide a backbone to the internet. Domain names are often used to present websites, or for sending and receiving e-mail. The domain name system allows for the translation from human-readable domain names to computer-readable IP addresses. Since the introduction of this technology in the early 1980s, additional technologies have been developed that extend or secure the DNS. Extra security settings have to be configured to ensure domain names are up to date with current cyber security requirements. To be able to use a domain name, it has to be registered at a licensed domain name registrar. Domain name registrations usually come with a yearly upkeep fee. In case a registration is terminated, the domain will be freed for new registrations. Organizations, companies and governments tend to own multiple, in some cases many, domain names. Overseeing these domain names can be challenging. Especially for governments, often with a highly-decentralized structure, centralized domain name management requires adequate policy-making. This thesis studied how domain name management is performed by the Dutch government.

The first contribution of this work is to identify three categories of cyber security risks that involve domain names from the perspective of domain name owners. In (sub)domain takeovers, adversaries gain control over a domain name that is supposed to be in control of the victim. The risk of impersonation and typosquatting involves adversaries that attempt to abuse domain names similar to those of their victim, with small differences like common typing mistakes or optically similar characters. The third category is non-compliance with current security standards. Several security standards need to be implemented at the domain name level and are required to allow the secure use of websites and e-mail.

The Dutch government published a wide range of policy documents that involve domain name management. As a second contribution, this work reconstructs the policy theory of domain name management policies in the Dutch government. In general, the topic is seen in three policy fields: archiving, communication and security & compliance. Policies about archiving deal with how domain names and their websites should be collected and stored to comply with legislation. Communication policies, like a central domain name policy, aim to make domain names clear and recognizable for citizens. Lastly, security & compliance policies are in place to ensure governmental domain names are resilient and compliant with current standards.

After analyzing governmental domain names in practice, several shortcomings are identified that may impose cyber security risks. These risks can be partly attributed to not adhering to policies, and partly to existing policies being insufficient to cover all identified risks. A third contribution of this thesis is the use of novel techniques to discover domain names that belong to the government.

Based on the previous findings, this thesis provides concrete recommendations for the Dutch government which can be used to improve its domain name management.

Keywords: domain names, cyber security, public administration, DNS, computer science, policy theory

ACKNOWLEDGMENTS

Dear reader,

This thesis represents the epilogue of my student career, which started in 2014 when I moved to Enschede. After two years of hard work, distractions, and ups and downs, I am proud to present the proverbial icing on the cake of my studies in computer science and public administration. I hope you enjoy reading it.

This work would not have been possible without the support of many others. First of all, I would like to thank my supervisors, René Torenvlied and Le Anh Long on behalf of MSc Public Administration and Jeroen van der Ham and Abhistha on behalf of MSc Computer Science. Thank you for your insights, feedback, and most of all patience throughout this research. I have enjoyed the many discussions on domain names and cyber security in a public context we have had. My thanks also go out to Matthijs Koot and Erwin Jansen of Secura B.V., for providing a pleasant working environment and insights on cyber security problems in practice. Then, I would like to thank Mattijs Jonker for his support in getting access to and conducting the data analysis on the OpenINTEL data set.

During my academic journey, I have had great support and encouragement from my family and friends, both in Enschede and Utrecht. My explicit thanks go out to my parents, Bernard and Enske, for motivating me to pursue my academic and non-academic interests and to Elin, for the limitless support during the past two years.

For further questions or discussions on the topic of domain name management, feel free to contact me at w.j.kobes@alumnus.utwente.nl.

Kind regards,

Wouter Kobes

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

ASCII	American Standard Code for Information Interchange
ccTLD	country code top-level domain name
CA	Certificate Authority
CAA	Certification Authority Authorization
CT	Certificate Transparency
DDoS	Distributed Denial-of-Service
DKIM	DomainKeys Identified Mail
DMARC	Domain-based Message Authentication, Reporting, and Conformance
DNSSEC	Domain Name System Security Extensions
DNS	Domain Name System
DoT	DNS over TLS
DoH	DNS over HTTPS
DPC	Dienst Publiek en Communicatie
FQDN	fully qualified domain name
gTLD	generic top-level domain name
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
IDN	Internationalized Domain Name
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
PQDN	partially qualified domain name

RDAP	Registration Data Access Protocol
RDDI	Rijksprogramma voor Duurzaam Digitale Informatiehuishouding
RFC	Request For Comments
SAN	Subject Alternative Name
SIDN	Stichting Internet Domeinregistratie Nederland
SPF	Sender Policy Framework
TLD	Top-level domain name
TLS	Transport Layer Security
VoRa	Voorlichtingsraad
WWW	World Wide Web

Part I

THESIS

1 INTRODUCTION

1.1 TOPIC INTRODUCTION

Since the introduction of the Internet and the World Wide Web (WWW) in the 1980s, the world engaged in an unprecedented level of interconnectedness [50]. These technologies played a large role in globalization, with the Internet accessible to about 60% of the world's population in 2020 [45]. Although the technical developments around the Internet are and always have been an ongoing process, many of the technical foundations still date back to the second half of the 20th century. Enhancements, additions, and new technologies have been added on top of existing technologies and protocols, allowing the Internet community to keep up with the newest technical innovations and increasing demand for connections and higher connection speeds. At the same time, the stacked approach of Internet development has ensured backward compatibility for legacy systems and services.

However, there are several occasions where the age of technology is causing problems. One well-known example is found within the Internet Protocol (IP). This protocol allows for the identification of and communication with different computers within a network. To achieve this, systems use unique addresses, known as IP addresses. IP is one of the fundamentals of the Internet and has two different versions that are currently in use. The first (and still) widespread version is the Internet Protocol version 4 (IPv4), introduced in 1981. One problem with IPv4 is that it allows for up to 4.3 billion unique IP addresses, which is insufficient for the enormous number of users and appliances that want to connect to the Internet. A solution for this problem was already accepted as a standard in 1995, by means of Internet Protocol version 6 (IPv6), which, besides other technical advancements, supports about a *trillion trillion trillion* unique IP addresses. The IP example is striking because, at the time of writing, the transition toward the use of IPv6 is far from complete. For illustration, the IPv6 adoption rate in the Netherlands in August 2021 was measured at less than 30% [81]. This example shows that the Internet remains highly dependent on older technologies, even when newer and better alternatives are available.

In this thesis, the focus lies on a different, yet related, technology, which is the concept of *domain names*. A domain name is a human-readable name that can point toward specific servers or services, commonly using an IP address. The first domain name dates back to March 1985, and, as the example above, still depends on the same technological principles today. Domain names and their applications currently all operate within the so-called Domain Name System (DNS). Since its original introduction, new technologies have been developed that extend and secure the DNS to keep up with current requirements.

A possible composition of a domain name is 'example.com', which consists of the domain name label 'example' and the Top-level domain name (TLD) 'com'. While domain name labels can be freely chosen, the choice of TLDs is limited to predefined labels and country codes. Although domain names serve various purposes within the Internet, their primary usage is providing access to websites in the WWW. Websites are served over the HTTP and HTTPS protocols and can be accessed by entering the domain name in a *web browser*. This means that domain names are the primary gateway for humans to access the WWW. Through the use of *hyperlinks*, websites can redirect users to other domain names. These connections make up the 'web' and can together be observed as the largest IT communication network apart from the Internet itself. Since domain names are primarily used by humans, implies that human fallibility is to be regarded when assessing the risks involving domain names. Furthermore, domain names that are not configurated to meet the latest security requirements, may impose additional risks for their users.

One fundamental property of domain name registrations is that they are temporary. Registrations need to be renewed and annual fees need to be paid. Whenever a domain name registration is canceled, the domain name usually becomes available for registration again after a grace period. The aforementioned hyperlinks, however, do not disappear whenever a registration of that domain name ends.

Registrations can be done by individuals, as well as companies and organizations. It is common for organizations and companies to uphold multiple domain name registrations. For example, reasons for domain name registrations are the maintaining of multiple trademarks, separation of various departments with different websites, having operations in various countries using different TLDs, or collaborations with other organizations. When the domain name portfolio of an organization grows larger, keeping track of these registrations becomes more difficult. Within larger organizations, multiple divisions could be registering domain names, without central communication or coordination. Maintaining oversight of the domain name registrations of an organization is key to both preventing and detecting domain name-related security risks. This problem will be the main focal point of this thesis, which will be referred to as the problem of *domain name management*.

The context in which the problem is analyzed is the Dutch central government. Due to its size and complex organizational structure, domain name management is not trivially maintaining a list of domain name registrations. Policies have to be in place to regulate domain name usage within its highly decentralized structure. Domain names of government should be recognizable and distinguishable from other domain names, while at the same time being optimally configured in terms of cyber security. At the same time, there are already several existing policy documents of the government related to domain name management that can be found online. Consequently, domain name management in the Dutch government is a topic that is worth researching.

1.2 PROBLEM STATEMENT, MOTIVATION AND JUSTIFICATION

While state-of-the-art literature does discuss domain names and related security risks, no research has been conducted on the specific problem of domain name management. This is remarkable since the problem is not novel. In the past few years, there have been various security incidents that can be (partly) attributed to improper domain name management [93]. Current research treats domain names as technologically isolated 'islands', where attention is directed to the security risks of the technical properties of the DNS, and how they could be abused [88]. A study in domain name management should include the risks of interactions between domain names, and the risks of having an incomplete or incorrect overview of domain name registrations. Furthermore, specifically those risks affecting domain name owners need to be considered, since security risks in the technical implementation of DNS are not caused nor solved by implementing proper domain name management. Risks affecting underlying DNS technologies are better suited to be addressed by Internet researchers and DNS operators. This distinction of risks is currently not made in the literature.

This thesis argues that proper domain name management is important from a security perspective, because domain names are entry-level targets for cyber attacks. Securing domain names requires a correct configuration, as well as regular maintenance of the services that a domain name hosts. Organizations might intend to maintain certain security standards for all of their websites and services, yet could turn out to still be vulnerable to cyber-attacks through unknown and unsupervised domain names. When missing from their domain name management, domain names could turn into these organizations' weakest links in terms of cyber defense. At the same time, when a domain name is wrongfully part of a management system, organizations expect the content and services on this domain to be managed by them. In this case, the domain could be abused in trust-based attacks, like the sending of SPAM or phishing attempts.

The problem of domain name management is not merely about its underlying technologies. Policies have to be adopted that deal with the managerial aspects of the problem. This makes the problem an interdisciplinary one, requiring academic attention for both the technological and policy issues alike. Current policies should be further studied for possible hiatuses in mitigations for relevant cyber security risks.

The motivation for conducting this research came forth from analyzing the resources of domain names for the Dutch government. This government makes use of a domain name management system, of which part is made public [76]. While investigating this domain name management system, several shortcomings were identified that led to cyber security risks. These initial findings sparked the research conducted in this thesis, into the Dutch government's approach to domain name management.

The Dutch central government is a suitable case subject for researching governmental domain name management, due to its organizational complexity and the diversity of the domain name portfolio. Firstly, the central government consists of a large number of departments, ministries, and agencies, which house various IT departments and work with various external IT suppliers. Secondly, there is a clear distinction between central and decentral governments like provinces and municipalities, whereas requirements for effective domain name management could transcend the separation of IT responsibilities. Lastly, the domain name portfolio is diverse in the sense that it contains domain names under multiple TLDs and domain name registrations are managed by various organizations.

The DNS and Internet Standards, in general, have received fairly little attention in academic literature studies in the earlier decades of the Internet. After the usage of the Internet increased, this attention increased. That DNS is a topic that requires academic attention, becomes clear in the 1995 work of Bellovin. This work led to the development of the Domain Name System Security Extensions (DNSSEC) standard, addressing critical security issues in the fundamental design of DNS. Notably, the paper was withheld for over four years, quoting the author: 'The paper was held back - not suppressed; no external agency applied any pressure, though there were certainly others who were happy it was not published at the time - because it described a serious vulnerability for which there was no feasible fix' [7]. As with the earlier example of IPv6, the DNSSEC standard also lacks full adoption by the Internet community, even though the attack vectors found by Bellovin can still be abused when DNSSEC is not implemented in 2021 [80]. For instance, the adoption of DNSSEC in the Netherlands was measured at 58 percent in 2023 [79].

The identification of security risks in current implementations of domain name management, accompanied by the scarce research conducted on Internet Standards and known vulnerabilities in DNS, justifies conducting this research. The problem requires academic attention and solutions to this problem may increase the level of security of the World Wide Web in general, as well as potentially aid governments and other organizations in improving their domain name management approaches.

1.3 RESEARCH QUESTIONS

To further study the topic of domain name management at the Dutch government, a research question and three subquestions have been formulated. The main research question is:

RQ: How should the Dutch government implement domain name management?

This thesis aims to provide concrete recommendations on how the Dutch government should address the problem of domain name management. In answering this question, a focus is laid on cyber security risks. The research will be based on both current literature and a review of the existing domain name management implementation.

In aid to answering the main research question, the following subquestions have been formulated:

SRQ 1: What are the possible security risks related to domain names?

Domain name management becomes a relevant problem when it can address security risks related to domain names. State-of-the-art literature will be reviewed to identify possible security risks related to domain names. The causes and potential impact of all identified security risks will be discussed. It is also addressed if and how proper domain name management can aid in mitigating these security risks. Where necessary, risks are illustrated with real-life examples in media outlets, if sufficient academic literature is missing.

SRQ 2: What is the policy theory of the Dutch central government concerning domain name management?

To provide guidelines on how domain name management should be implemented, it is key to observe the current approach regarding domain name management. For this, the policy theory behind existing policies will be reconstructed. A policy theory outlines the intent and thought process behind the policies, instead of merely evaluating policy outcomes. This rationale can be used to assess if and to what extent the policy mitigates the security risks that have been identified.

SRQ 3: How is domain name management executed within the Dutch central government in practice?

The current policies on domain name management are already executed in practice by the Dutch central government. To assess the performance of these policies, their policy outcomes have to be evaluated. To this extent, technical methods are used to obtain a governmental domain name overview. Out of the identified shortcomings in the current approach, common risks will be extracted. These risks are used in the recommendations on how the current policies could be improved in the future.

1.4 STRUCTURE

The remainder of this thesis is structured as follows. In chapter 2, a background study into the technical context of domain names and the Domain

Name System is provided. Then, the possible security risks involving domain names are researched in chapter 3. A review of current policies regarding domain name management by the Dutch central government is conducted in chapter 4. The effectiveness of this policy is measured, by analyzing policy outcomes in practice in chapter 5. Finally, the main conclusions, recommendations and discussions are given in chapter 6, followed by the bibliography of sources on which this thesis was based.

2 BACKGROUND ON DNS

2.1 INTRODUCTION

In this chapter, a background study into domain names and the DNS is conducted. Understanding these concepts is key to identifying the main incentives and challenges for domain name management. Next to a description of the operation of DNS, works are included that describe alternatives for or extensions to the DNS, which might be part of the specification in the future.

2.2 METHODOLOGY

The methodology for the selection of literature is as follows. For domain names and the DNS, the Internet Standards defining these concepts were studied. These Internet Standards are described in Requests For Comments (RFCs). RFCs play a fundamental part in Internet governance as they are the vessel to propose, discuss and eventually accept new Internet Standards. The processes around RFCs are supervised by the Internet Engineering Task Force (IETF) [44]. The format and procedures around RFCs are different from regular academic literature, yet new Internet Standards are not accepted before experts from both academic and industrial backgrounds have had the opportunity to review and improve the proposals. These strict and transparent procedures make the documents suitable sources for this background review. Technical implementations in practice might differentiate from the definitions in the RFCs, yet this cannot lead to significant functional differences between DNS implementations. Table a.1 in Appendix a contains an overview of RFCs related to DNS that contributed to the specification and the description in this chapter.

The background review is extended with various works of academic literature, which were selected from the academic search engines Scopus and Google Scholar. The used keywords were ‘DNS’, ‘domain names’, ‘domain name system’, combined with ‘alternatives’, ‘extensions’ and a combination thereof. Works were selected that introduce the technical concepts or discuss relevant extensions and alternatives for the DNS. Any relevant works cited by these papers were selected as well. Excluded are works on brand protection and market analyses, dated articles on potential DNS alternatives that did not result in any traction by the Internet community, and technical papers that address technical fundaments outside the scope of this thesis. Lastly, works that specifically address security risks related to domain names, will be discussed in chapter 3.

2.3 RESULTS

The concepts of domain names and their overarching DNS find their origin in two accepted Internet standards, RFC 1034 and RFC 1035 [62, 63]. These two standards have been in place since 1987, having received only minor updates through newer accepted RFCs. This means that the core concepts of DNS and domain names have not significantly changed over time.

In this section, firstly the technical description of domain names will be continued, as was stated in the introduction. Then, the focus is laid on the topic of DNS in the second subsection.

2.3.1 *Domain names*

In the introduction, 'example.com' was used as an illustration of the concept of domain names. It was explained, that this domain name consists of the domain name label 'example' and top-level domain 'com'. In its current form, the domain name is a partially qualified domain name (PQDN) as well as an *apex, bare or naked domain name*. The domain name is partially qualified because a fully qualified domain name (FQDN) should end with a trailing dot ('example.com.') to indicate the so-called empty 'root' label. The trailing dot is usually omitted in user interfaces to save typing [62], however, it should be taken into consideration that software may accept both PQDNs and FQDNs and may treat them differently.

The example domain name is apex, bare or naked, since it does not contain a *subdomain*. A subdomain implies the use of an additional sublabel in front of a domain name. For instance, in the domain name 'ex2.example.com', 'ex2' is a second-level subdomain label for the domain name 'example.com'. It is also possible to have higher-order subdomains, e.g. a third-level subdomain 'ex3' in the previous example results in 'ex3.ex2.example.com'. Subdomains are hierarchically subordinate to their higher-order domain names, yet may point towards different services and websites. A common subdomain label is 'www', which was first introduced as the subdomain where websites could be presented on the World Wide Web. The original idea was that different services would operate behind different subdomains. However, due to the dominance of the World Wide Web over other services, it is now common to serve websites also on the apex domain.

2.3.1.1 *Top-level domain names and ICANN*

TLD registrations and the application process are governed by the Internet Corporation for Assigned Names and Numbers (ICANN). Various types of TLDs exist, whereas the most prominent are country code top-level domain names (ccTLDs) and generic top-level domain names (gTLDs). ccTLDs consist of two characters and are allocated to regions, countries, and territories that are represented by a country code. If a country ceases to exist, the accompanying ccTLD is also discontinued, for example, the ccTLD '.an' of the Netherlands Antilles was removed in 2015, five years after the country was dissolved [99]. A second example is the top-level domain of Tuvalu, which is a small island country in the Pacific Ocean. Since its country code is 'tv', the TLD is frequently used in the context of television broadcasts around the world. However, the country is under threat of permanent flooding due to climate change, which may cause the country to disappear into the ocean. It has been announced that the TLD could be dissolved in this case [31].

gTLDs are top-level domains of at least three characters that serve various clients and purposes. In the early years of the internet, only a few gTLDs were assigned, serving specific purposes, for example, '.com' for companies and '.edu' for educational institutions. In the last decade, gTLD registration has been opened for custom labels and since then over 1000 gTLDs were registered for trademarks, cities, popular keywords, et cetera. Registration of a gTLD is subject to significant license fees, as well as high-standard technological requirements [41].

2.3.1.2 *Domain name registrations*

Within a top-level domain, domain names can be registered. Every TLD is assigned to a responsible organization, which is known as the *registry* of that TLD. Registries may provide licenses to *registrars*, allowing them to register domain names within that specific TLD. Clients of registrars can be both companies and individuals that want to register domain names. These clients are the final domain name owners, or *registrants*. Registries may set specific requirements for registration, for instance requiring citizenship of the country. They can also restrict their TLD for public registration entirely.

The registry maintains a database of the domain name registrations along with information about the registrant, often including contact details. Most TLDs make it possible to access part of this registration information, which allows end-users to verify who owns a specific domain name, and how to contact them. This information is usually accessed through the WHOIS protocol, which represents the information in a human-readable fashion. An upcoming improvement of the WHOIS protocol is the Registration Data Access Protocol (RDAP), which defines a computer-readable format for domain name registration information, among others. A recent study [53] has shown that the amount of information disclosed through these protocols has been reduced since the European privacy regulation (GDPR) became effective in 2018. This means that especially for domain names registered by individuals, personally identifiable information can no longer be retrieved through these protocols.

Certain TLDs only accept registrations under predefined domain names, instead of directly within the TLD. These domain extensions are also known as second-level domain names or public suffixes. Domain names registered by commercial entities in the United Kingdom, for instance, usually end with '.co.uk', instead of with '.uk'. The Mozilla Foundation maintains a public suffix list as a community resource[1]. In this case, '.co.uk' acts as the de facto top-level domain, even though technically 'co' is a domain name within the

[1] https://publicsuffix.org/

'uk' TLD. The list is for instance used by internet browsers to apply extra security and privacy measures to public suffixes, which are also applied to regular TLDs. In recent years, the registry of '.uk' also started accepting registrations directly within the ccTLD.

In case of termination of a domain name registration, the domain name goes into a grace period, or quarantine. The length of this period is to be determined by the registry. During this period, only the previous registrant can request to restore the registration. Once the period has ended, the domain name becomes available for regular registration again.

2.3.1.3 *DNS Resource Records*

The pointing of a domain name to a specific service is done using DNS Resource Records. Resource Records are set for a specific domain or subdomain and come in various types. Every type serves a specific purpose in the Domain Name System (DNS). For instance, Resource Records of type 'A' point to IPv4 addresses of that domain name and type 'AAAA' to IPv6 addresses. Most Resource Records can be specified multiple times, which might serve useful for purposes of load-balancing or redundancy. The free format record type 'TXT' supports plaintext strings as values, which is leveraged by the specifications of several Internet Standards. Resource Records are defined with a Time-To-Live (TTL) value, allowing a user to cache values for a specified time and decreasing the load burden on the DNS. Records are also assigned a CLASS value, which is commonly 'IN' for the Internet. Listing 2.1 contains a snippet of the DNS zone of 'example.nl', where every row contains the domain name, TTL value, CLASS, type and lastly the record's value.

```
example.nl.    3600   IN   SOA    ex1.sidnlabs.nl. hostmaster.sidn.nl. 1085 [...]
example.nl.    3600   IN   CAA    0 iodef "mailto:abuse@sidn.nl"
example.nl.    3600   IN   CAA    0 issue "sectigo.com"
example.nl.    3600   IN   MX     0 .
example.nl.    3600   IN   TXT    "v=DKIM1; p="
example.nl.    3600   IN   TXT    "v=spf1 -all"
example.nl.    3600   IN   TXT    "You may use this [...] without prior consent."
example.nl.    3600   IN   AAAA   2a00:d78:0:712:94:198:159:35
example.nl.    3600   IN   A      94.198.159.35
example.nl.    3600   IN   NS     ex1.sidnlabs.nl.
example.nl.    3600   IN   NS     ex2.sidnlabs.nl.
```

Listing 2.1: Sample of the zone definition for example.nl

2.3.1.4 *Internationalized Domain Names*

Domain names are represented in the American Standard Code for Information Interchange (ASCII) character encoding. This encoding is limited to the use of Latin-script letters in the English alphabet and Arabic numerals, and thus does not support the many alphabets used in the world, nor any diacritics. A more extensive character encoding is Unicode, which incorporates a wide range of scripts worldwide, as well as symbols and emoticons or emojis. To be able to use Unicode in the ASCII-encoded domain names, an additional standard was defined in RFC 3492 named Punycode [11]. Punycode provides the possibility to represent Unicode strings in ASCII, thus allowing registration of domain names in local scripts. A domain name using Punycode is called an Internationalized Domain Name (IDN).

While Punycode, in theory, supports the entire Unicode character encoding, the usage of symbols in domain names is limited by RFC 5892 through the IDNA2008 protocol [21]. For instance, this protocol restricts the usage of emojis in domain names [42]. In practice, however, various top-level domain names support Punycode while not respecting the IDNA2008 protocol, making it possible to register domain names that contain emoticons. An example of this is the ccTLD '.ws' of Samoa [17].

In 2010, ICANN enabled the first IDN top-level domain names, first giving priority to translations of existing ccTLDs. Since 2012, it is also possible to register new gTLDs that make use of IDN characters at ICANN [40].

2.3.2 *Domain Name System*

The Domain Name System, or DNS, is the core system that translates a domain name into the corresponding Resource Records, most significantly IP addresses. If a user wants to use a domain name to access a website, for instance, the user issues a DNS request to retrieve this IP address. Once the IP address is retrieved, the user sets up a connection to access the website. Servers that hold Resource Records for the DNS are called *name servers*.

In simplified form, this resolution from domain name to IP address is done as follows. Consider a user that requests the IP address for the domain 'example.com'. The DNS request of the user is issued to a *DNS resolver*, which is a server dedicated to handling DNS requests. For instance, it is common for Internet Service Providers (ISPs) to provide a DNS resolver. The user's Internet modem is configured with the IP address of their resolver of choice. The DNS resolver, in turn, has stored the IP addresses of one or multiple *root name servers*. The root name servers are the top level of the hierarchical DNS. The resolver relays the user's query to one of the root name servers. The root name server does not know the IP address of 'example.com', but *does* know the name server of the top-level domain 'com', and returns the IP address of this name server. The DNS resolver relays the query to the 'com' name server, who still does not know the IP address, but does know the domain of the name server that knows it, and returns that to the resolver. This process is repeated until the name server that does know the IP address of 'example.com' is queried. This name server is called an *authoritative name server* for 'example.com'. The authoritative name server returns the IP address for 'example.com' to the DNS resolver, which in turn returns this to the user. The DNS query is completed. In practice, DNS resolvers implement a caching mechanism that allows them to store DNS responses for a specified time. This reduces the number of queries and thus the load of the DNS in general.

The decentralized design ensures that every name server only has to hold a portion of the information that the DNS contains in total. Every name server has stored its information in a so-called *zone file*, and the root name servers of the hierarchical structure are defined in the *root zone file* [38]. The root zone file is managed by Internet Assigned Numbers Authority (IANA), a department of ICANN.

2.3.2.1 *DNSSEC*

As mentioned in the introduction, literature has identified significant security flaws in the original design of DNS. The original design was vulnerable to spoofing attacks that involved the injection of malicious Resource Records. In turn, this could be abused to intercept website connections and sent e-mails. These flaws have since then been addressed through improvements in wide-used DNS implementations, as well as through a new Internet Standard called DNSSEC [5]. This standard is currently defined in RFC 4033 [4]. DNSSEC allows registrants of domain names to digitally sign the information included in the Resource Records of those domain names. To this extent, it uses a public-private key infrastructure, in which every layer of the DNS has to broadcast a signature to its parent layer. The result of this effort is a top-down chain of trust, which means that DNS resolvers can verify the content of received Resource Records. The top-down approach of DNSSEC implies that it can only be enabled for individual domain names for which the TLD is signed. In June 2022, 1372 out of 1487 existing TLDs were signed using DNSSEC [43].

It should be noted that DNSSEC only attributes to the security of the DNS if two conditions are met. Firstly, the registrant of a domain name has to add a signature to its domain name, and all hierarchical layers above it have to be signed as well. This includes the TLD and the root node. Secondly, the DNS resolver of the end-user must verify if the signatures match the content of the received Resource Records, and reject the response otherwise. As DNSSEC is far from fully deployed in the DNS, the exploitation of several vulnerabilities remains theoretically possible [79].

2.3.2.2 *DNS Alternatives*

The core concepts of the DNS have not been fundamentally altered since its introduction in the 1980s. There have been several works that discuss the core design of the DNS and propose alterations or alternatives to the current DNS. These proposed alternatives are discussed below and it is explained in what way they differ from the current implementation.

DNS over TLS (DoT) [14, 35] is an adapted version of the current DNS implementation, which sends DNS traffic over an encrypted channel. A significant downside of regular DNS, is that its traffic is by default unencrypted. This means that even when an end-user only uses the encrypted Hypertext Transfer Protocol Secure (HTTPS) protocol to access websites, the DNS can still reveal which domain names are requested by the user [95]. When using DoT, the content of the DNS traffic can no longer be viewed by external parties.

Another alternative that aims for improved client privacy is DNS over HTTPS (DoH) [8, 32]. DoT and DoH are similar, yet have been developed independently. DoH makes use of the Hypertext Transfer Protocol Secure (HTTPS) protocol as its 'vessel' over the internet. HTTPS also relies on TLS for encryption, just like DoT. The difference between the two alternatives is that in DoH, DNS traffic will be indistinguishable from regular HTTPS requests. DoT, on the other hand, is transmitted over a different port-channel and thus could be detected and filtered from other traffic. However, DoT operates with less overhead as it does not have the overhead of the HTTPS protocol. In any case, if implemented correctly, both alternatives would be effective solutions to the information leakage of the current DNS. Both alternatives have been accepted as Proposed Standards by the IETF community and thus can be implemented in practice.

More recent alternatives that have been considered, are variants of leveraging DNS on blockchain technology [51, 60]. Blockchain is considered a potential solution for future DNS implementations, due to its decentralized nature without the need for trust in a central institution. In 'classic' DNS, there are various stakeholders that require trust, for example, the registries of every TLD, the ICANN as distributor of TLDs, and the IANA as the maintainer of the root zone file. Another use case could be leveraging the immutability of the blockchain to prevent tampering with Resource Record values. End-users could verify these values directly on the blockchain ledger, without having to verify signatures as with DNSSEC. These alternatives have only been theorized in the last few years, and are far from being adopted as standards. It is, therefore, unlikely that blockchain technology will be implemented in the next few years in the context of DNS, yet it might play a larger role in the future.

2.4 CONCLUSION AND DISCUSSION

In this chapter, the technological concepts of domain names and the DNS have been introduced and explained. Knowledge of the technological concepts is useful to perceive the complexity of domain name management and can be used to design solutions that conform to the technical properties of DNS.

It has been identified that alterations and alternatives to DNS have been proposed or are currently in development. This means that the problem of domain name management as discussed in this thesis, might be of different concern in these alternative systems.

However, while alternatives may make use of different technologies or protocols, none of the alternatives appear to propose a situation where domain name management is no longer required. It is expected that the problem remains relevant, even if another implementation would become the new standard.

3 SECURITY RISKS IN DOMAIN NAMES

3.1 INTRODUCTION

Cyber security risks and incidents are actual topics in recent years. Media outlets regularly publish articles about cyber attacks and data leaks, and security researchers report about newly found vulnerabilities in existing technologies. Domain names and the DNS are no exception to this, regularly reported either as the cause of a vulnerability or the target of a cyber attack. Domain name management could play a role in the mitigation of security risks. Having an overview of the domain name portfolio allows for continuous security monitoring and scanning of all connected systems. Irregularities and misconfigurations within these systems or the DNS could be more easily detected, decreasing the chance of successful cyber attacks.

Before determining how domain name management could be employed for this end, it should be determined which risks are related to domain names and the DNS. To this means, a review of relevant literature is conducted. From this review, possible risks related to DNS and domain names are abstracted. The risks are classified into different categories and the potential impact of every risk is discussed. This will be used to answer the following sub research question: *SRQ 1: What are the possible security risks related to domain names?*.

3.2 METHODOLOGY

For the identification of risks, a review is conducted of academic literature that discusses risks related to domain names and DNS. A literature review is deemed sufficient, since the DNS is a mature technology that has received sufficient attention from academics in recent years. It is deemed unlikely that other methods, like expert interviews, would expose other risks than those already discussed in the literature. Nevertheless, it remains possible that there reside other risks in the DNS that are currently undiscovered.

The documents for the review were selected from Scopus and Google Scholar using multiple query words, 'DNS', 'risk', 'vulnerability', 'weakness', 'domain name', and combinations thereof. Several documents detailing security risks were already identified during the review conducted in chapter 2. An additional search is conducted on Google Search, to identify relevant news articles in media outlets on this topic. Lastly, several of the RFCs on DNS contain security considerations or propose security standards. Where relevant, those RFCs are included from Table a.1 in the review.

An initial observation of the obtained literature is that current research does not distinguish between risks that are affecting domain name holders and risks that are affecting the DNS. In other words, not all risks identified in literature can be solved by, nor are directly relevant to, domain name holders. Since the issue of domain name management relates to the perspective of the domain name holders, only those risks that affect them should be included. It is expanded upon which risks are and which are not included. Furthermore, several real-life examples observed in media outlets are included that illustrate the identified risks.

The papers that discuss risks for domain name holders, are found to be categorizable into three risk areas. These categories are based on the cause of the associated risks, and the actor that is directly affected. An overview of this categorization is shown in Table 3.1. The three categories will be shortly addressed, after which the excluded risks are discussed.

The first category is *(Sub)domain takeovers*. This categorizes risks that would gain an adversary (partial) control over a domain name held by the victim. The cause can both be of technical origin [74], improper domain name management [87, 93] or excessive trust in external domain names [37]. For all causes holds, that the origin of these risks lies with the domain name owner, which made a mistake in its domain name management.

Category	Vulnerable actor	Attack vector
(1) (Sub)domain takeovers	Domain owner	Owner lets domain name expire or configured outdated records
(2) Impersonation and typosquatting	Domain users	User is fooled to use malicious domain name
(3) Security Standards non-compliance	Domain owner	Owner does not implement optional yet security-wise fundamental standards

Table 3.1: Categorization criteria for the three risk categories related to domain names from the perspective of domain name owners

The second category is *Impersonation and typosquatting*. Risks in this category all relate to an adversary impersonating a victim through the use of domain names. The main distinction with the first category is that the domain names of the victim are not directly targeted, but instead the users that make use of the victim's services. Impersonation can be done through the reliance on user error when typing domain names [47, 64, 68, 86], abusing optically indistinguishable characters in Internationalized Domain Names [88], or the use of other generic domain names that do or do not contain a reference to the victim [47]. While the cause of the risks lies with domain name users, this also affects domain name owners. Users are under the impression they interact with a genuine website, and thus the reputation of the domain name holder is at stake.

The third and final category concerns *Security Standards non-compliance*. This entails non-compliance to any security standard that is addressed at the domain name level. Examples are the use of domain names in Transport Layer Security (TLS) certificates for encrypted communication [23, 82] and security standards defined using DNS Resource Records [4, 12, 30, 48, 49, 66]. This category differs from the first category, because they are not caused by an owner's mistake, but instead by the lack of implementation of an optional or additional security mechanism. Where it could be argued that a domain name owner should be blamed for not adhering to current standards, its cause is fundamentally not stooled on a configuration mistake as with the first category.

Then, several risks affect the DNS technology, but do not affect or cannot be addressed by individual domain name owners. These are not included in this risk review, since it is focused on domain name owners. These include risks affecting the availability, where the DNS is abused for Distributed Denial-of-Service (DDoS) attacks using DNS amplification [3, 55, 56, 83]. Whereas DDoS attacks could be targeted toward individual domain name owners, mitigations against this kind of attack have to be implemented on the DNS or IP routing levels. Another out-of-scope security risk is the occurrence of junk records in zone files [85]. These kinds of records indicate misconfigurations at the operator of the zone file, rather than the domain name owners. The last type of risk excluded in this chapter is anything related to privacy concerns for end users, where the end user is not conforming to the most recent security standards [83]. In those scenarios, end users could remain vulnerable to the risks in category 3 even when the domain name owner has implemented all relevant security standards.

The risks will also be discussed from an operational security perspective, to provide more insight into how these risks could be abused in practice.

3.3 RESULTS

The identified risks have been divided into three categories. These categories are separately discussed in the following subsections.

3.3.1 *(Sub)domain takeovers*

There is spoken of a domain or subdomain takeover, when an adversary can take over a domain name registration or can make an existing domain name registration connect to the adversary's service. There are various methodologies on how this could be achieved.

In case of an (unintended) expiration of a domain name registration, any adversary could register the domain name once the quarantine period has ended. This may have various implications. In the first place, if users of the previous registrant's website still assume the domain name is owned by that registrant, the user can be tricked into submitting personal details to the new website. Secondly, if the taken-over domain was usually receiving mail, it is possible that unaware end-users still send e-mails to this domain, which could result in data leaks [87]. Lastly, if e-mail addresses using the domain were used for authentication services on other domain names, it can be possible to take over these accounts, for instance by using the password reset functionality. This method has also resulted in data leaks in practice [93]. A final possibility is that the taken-over domain was used on other domain names, for instance for including media content or functional scripts. In this case, the adversary could inject arbitrary information into existing websites, possibly targeting the users of these other websites [37]. Webpages containing a hyperlink toward the domain name would now link to the adversary's website. Hyperlinks that stop working due to domain name expiration degrade the functional purpose of the website. This phenomenon is also known as *link rot* [59].

A more complex cause of domain name takeover is when the registrant's account at the accompanying registrar is compromised. This would give the adversary access to the tools to change the domain name's Resource Records, or even transfer it to a different registrar. In even more extreme situations, a registrar could be compromised in itself, which would give the adversary access to all domain names registered under that registrar.

Secondly, there is the possibility of domain and/or subdomain takeovers based on outdated Resource Records [74]. In this situation, the domain name registration is still held by the original registrant. However, the specified Resource Records of that domain name are pointing toward IP addresses or other domain names that are no longer under the control of the registrant. An adversary can take over these (sub)domains by taking control of the specified IP address or domain name, and start accepting traffic from the specific (sub)domain. This way, end-users could be tricked in a similar fashion as was observed with expired domain name takeovers.

The possible impact of (sub)domain takeovers depends on what the original use case of the domain name was, and how many users were dependent on the domain name. Takeovers pose a significant risk as long as the original domain owners do not notice the issue, and still assume they are in control of the domain name. Furthermore, end users cannot fully verify whether a domain name is taken over, since the domain name appears genuine and has not been altered from earlier genuine usage.

3.3.2 *Impersonation and typosquatting*

The second category of risks is composed of impersonation and typosquatting. In comparison to the preceding category, this category of risks is based on the human fallibility of the end-user. This implies that given that an end-user has sufficient technological knowledge, they would be able to differentiate impersonated use of domain names from the genuine organization that is being impersonated.

Impersonation refers to the usage of domain names by an adversary with the purpose to impersonate an existing website or organization. For example, the adversary could register the same domain name label under a different Top-level domain name. Furthermore, they could register a domain name that contains the impersonated organization's brand name together with generic keywords such as 'services' or 'mailing', which might not be recognized by the user as an impersonation attempt. This approach is known as combosquatting [47]. Even domain names that are unrelated to the impersonated organization could be used in impersonation attacks. The legitimacy of these domains could be increased by presenting a copy of the real organization's website or sending out e-mails using the organization's mail template.

Several techniques exist that rely on the misspelling of domain names by users. This is generally known as typosquatting [64, 86], where an adversary registers variants of the real organization's domain name, containing common spelling mistakes, like the change of one character to an adjacent character on the keyboard. Other techniques are: the abuse of phonetic variations, or soundsquatting/homophone-based squatting [47, 68], the abuse of bit-flip errors in computer programs, or bitsquatting [67], and the abuse of visually similar characters, or homograph-based squatting [33]. End-users that make these kinds of typing mistakes, would end up navigating to the adversary's domain name, or sending an e-mail to the misspelled domain name. An example of each of these types is included in Table 3.2.

Example domain	**Type of squatting attack**
utwente.nl	Original domain, no attack
utwentee.nl	Typosquatting
utwenve.nl	Bitsquatting
youtwente.nl	Soundsquatting/Homophone-based squatting
utwente-mailings.nl	Combosquatting
utwen1e.nl	Homograph-based squatting
xn–twente-okh.nl (utwente.nl)[1]	IDN-based homograph-based squatting

Table 3.2: Examples of possible squatting attacks on domain names [88]

A more advanced method of homograph-based squatting is abusing Internationalized Domain Names (IDNs). Various Unicode characters share the same optical appearance but are seen as different by computer interfaces [88]. For example, the Greek omicron, which is visually identical to the Latin letter 'o' in most fonts, could be used in this kind of attack. By swapping one or more characters of an existing domain name with an optically identical character, the adversary would be able to register a domain name that is optically identical to the impersonated organization's domain name.

The possible impact of impersonation attacks is highly comparable to those of domain name takeovers. The difference lies in the need for an end-user mistake, either by making a typing mistake when entering a domain name, or by not recognizing a domain name that is normally used by the

[1] The .nl TLD currently does not offer IDN domain name registration, making this a theoretical example.

organization it impersonates. Organizations can implement preventive measures for impersonation by using monitoring services that detect the usage of their brand names in domain name registrations, or by preventive registration of common typing mistakes of their domain names.

3.3.3 *Security Standards non-compliance*

Various accepted Internet security standards are directly related to, or implemented at the domain name-level. Non-compliance to these standards does not imply that the domain name is directly vulnerable to cyber attacks. However, it may enable adversaries to exploit the known cyber risks that these standards aim to address. The main security standards are discussed, as well as the possible disadvantageous outcomes of non-compliance.

First, the role of domain names in TLS certificates is regarded [82]. TLS certificates are digital certificates that can be used to set up secure internet browser connections and transactions. Certificates are provided by Certificate Authorities (CAs), which are organizations provided with the trust to issue these certificates according to high security standards. Certificates can be issued for domain names, allowing for secure communications between users and the domain name. One common use is through the HTTPS protocol [23], in contrast to the insecure HTTP protocol. Domain names that do not provide a TLS certificate do not support secure connections, which could expose the contents of the information exchanged between the user and the domain name. If a domain name presents a certificate that is issued for a different domain name, or a certificate that is not provided by a recognized CA, browsers or connection software usually warns the user of an insecure connection.

TLS certificate compliance is reasonably high on the Internet, and the usage of HTTPS instead of its insecure alternative is becoming more common. Using the CAs as trust anchors for TLS certificates has been going well in general, with incidental exceptions like the DigiNotar incident [61].

Then, the security standards that are to be adopted through and configured in DNS Resource Records are considered. Firstly, the earlier discussed DNSSEC extension to DNS [4]. This standard addresses security issues in the fundamental design of DNS. Non-compliance to the standard can theoretically make the users of these domain names vulnerable to the known vulnerabilities in the DNS design, most significantly to cache poisoning attacks [2]. Even though other measures have been taken to make these attacks less likely to be successful in practice, new attack vectors are still being discovered in the present day [80]. Adoption of the DNSSEC standard is possible as long as the higher-order parts of the domain name (in most cases the TLD and the root zone) have been signed with a DNSSEC signature.

Then, there are three security standards related to sending e-mails from a domain name, which are SPF, DKIM, and DMARC [12, 48, 49, 66]. The standards restrict which servers are allowed to send e-mail on its behalf, and prevent tampering with the e-mail once it has been sent. Implementing these standards prevents the domain names from being used in unauthorized e-mail sending, since receiving mailboxes can verify if the sender was authorized to use a specific domain name. The standards are required since the e-mail standard by default does not implement any authorization methods. Thus, in theory, e-mail can be sent from any arbitrary domain name not implementing these standards. It has been observed that non-compliance increases the chance of usage of those domain names in SPAM attacks [27].

Lastly, there is the Certification Authority Authorization (CAA) standard [30]. This standard can be used to specify which Certificate Authorities are authorized to issue TLS certificates for the specific domain name. By authorizing a subset of CAs, the chances of misissueing certificates decrease, and the domain name owner will be notified in case of an invalid certificate request. It should be noted that the standard does not cover situations where a CA in itself would become malicious, as in the case of DigiNotar [61], since it could simply ignore the value of the CAA record.

The possible impact of non-compliance to any of the preceding security standards depends on which of the standards are not complied with. Every standard addresses different security risks and although it is best to comply with all accepted standards, non-compliance will not always directly make a domain name vulnerable to cyber attacks. A benefit of security standards is that the level of compliance can be relatively easily measured, for instance by checking the presence of the required DNS Resource Records. This makes it easier to assess the risks that a domain name may still be exposed to in terms of security standard non-compliance.

3.4 OPERATIONAL SECURITY PERSPECTIVE

The identified risk categories illustrate ways in which domain names may cause cyber security issues. However, they cannot be directly translated into real-world attack scenarios by criminals. In this section, it is aimed to provide insight into which ways these risks could be operationalized in cyber attacks.

The principal scenario where all three risk categories could be abused is in social engineering attacks. Social engineering is an attack methodology that aims to abuse human fallibility. Domain names can play a large role in this methodology, as adversaries could for instance host a malicious website on either taken-over (category 1) or impersonating domain names (category 2) that are similar to a genuine website. Users could be tricked to provide sensitive information that may be used to gain further access to an organization's systems. Apart from websites, domain names could also be spoofed in e-mail phishing campaigns. Especially when domain names have not configured e-mail security standards (category 3), this is a likely scenario. In these cases, the abuse of domain names is purely for obtaining an entry point to underlying systems.

A secondary scenario that could lead to a successful cyberattack is when domain names are misconfigured (category 3) to the extent that they expose sensitive information. Misconfigurations could make underlying services more vulnerable to direct attacks, for instance when internal services are exposed to the internet by accident. Also in this scenario, domain names would act as an initial entry point for these attacks.

The last scenario was already discussed in category 1, when expired domain names are re-registered to intercept sensitive e-mails that may still be sent to old contacts. It is known that this attack scenario has been successfully executed in practice [93].

Summarizing, it can be said that the identified risks can be abused in practice. However, they are most likely to serve as an initial point of entry, after which other security issues of underlying systems need to be abused to make an attack a success. Nonetheless, it is still important to mitigate these risks, as they limit the points of entry an attacker may have.

3.5 conclusion and discussion

Having regarded the literature that identifies risks tied to domain names, an answer can be provided to the first subquestion: *SRQ 1: What are the possible security risks related to domain names?*. Three risk categories have been derived from the literature that directly relate to domain names, (1) (Sub)domain takeovers, (2) Impersonation and typosquatting, and (3) Security Standards non-compliance. Where necessary, these risk categories have been illustrated with real-life examples from articles by media outlets. The typology of risk categories will aid in further assessing the problem of domain name management, as it is now possible to test current domain name management approaches on their mitigation of these risks.

Discussing the results presented in this chapter, it has to be considered that the amount of academic literature on the subject is limited. This led to the inclusion of at least one work that is still in press at the time of writing. Furthermore, it could be that not all possible risks have been discussed in academic literature yet, and thus are missing out from the proposed typology. The investigation of articles in media outlets, however, has not shown risks that could not be attributed to one of the three categories. Lastly, it remains possible that undiscovered risks reside in the DNS technology that might affect domain name holders as well.

A second point of discussion is the typology of the third category, non-compliance with security standards. One could argue that it is not non-compliance that poses a risk, but instead, the various security vulnerabilities that those standards aim to address. However, it is argued that in the context of domain name management, it is more relevant to address risks from the perspective of the domain name owners. In this case, the risk for domain name owners lies in not implementing the latest security standards for their domain names. Risks that reside in the technical implementation of DNS, like amplification attacks, are more relevant for DNS operators and internet researchers instead.

DOMAIN NAME MANAGEMENT IN DUTCH CENTRAL GOVERNMENT

4.1 INTRODUCTION

In this chapter, a close look is taken into the policies concerning domain name management at the Dutch central government. The aim of this chapter is to understand the reasoning behind the current policies and determine whether the policies are expected to effectively mitigate the risks as identified in chapter 3. Consequently, this will provide an answer to the second sub research question: *SRQ 2: What is the policy theory of the Dutch central government concerning domain name management?*. The approach to analyzing the past and current policies is by abstracting the policy theory behind domain name management [34]. Since several developments are still ongoing in the policy field of domain name management, an outlook on possible future policies is included as well. Furthermore, it is aimed to visualize the current policy arrangement of domain name management at the Dutch government.

4.2 METHODOLOGY

4.2.1 *Policy theory*

The policy theory behind a certain policy outlines all expectations, assumptions and circumstances that contributed to the establishment of that policy [34]. This provides a more complete view of the policy process than just measuring policy outcomes, because the policy outcomes may significantly differ from those that were predicted and intended. This is conducted in a systematic manner, based on all the resources available that relate to the specific policy. A reconstruction of the policy theory can be used to understand why certain choices have been made, and in turn, be leveraged for proposing policy improvements or additions. Reconstructing policy theories is a common tool used by policy studies in various policy fields. It is a method of policy analysis, whereas the approach is generally client-oriented [94].

Searching for applications of policy theory in the field of cyber security and government, two recent works are identified. These were found on Scopus and Google Scholar using keywords 'cyber security', 'policy theory', 'government', their Dutch translations and combinations thereof. One work evaluates the policy process of the Dutch Cyber Security Agenda of the Dutch central government [9]. This agenda is a single document outlining all areas within cyber security on which the government is intending to improve. The policy theory focuses on the structure of the original agenda, but also on the outcomes after the agenda was executed. The second work is a reflection on the HackShield initiative, a regional collaboration between governments in the province of Noord-Holland, Netherlands. This initiative involves a cyber security game for children between 8 and 12 years, to increase cyber awareness among them. The work focuses on how the initiative was organized and how the plans were executed in practice. A thorough evaluation of concrete outcomes is not included and is left open for future research.

The setting of the two identified works indicates that the application of policy theory reconstruction can also be relevant for a study into domain name management at the Dutch central government. Since this thesis is not only focused on policy outcomes of domain name management, but also on providing guidelines on how domain name management should be conducted, this method is suitable. No alternative methods have been identified that could provide a similar complete picture of the policy process.

Therefore, the policy theory of domain name management will be constructed. Whereas a policy theory is regularly applied to a single bill or policy document, this reconstruction will entail a more abstract policy field instead. This is still feasible since the field is small and concerns a single technological topic. Furthermore, while domain name management may be addressed from different perspectives in various documents, the cyber security risks identified in chapter 3 can be relevant to each of these perspectives.

The research will be limited to policy documents that are available in the public domain. Following the systematic of [34], the reconstruction will focus on the policy arrangement around the subject of domain name management. To better understand the context in which the policymaking is done, historical policy documents will be included as well.

4.2.2 *Selection of domain name management policies*

To identify the policy documents required to reconstruct the policy theory, various governmental websites and general search engines have been queried. The used Dutch key words were: 'domein*', 'domeinnamen', 'Rijksoverheid', 'overheid', 'domeinnaambeleid', 'domeinnaamregister', 'domeinnaam', 'domeinnaambeheer', 'website' and combinations thereof. This resulted in the selection of 29 documents and web pages that contained policies or supporting documents related to domain names. These documents are categorized as follows (with the number of occurrences): Policy document (4), Guideline for domain name procedure (5), Register (3), News article concerning the topic (5), Letter to Parliament (3), Research report (5), Assessment framework (1), Formal decision (2), Auxiliary website (1). The policy documents were separatable into three time sets. These time sets will be discussed separately in the construction of policy theory.

The first set contains documents that contain historical policy approaches on domain names, which are no longer implemented at this point. They are still valuable for the reconstruction of the policy theory, since they give insight into how domain name management policies have developed over time and in what manners these policies were lacking.

This is followed by a set of current policy documents, representing the policies that are currently in place. These policies are oriented in three different policy fields: communication, archiving and security & compliance. Whereas overlap between the policies is observed, there are also clear distinctions in these policy approaches. These will be further discussed in the results section.

Lastly, there is a set of policy documents related to ongoing changes and policy research. The government is aware that current policies might contain weak spots, and an investigation is ongoing into improvements and alterations. This could give an outlook on future adaption to the domain name policy of the Dutch government.

4.2.3 *Goal-Means tree analysis*

Whereas the method of policy theory is applied for interpreting policy documents on domain name management, an additional method is applied for visualization of the policy field. For this, the Goal-Means Tree Analysis method [92] will be applied. This method originally emerged in the Dutch Commission for Policy analysis in the 1970s and could be used to visualize single policy texts based on their eventual goals as well as the means to reach them. The method was extended by Van de Graaf and Hoppe to support multiple policy texts in a single tree, renaming it to 'goals-mean interpretation' [91]. Vanhaeght supplies an English translation of this method.

A goal-means tree is constructed in seven steps. The first four steps concern extracting the required information from policy documents. This means defining the means, objectives or sub-goals and goals of the policy documents. In steps five to seven, the obtained information is structured in the tree structure including the existing links. The links are either defined by logic implication (o), determined by the author of the text (-) or explicitly mentioned in the policy documents. The tree is structured to flow from means, via objectives, to the final goals.

The constructed tree can easily show the relationships between various policy documents and policy goals. A limitation, however, is that the tree will display policy only as a plan, not as a process. Moreover, the means-end relationships are from the perspective of the policymaker only, and do not include the goals or means of other stakeholders.

4.3 RESULTS

4.3.1 *Historical context*

This chapter aims to describe the policies that relate to domain name management and the reasoning behind them. The historical context around this issue explains how the current situation came to exist, and is therefore key to understanding the policy theory surrounding domain names.

The first '.nl' domain name dates back to 1986, yet the real traction of the Internet and domain names started in the second half of the 1990s. Naturally, most governmental departments predate this and thus were confronted with an unknown and upcoming technology [57]. While the popularity and usage of the WWW increased towards the century change, the registration of commercial and personal domain names became available and gained traction. This did not go unnoticed by public officials and incumbents in government. In a written consultation in 2000, questions are raised about individuals that have registered domain names using the names of governmental bodies. The parliamentary committee states that all governments should have possession of websites using a unique domain name. The responsible ministers are explicitly asked whether they intend to make domain names an official communication and identification method by law. Furthermore, the ministers are urged to inform decentral governments to register a domain name as soon as possible. [90]

The minister responded that indeed governmental organizations should possess a domain name, but leave the responsibility to the individual departments. They state that the Stichting Internet Domeinregistratie Nederland (SIDN) reserved domain names according to a list of municipalities, so these could not be registered by others. However, the list turned out to be missing 25 municipalities, which led to the domain names of these municipalities being registered by third parties, even though domain names were only available to businesses and organizations at the time. [90]

Lastly, the minister identified in their response that the increasing number of governmental domain names is problematic for end users, since they cannot easily distinguish them from other domain names. He states that the general website 'overheid.nl' ('government.nl') would serve as a portal to all governmental domains. Furthermore, subdomains on this domain would be used to direct to all governmental bodies, in essence stating that 'overheid.nl' will perform as a public suffix. For unknown reasons, this was not implemented as such, although the 'overheid.nl' website still serves a portal role in routing to other governmental bodies[1].

In terms of policy theory, this historical background gives significant insight. For starters, the problem of domain name takeovers appears to be as old as the WWW itself, and governmental bodies were rushed by incumbents and public officials alike to register a representative domain name as soon as possible. Furthermore, it is striking that the registration of domain names at the time was left to individual governmental organizations, instead of choosing a centralized approach. This would automatically result in no central oversight over the domain name portfolio of the Dutch government.

It should be noted that the role that websites played in the early 2000s is very different from today. Then, websites provided mostly static informa-

[1] https://www.overheid.nl/

tion about the governmental body and would list contact information. The organizational importance of websites was of less importance, whereas now, most governmental bodies conduct a significant part of their service provision online. For illustration, Figure 4.1 contains an archived version of the 'overheid.nl' website in 2000, around the time of the written consultation.

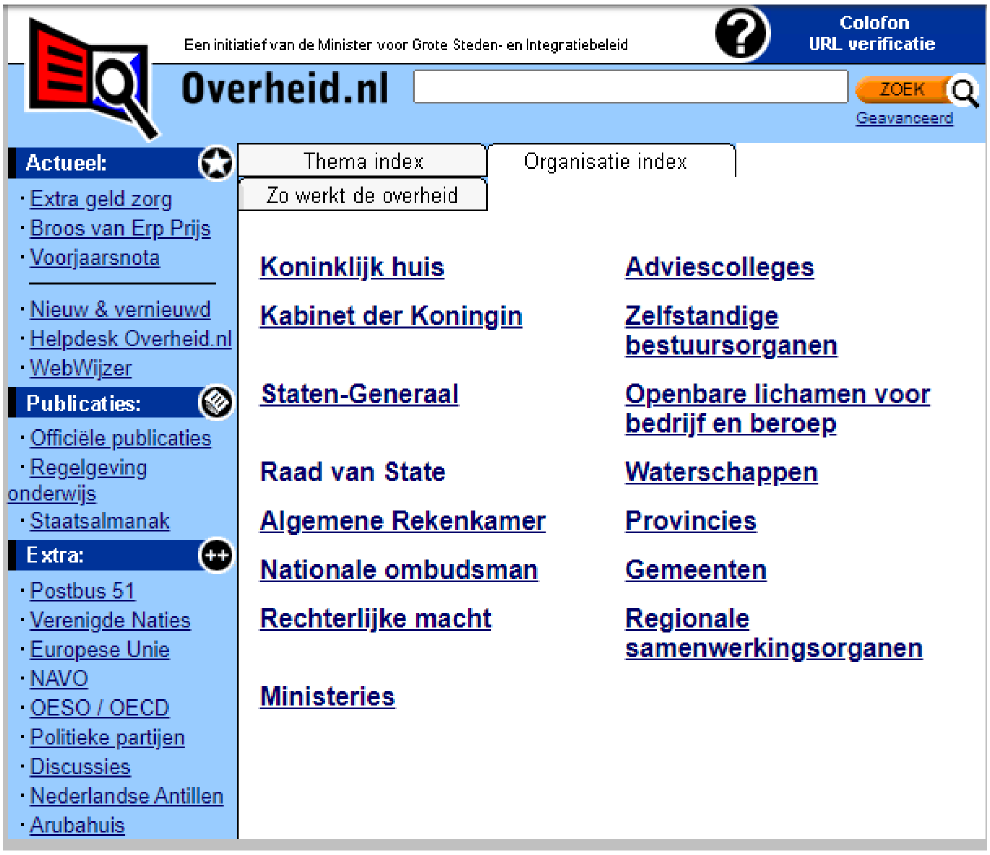

Een initiatief van de Minister voor Grote Steden- en Integratiebeleid

Figure 4.1: Screenshot of an archived version of the website overheid.nl in May 2000

4.3.2 *Current policy fields*

The research jumps in time toward policy documents that contribute to the situation of today. These are discussed separately per identified policy field.

Communication

Most of the identified policy documents were related to the field of communication. Here, communication indicates the interaction between citizens and the government. Policy documents focus on how communication using domain names can be conducted clearly and recognizably. It will be of no surprise that most of these policy documents are from the hand of the Service for Audience and Communication (*Dienst Publiek en Communicatie* (DPC)), which is part of the Ministry of General Affairs.

The principal policies regarding domain name management of the Dutch central government are located on the website regarding governmental communication [15]. The website contains the government's domain name policy which in part addresses domain name management [75]. This policy describes how domain names can be registered, for which purpose domain names may be registered and which TLDs may be used. The policy lists to be focused on the following four goals: (1) unambiguity of policy making and use, (2) transparency of policy making and use, (3) protection of the government's legal position, and (4) maintaining manageable costs. An earlier version of the domain name policy was found, which showed that the policy goals have been the same since at least 2014 [28].

The policy contents are set by the Information Council (*Voorlichtingsraad (VoRa)*), which is an advisory body that consists of the directors of communication of all ministries. The document refers to two official decisions which serve as the legal base of the domain name policy, a VoRa decision of 9 June 2011 and the ICBR (*Interdepartementale Commissie Bedrijfsvoering Rijksdienst*) decision of 15 March 2011. The exact contents of these decisions were unfortunately irretrievable.

According to the policy, all domain name registrations of the central Dutch government have to be handled by the DPC. The policy allows only registrations within the 'nl', 'eu' or 'com' TLDs, or within the 'aw' (Aruba), 'cw' (Curaçao) and 'sx' (Sint Maarten) TLDs when designated for one of the other countries within the Dutch Kingdom. Domain names are not rented, bought, or sold by the government. Domain names of dissolved websites are retained and pointed toward an archived version of the website. The policy does not mention, however, if and how these rules are enforced. When these rules are insufficiently enforced, the intended policy outcomes may not be achieved.

Then, the policy outlines three considerations of why this domain name policy is of benefit for communication purposes: (1) it is clear for citizens which communications originate from the central government, (2) the method remains overseeable and manageable for the government, and (3) it is prevented that already-taken domain names have to be obtained.

The domain name policy ends with references to two registers of domain names, which are the domain name register and the website register. The domain name register is a non-public register containing all domain names that are owned by the Dutch central government. Inquiry at the DPC revealed that this register contains over 9000 apex domain names, and over 16000 domain names when including higher-order domains[2].

The website register is a public subset of the domain name register, containing all domain names that serve a website of the Dutch central government [76]. The register is updated monthly and published as a spreadsheet. A screenshot can be found in Figure 4.2. This register contained 1718 domain names in November 2021 that are serving websites.

Another overarching communication policy is identified in the central government-wide assessment framework for online utilities set by the VoRa

[2] Personal Communication. Information obtained on 27 August 2021

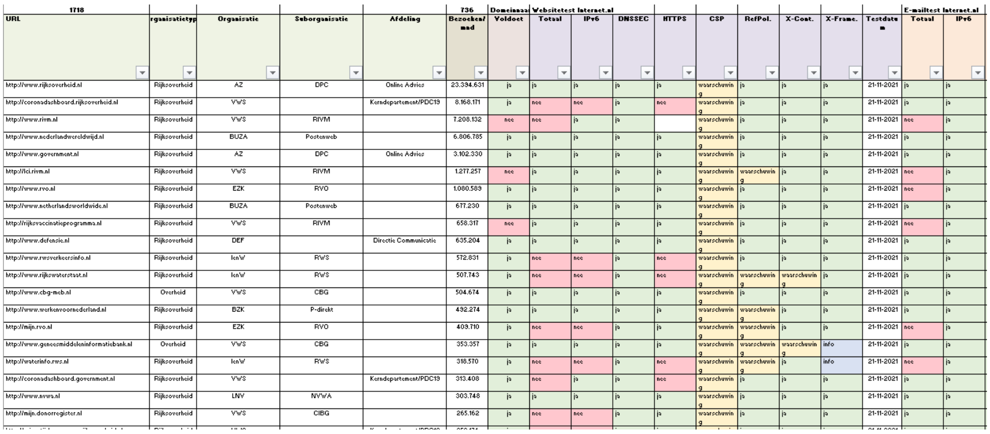

1718					736	Domeinnaa	Websitetest Internet.nl									E-mailtest Internet.nl	
URL	rganisatietyp	Organisatie	Suborganisatie	Afdeling	Bezoeken/mnd	Voldoet	Totaal	IPv6	DNSSEC	HTTPS	CSP	RefPol.	X-Cont.	X-Frame.	Testdatum	Totaal	IPv6
http://www.rijksoverheid.nl	Rijksoverheid	AZ	DPC	Online Advies	23.394.631	ja	ja	ja	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://coronadashboard.rijksoverheid.nl	Rijksoverheid	VWS		Kerndepartement/PDC19	8.168.171	ja	nee	nee	ja	nee	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://www.rivm.nl	Rijksoverheid	VWS	RIVM		7.208.132	nee	nee	ja	ja		waarschuwing	ja	ja	ja	21-11-2021	nee	ja
http://www.nederlandwereldwijd.nl	Rijksoverheid	BUZA	Postenweb		6.806.785	ja	ja	ja	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://www.government.nl	Rijksoverheid	AZ	DPC	Online Advies	3.102.330	ja	ja	ja	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://lci.rivm.nl	Rijksoverheid	VWS	RIVM		1.277.257	nee	ja	ja	ja	ja	waarschuwing	waarschuwing	ja	ja	21-11-2021	nee	ja
http://www.rvo.nl	Rijksoverheid	EZK	RVO		1.080.589	ja	ja	ja	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	nee	ja
http://www.netherlandsworldwide.nl	Rijksoverheid	BUZA	Postenweb		677.230	ja	ja	ja	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://rijksvaccinatieprogramma.nl	Rijksoverheid	VWS	RIVM		658.317	nee	ja	ja	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	nee	ja
http://www.defensie.nl	Rijksoverheid	DEF		Directie Communicatie	635.204	ja	ja	ja	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://www.rwsverkeersinfo.nl	Rijksoverheid	IenW	RWS		572.831	ja	nee	nee	ja	nee	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://www.rijkswaterstaat.nl	Rijksoverheid	IenW	RWS		507.743	ja	nee	nee	ja	nee	waarschuwing	waarschuwing	waarschuwing	ja	21-11-2021	ja	ja
http://www.cbg-meb.nl	Overheid	VWS	CBG		504.674	ja	ja	ja	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://www.werkenvoornederland.nl	Rijksoverheid	BZK	P-direkt		492.274	ja	ja	ja	ja	ja	waarschuwing	waarschuwing	ja	ja	21-11-2021	ja	ja
http://mijn.rvo.nl	Rijksoverheid	EZK	RVO		409.710	ja	nee	nee	ja	ja	waarschuwing	waarschuwing	ja	ja	21-11-2021	nee	ja
http://www.geneesmiddeleninformatiebank.nl	Overheid	VWS	CBG		353.357	ja	ja	ja	ja	ja	waarschuwing	waarschuwing	waarschuwing	info	21-11-2021	ja	ja
http://waterinfo.rws.nl	Rijksoverheid	IenW	RWS		318.570	ja	nee	nee	ja	nee	waarschuwing	waarschuwing	ja	info	21-11-2021	nee	ja
http://coronadashboard.government.nl	Rijksoverheid	VWS		Kerndepartement/PDC19	313.408	ja	nee	ja	ja	nee	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://www.nvwa.nl	Rijksoverheid	LNV	NVWA		303.748	ja	ja	ja	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	ja	ja
http://mijn.donorregister.nl	Rijksoverheid	VWS	CIBG		265.162	ja	nee	nee	ja	ja	waarschuwing	ja	ja	ja	21-11-2021	ja	ja

Figure 4.2: Screenshot of the governmental website register, containing domain names, their responsible organizations and their compliance to internet standards

in 2020 [77]. This framework contributes one significant rule of thumb in relation to domain names, namely that new projects or departments should retain as much as possible towards the existing communication channels. In other words, if it is not necessary to create a new website and domain name, stick to one that already exists.

Finally, a research report on the opinion of citizens concerning the governmental service provision was found [71]. In this research, 52% of the respondents said to prefer a single website to handle all interactions with the government. This could be interpreted as the use of a uniform domain extension, for instance, a TLD or public suffix.

In terms of policy theory, a key point of attention is the lack of attention to cyber security risks in the domain name policy. The policy does not point out that there are risks associated with domain names and domain name registrations. The only exception is the mentioning of the domain guarding service that is provided by SIDN [13].

Furthermore, it cannot be derived from the policy whether the policy is actively enforced, and what the consequences would be for not following this policy. The policy may be intended as a guideline, putting the responsibility on the individual departments.

Lastly, the policy listed that it should contribute to the recognizability of governmental domain names to citizens. This statement might be flawed, considering the domain name register contains 16000 different domain names. Of those, only a fraction is disclosed to citizens through the website register. Thus, it could be argued this policy goal is not reached in the current approach.

Archiving

The second policy field is on archiving. Governmental domain names are considered to be official publications, and therefore have to comply with the laws on archiving [75]. The Governmental Program for Sustainable Digital Information Management (*Rijksprogramma voor Duurzaam Digitale Informatiehuishouding (RDDI)*) is a department focusing on information management within the government. They have published a guidance document in 2021 on how governmental organizations can create insight into their own domain name portfolio [29]. They come to similar conclusions in terms of the complexity of this problem and provide several tools that can help in organizing one's domain name portfolio. This is part of their larger focus on the implementation of archiving regulations. The complete archiving of governmental websites, and providing public access to these archives, is part of this project.

In an additional document on web archiving from 2020 [73], they list the set of requirements to which all governmental websites should confirm. This provides a more detailed overview of all requirements, although no other requirements are mentioned than on the website of the DPC.

The Dutch National Archive has published a detailed guideline on how website archiving should be applied [65]. This includes a description of the exact requirement of the relevant archiving laws, as well as how an archiving methodology could be applied in practice. Adhering to the guideline itself is stated to be "voluntary but not without obligation".

The policy theory behind website archiving is straightforward since this is simply a fulfillment of the legal requirement to archive governmental publications. The documents provide a clear and complete overview of how archiving should be addressed.

Security & compliance

The third policy field is that of security & compliance. With compliance, adhering to mandatory standards is meant, both focused on security and otherwise.

Looking back to the website register discussed previously [76], it is supplemented with scoring results of a website test and an e-mail test on the compliance of mandatory Internet Standards, provided by Internet.nl [19].

Internet.nl is an initiative of the Dutch government and non-profit partners from the Dutch Internet community. The scans include both security standards and standards focused on accessibility. Furthermore, the register lists the responsible governmental department for every domain name, the number of monthly visitors and the platform used for a subset of the domain names. These scan reports give a clear and transparent view of the adoption of internet standards on governmental websites.

Another resource on compliance can be found on the website of the Dutch Standardization Forum (*Forum Standaardisatie*) [24]. This Forum is an independent authority, housed within the Ministry of the Interior and Kingdom Relations and is in charge of determining the use of which standards, including Internet standards, should be mandatory for or recommended to governments. This includes all layers of government. New standards can be submitted for consideration to this Forum by both government and commercial parties. Accepted standards are mandated through a "comply or explain" principle. In March 2020, a report was presented to parliament measuring the adoption of internet standards by governmental domain names over time [26]. It is concluded that the average adoption rate of internet web standards has increased to 94%, coming from 35% in 2015. Internet standards related to e-mail are behind at 81% adoption on average.

The Dutch Standardization Forum also has published a magazine on the management of domain names, in which they conclude that an effective approach for government-wide management is missing [25].

A dedicated website exists for the registration of accessibility certificates [52]. Accessibility certificates are given out by external auditors, rating the accessibility of a governmental website. If a website is developed in an accessible manner, the website can be accessed by people with disabilities, for instance visually impaired people using a Braille reader. The website contained 3710 accessibility certificates in August 2022, for both central and decentral governments.

Concerning security and domain names, a letter from March 2021 from the Dutch cabinet was sent to parliament concerning the state of information security [16]. This letter contained a section on domain names specifically, in which the uncontrolled growth of the number of governmental domains was described as a problem. Not only was it observed as a problem of recognizability, but it also affects the manageability of communication, information security, information management and IT management [16]. As a potential solution, the possibility of a uniform domain extension is mentioned again.

Looking at the policy theory of this policy area, there appears to be a disconnect with the other two policy areas. The only observed overlap is the inclusion of internet standard scan results in the governmental website register and the application of the domain guarding service by SIDN. This agrees with the earlier remark on the domain name policy, and how it did not contain any focus on cyber security risks and risk mitigation. The policies drafted in relation to domain name management are not done coherently, spanning over the three identified policy areas, but rather from the perspective of one area only.

4.3.3 *Goal-Means tree analysis*

The policy theory of the current domain name management policies as described in the preceding section has been mapped in a goal-means tree. To this end, the seven steps described in [92] have been applied to the discussed policy documents. In this phase, any goals that relate to financial savings have been excluded. The resulting goal-means tree is included in Appendix b. The composition of the tree clearly shows the three identified policy fields, as well as the overlap between them. Now, the tree composition will be further discussed.

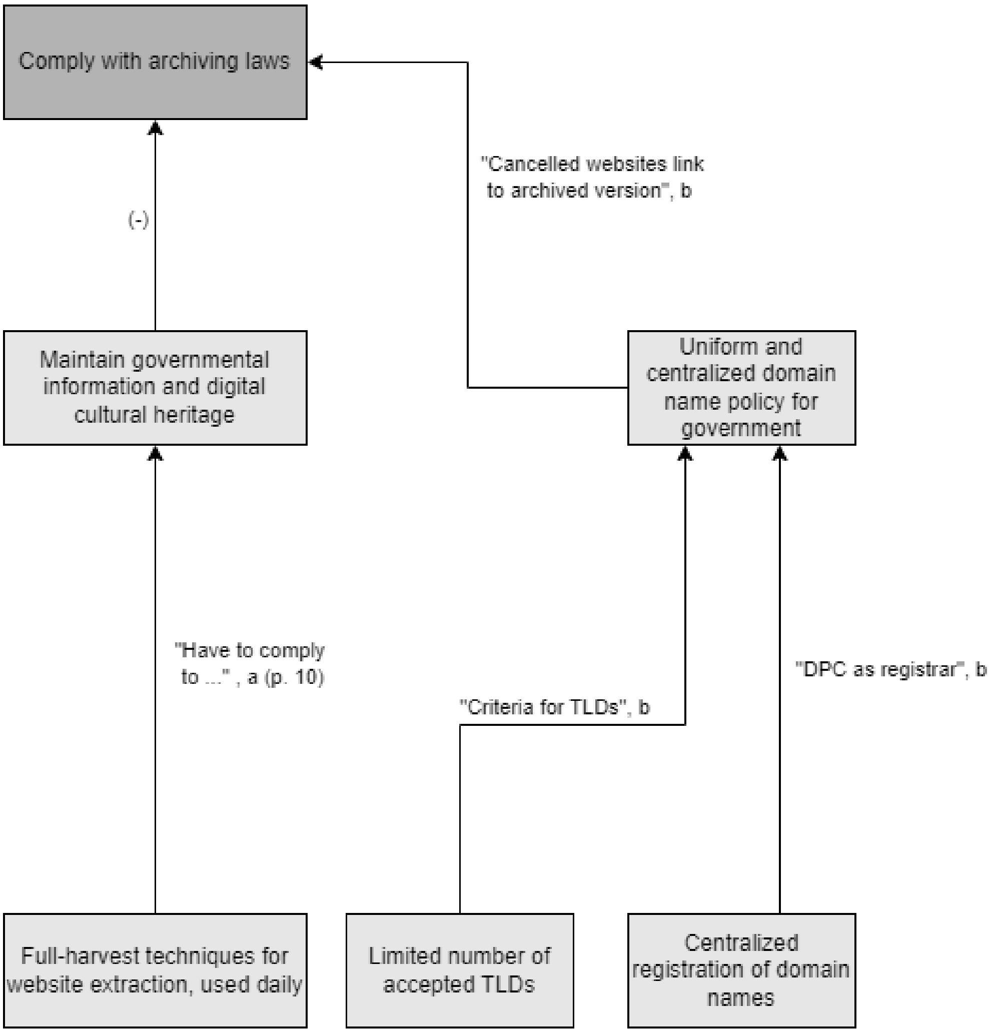
Comply with archiving laws
"Cancelled websites link to archived version", b
(-)
Maintain governmental information and digital cultural heritage
Uniform and centralized domain name policy for government
"Have to comply to ..." , a (p. 10)
"DPC as registrar", b
"Criteria for TLDs", b
Full-harvest techniques for website extraction, used daily
Limited number of accepted TLDs
Centralized registration of domain names

Figure 4.3: Policy field on archiving as a goal-means subtree of domain name management at the Dutch government

The tree is structured bottom-up. The bottom row represents the identified means that can be utilized to reach policy goals and objectives. The upper two rows containing dark gray boxes, contain the final policy goals. In between, the identified sub-goals or objectives can be found. The entire policy field is connected through explicit and implicit references. Explicit links contain the source from which the link is obtained. Implicit links, labeled with (-), are determined by the author based on the policy description, or represent the separation of means and objectives that were obtained from the same source. Lastly, two boxes have a dotted border. This is to indicate that this mean and objective are currently not in place, but are actively discussed in ongoing policy developments. These will be further discussed later in this chapter.

Looking at the goals, one can observe that the three policy fields, communication, archiving and security & compliance, are subordinate to the goal of domain name management. This is to illustrate that domain name management is an implicit goal of the policy documents, and confirms the earlier observations. In that sense, domain name management could also be seen as a (required) means to reach these goals. Separating the tree into the three policy fields, however, displays a clearer picture of the focus areas of individual policy documents. The content of the tree will be further discussed based on separated subtrees in Figures 4.3, 4.4 and 4.5. Overlap between the policy fields and their connections have been preserved in these subtrees.

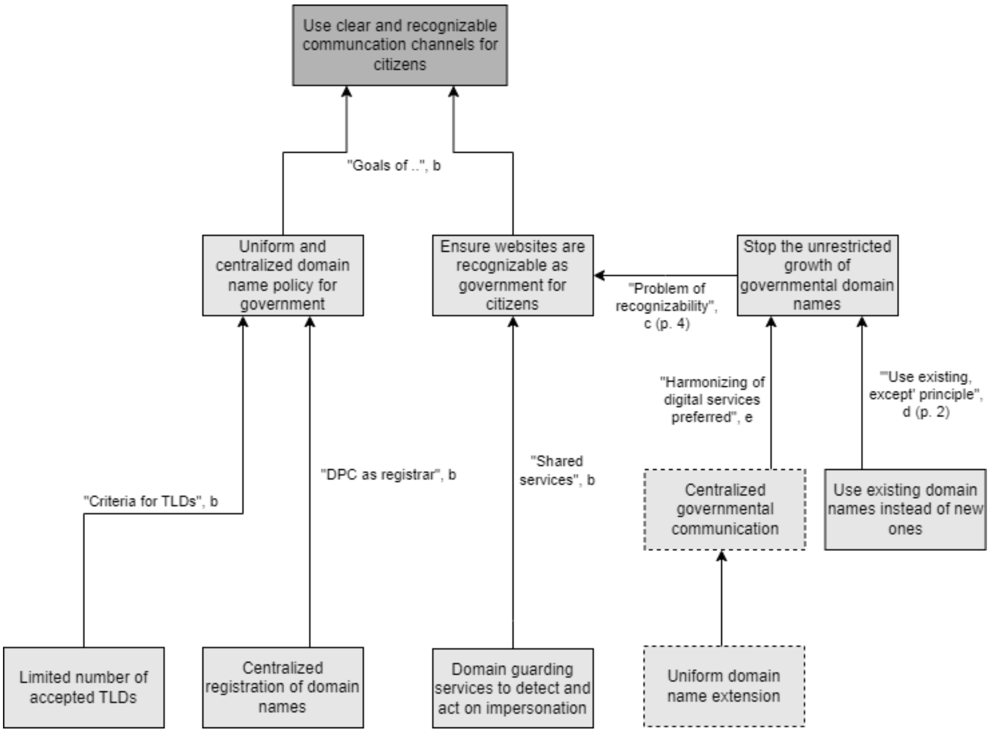
Use clear and recognizable communcation channels for citizens
"Goals of ..", b
Uniform and centralized domain name policy for government
Ensure websites are recognizable as government for citizens
Stop the unrestricted growth of governmental domain names
"Problem of recognizability", c (p. 4)
"Harmonizing of digital services preferred", e
"'Use existing, except' principle", d (p. 2)
"DPC as registrar", b
"Shared services", b
"Criteria for TLDs", b
Centralized governmental communication
Use existing domain names instead of new ones
Limited number of accepted TLDs
Centralized registration of domain names
Domain guarding services to detect and act on impersonation
Uniform domain name extension

Figure 4.4: Policy field on communication as a goal-means subtree of domain name management at the Dutch government

Identified means and goals related to archiving have been listed in Figure 4.3, with the final policy goal to comply with all relevant archiving laws. Domain names and government websites represent a small fraction of governmental communication that needs to comply with archiving laws. For websites, full-harvest techniques can be used as means to archive them. This means that every page on a website is 'harvested' and stored every day, allowing for viewing the website content for any date in the past. A connection can be observed with the policy field through the central domain name policy [75], which states that once websites are decommissioned, the website will link to its archived version instead.

The subtree on the policy field of communication in Figure 4.4 is more expansive. The final policy goal is the use of clear and recognizable communication channels toward citizens. In terms of domain names, it is aimed to achieve this by making domain name registrations as uniform as possible, while at the same time trying to restrict the number of domain names that are used. Combined with measures to detect malicious registrations that attempt to impersonate governmental domain names, this contributes to the general recognizability of governmental domain names. For future policy, more harmonization between digital services is preferred. A technical measure that is proposed for this is the use of a uniform domain name extension.

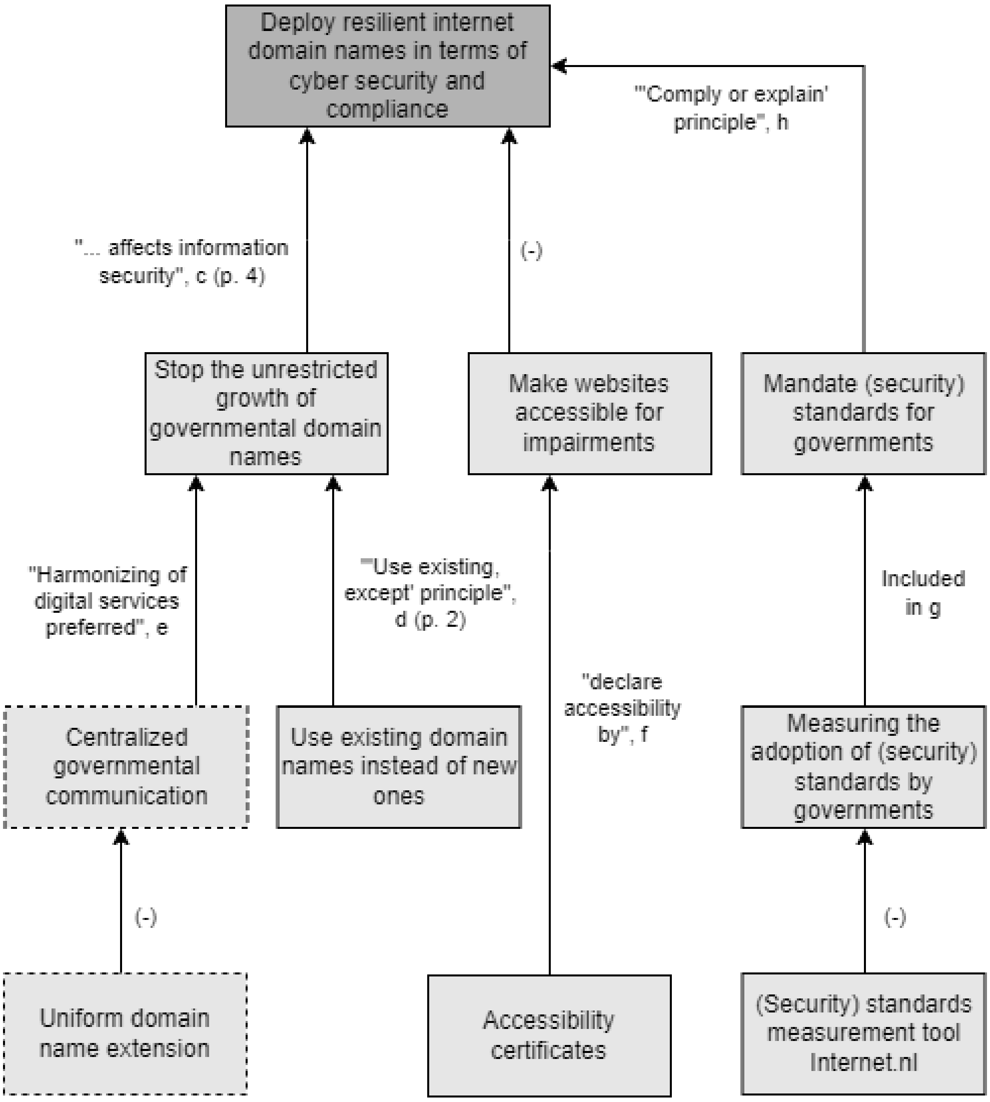
Deploy resilient internet domain names in terms of cyber security and compliance
"'Comply or explain' principle", h
"... affects information security", c (p. 4)
(-)
Stop the unrestricted growth of governmental domain names
Make websites accessible for impairments
Mandate (security) standards for governments
"Harmonizing of digital services preferred", e
"'Use existing, except' principle", d (p. 2)
Included in g
"declare accessibility by", f
Centralized governmental communication
Use existing domain names instead of new ones
Measuring the adoption of (security) standards by governments
(-)
(-)
Uniform domain name extension
Accessibility certificates
(Security) standards measurement tool Internet.nl

Figure 4.5: Policy field on security & compliance as a goal-means subtree of domain name management at the Dutch government

The last subtree in Figure 4.5 is on security & compliance. This subtree overlaps with communication on the objective of restricting further growth of governmental domain names. Next to the problem of recognizability that plays a role in communication, the deployment of more domain names also makes it harder to ensure all of these are conforming to (security) standards. To achieve resilient domain names in general, the implementation of specific security standards is mandatory. This adoption is monitored through the use of the measurement tool Internet.nl [19]. Lastly, government websites are mandated to be accessible to citizens with various impairments. As a means to prove compliance, websites can be certified on accessibility.

4.3.4 *Mitigation of security risks*

Now the domain name management policy resources and the corresponding policy theory have been discussed, it may be interesting to compare the policies with the identified risks in chapter 3. It was identified that policies relating to domain name management came forth from different policy fields, which were not all focused on the topic of cyber security. Since domain name management is not a problem that is covered by any RFC standard or related works of literature, policy implementations in practice are based on the organizations' insights and expertise. For each of the three risk categories of the previous chapter, a description is provided of how the current policy implementations aim to address them.

First, a risk exists of domain and subdomain takeovers. Mitigation approaches appear to be limited in current policy documents. The governmental website register does not provide information on the current registration status or registrant, only a responsible governmental organization is listed. Since a large part of the domain names is registered at DPC directly, and the DPC is also responsible for publishing the website register, it is concluded that those domain names are unlikely to be vulnerable to a full-domain takeover. However, for domain names that are registered externally, there is no policy (publicly) defined on how the registration should be terminated. Also, in terms of subdomain takeovers based on outdated Resource Records, the available policy documents do not indicate any mitigations in place for this. This observation is supported by the goal-means tree, whereas mitigations for this risk seem to be missing entirely.

The second risk concerned impersonation and typosquatting. On this matter, several mitigative measures are observed in policy documents. The DPC's domain name policy indicates that domain names are in principle not registered for defensive purposes, except in cases where abuse could lead to the loss or exposure of sensitive information [75]. On the website of the DPC it can be found that they utilize a domain name guarding service by SIDN [13]. This service guards domain name portfolios by identifying domain name registrations intended for typosquatting or impersonation. When malicious domain name registrations are detected, the government could submit a take-down request to prevent future abuse. This has been included as a means in the goal-means tree to ensure real government websites are better recognizable. The use of this service by governmental bodies is not mandatory, but recommended. Individual civilians that want to verify if a certain domain name is indeed governmental, could use the website register of the DPC [76]. Since this register is supposed to contain all governmental website domains, it is a suitable verification method. However, since not all governmental domain names are serving a website, it is not always possible to use this method. For instance, when a governmental domain name is merely used for e-mail services, it is currently not possible to verify the domain name using this register. Furthermore, since the website register is not published in a format that will be understandable for the average citizen, it is doubtful whether it will be used for verification purposes.

Lastly, the risk of security standard non-compliance. A lot of attention is being given to this risk in the policy documents. The *Forum Standaardisatie* has included the security standards discussed in chapter 3 in their list [24]. This means that their use is either mandatory or recommended for governmental domain names. The use of HTTPS, DNSSEC, and the three e-mail standards have all been made mandatory, which means that governmental institutions are required to implement them or explicitly explain why they are unable to do so. The CAA standard has been included in the recommended standards list, which means that while use is highly recommended, organizations are free to make their own considerations. As shown earlier in Figure 4.2, the published website register contains an overview of standard compliance for every domain name created using the Internet.nl measurement tool [19]. This tool is also included as a means in the goal-means tree to achieve more resilient domain names. On inquiry, the DPC has confirmed that a similar compliance test is done for the domain name register in its entirety. These tests include all aforementioned standards, apart from the CAA standard, although this is planned to be added in the future. In terms of risk, security standards non-compliance appears to be sufficiently addressed in the current approach. One remark still to be made, however, is that there is no active enforcement of these policies as of now.

4.3.5 *Outlook to future policy adoption*

The past and present with regard to domain name management at the Dutch central government have been considered, as well as their performance in mitigating identified security risks. However, developments on domain name management are still ongoing and other policy options are being considered. This subsection aims to describe which policy options are currently being considered by the Dutch government.

Throughout this chapter, the possibility of a uniform domain extension, for instance, using a TLD or public suffix, has already been mentioned. Interestingly, it was identified that a public suffix was even already considered in the year 2000, with the internet just upcoming [90]. In 2012, the step towards a governmental TLD increased, when the government applied at ICANN for the gTLDs 'politie' and 'overheidnl' [46]. However, the 'overheidnl' application was subtracted. The '.politie' gTLD is still in possession of the Dutch Police Force at this time, yet has not been used for public purposes yet.

The subtraction did not mean that there was no more interest in a uniform domain extension. Simultaneously, a research report commissioned by the government into the use of a gTLD was published in 2012 [20]. One of the main conclusions was that the added value of a designated TLD would be limited in comparison to a public suffix like '.overheid.nl' or '.gov.nl', both in terms of security and recognizability.

Another research got published in 2019, which contained a study on foreign governmental domain name suffix approaches [70]. In a follow-up study published in 2020 [36], it is concluded that a uniform domain extension would remain a highly costly choice in comparison to a continuation of the current approaches. Nonetheless, steps have been taken since then to further explore the possibilities, with most recently the formal adoption of 'gov.nl' as a public suffix in the public suffix list in May 2022[3]. Thus, it is not unlikely that more governmental domain names may be placed under the 'gov.nl' domain in the near future. Its potential location in the policy arrangement on domain name management at the Dutch government can be seen in Figure b.1.

Reflecting on the explorative actions by the government over the past years, one observation is missing in the conducted research. Once a move towards a uniform domain extension has been made, governmental domains in the other TLDs remain to exist. This means, that even though the management burden on this uniform extension will be marginally lower, the former set of domain names would still have to be managed. It is unlikely that the government would decide to cancel all of the domain registrations, since that would create large risks of impersonation attacks with the former governmental domain names. This is an issue that is not sufficiently addressed yet, and should be solved before a full move can be made.

4.4 conclusion and discussion

This concludes the reconstruction of the policy theory on domain name management of the Dutch central government. It has been identified that policies concerning domain name management originated from three different policy areas: communication, archiving and security & compliance. Policies were not constructed incorporating all perspectives, but rather focused on the core expertise of the policy-making department. This resulted in the fact that the constructed policies were not considered capable of fully mitigating the earlier identified cyber security risks. Similar observations have been made from the goal-means tree analysis that visualized the policy field of domain name management.

Furthermore, it was observed that developments are ongoing in the approach to domain name management. There are considerations to utilizing a public suffix domain for recognizability and easier domain name management. However, it had to be concluded that such a solution would not resolve the domain name management problem entirely, since a lot of the already registered domain names would have to be upheld and thus maintained. No solution has been proposed to this problem as of now.

Lastly, the policies that imposed rules and restrictions regarding domain names did not include an approach to policy enforcement. Therefore it could not be assessed if policy deviations would be detected and acted upon. A

[3] https://github.com/publicsuffix/list/pull/1558

lack of policy enforcement could lead to less policy adoption, potentially leading to negative policy outcomes.

There are three points for discussion. In this chapter, it was attempted to reconstruct the policy theory behind domain name management policies using available policy documents. The number of available documents, however, was limited, and possibly not all relevant documents were published publicly. Therefore, this policy theory might represent an incomplete perspective on the entire policy field of domain name management policies. For future research, it could be considered to submit a formal disclosure request for all documents related to domain name management.

Secondly, policy theory was constructed of policy documents surrounding the concept of domain name management. This concept involved multiple relevant policies, and thus, in essence, multiple policy theories. It could be argued that this manner of conducting policy theory would require additional validation, since the documents originated from different policymakers.

Lastly, the identified policy goals were visualized in the goal-means tree. This tree contained links between the identified means, objectives and goals of the policy documents. Some of these links were not explicitly obtained from the policy documents themselves, but were determined by the author based on the policy texts. This means that these links could be affected by researcher bias and may require additional validation.

5 EVALUATION OF DOMAIN NAME MANAGEMENT IN PRACTICE

5.1 INTRODUCTION

In the previous chapter, the policy theory of domain name management at the Dutch central government was reconstructed. The goal of this chapter is to evaluate the policy outcomes in light of both the cyber security risks identified in chapter 3 and the reconstructed policy theory in chapter 4.

Policy outcomes in this context are domain names that are owned by the government, as well as the method in which they are registered and maintained. Analyzing policy outcomes, therefore, is conducting a technical analysis of governmental domain names. Since not all information on governmental domain names is publicly available, several methods have to be selected that may reveal governmental domain names. Once a list of domain names is compiled, they can be further analyzed for possible anomalies that indicate policy failure or policy gaps.

The starting point of this analysis is the website register that is published by the government and updated monthly [76]. First, the selection of methods and methodology are discussed. Then, the technical analysis is conducted, after which the most notable findings are listed. These findings are then discussed in light of the conclusions of chapters 3 and 4. This will provide an answer to the final sub research question: *SRQ 3: How is domain name management executed within the Dutch central government in practice?*.

5.2 METHODOLOGY

5.2.1 *Method selection*

To obtain governmental domain names for the evaluation, methods have to be applied based on the starting set of domain names. For this, related research has been searched on Google Scholar and Scopus, using keywords 'domain name', 'discovery', 'detection', 'recognition', 'abuse', and combinations thereof. This provided the following related studies in domain name discovery techniques.

In the field of domain name discovery, multiple studies have been conducted in the detection of malicious domain names. For starters, several studies try to predict domain names used for the various typosquatting attacks described in category 2 of chapter 3 and Table 3.2. One approach is to use machine learning techniques, for instance for the detection of typosquatting based on the number of words [58], or based on multiple domain name features [54, 84]. Another approach is the generation of potentially malicious domain names based on various typosquatting models [1]. A third approach is using passive DNS measurements as input for the detection of malicious domains [6, 89]. The last identified approach makes use of Certificate Transparency (CT) logs, a database where all issued TLS certificates are listed [18, 22]. TLS certificates are always issued for specific domain names, making this an interesting source for domain name discovery. Whereas these studies have used secondary data sets as input variables, for instance, known abuse databases or generated domain names, other research has focused on primary sources, being registration databases of registries containing WHOIS data [72] or passive DNS measurement data from authoritative name servers of TLDs [98]. Apart from impersonation attacks, a focus has also been laid on the detection of domain names that are automatically generated to avoid blocklisting [10]. The detection mechanisms used in these studies are not directly applicable to this research, since the purpose is to discover benign domain names instead. Some methods, however, could be adapted to be usable for this purpose.

On benign domain name detection, fewer studies appear to be conducted. Nonetheless, those works that have been identified make use of several other methods that may prove useful in this research. For starters, the work of Yi and Scholz uses an automated web crawler to gather hyperlinks from known organizational domain names, which is used to measure network relations between different organizations. Since governmental bodies may be expected to link toward other governmental entities, this method could prove useful in the detection of new governmental domain names. Another method is the detection of parked domain names using DNS fingerprinting [97]. In this approach, they make use of the DNS Resource Records of known parked domain names to detect other parked domain names. They also use a different data source for DNS data, whereas active DNS measurement data is used from the OpenINTEL project [78]. Active DNS measurement relies on primary sources, namely registration databases of registries, to obtain actual DNS Resource Records of entire TLD zones. Passive measurement, instead, is capturing DNS requests that are handled by authoritative name servers. This makes passive DNS a secondary source and thus may be less complete. One downside to active measurement, however, is that it usually does not include subdomains because these are not part of the registries' zone files [88].

Several of the methods observed in related work can be applied in this thesis, let it be altered to the detection of benign domain names. The selected methods for this thesis are the use of a web crawler, matching on active DNS measurement data and matching on Certificate Transparency logs. The web crawler will be employed for discovery through hyperlinks. As input data, the website register [76] will be used, after its contents have been analyzed and validated using DNS and WHOIS data. Secondly, DNS fingerprints obtained from the detected domain names will be used to match similar domain names in the active DNS dataset of OpenINTEL. There is chosen for active DNS data as this is expected to be more complete than passive data. Then, all obtained domain names so far are then matched with CT logs. This would reveal (sub) domain names that share their certificate with the known domain names, but were hyperlinked towards or not matched based on their DNS fingerprint. Conveniently, the OpenINTEL infrastructure also contains a mirror of CT logs. The last method used is the inclusion of domain names listed in the zone file of .politie, the TLD of the Dutch police. This zone file can be accessed directly, through the ICANN Centralized Zone Data Service (CZDS) [39].

The order in which these methods are applied can influence the outcome of this technical analysis, since a higher number of domain names is likely to provide more results. Optimally, the chosen methods are applied repeatedly, to increase the input size for every method. However, due to limited access to the OpenINTEL dataset, it has been decided to execute the methods once and in the order described above. This means that the results of the web crawler will be used to create DNS fingerprints, maximizing the input on the OpenINTEL dataset analysis. CT logs can best be used last, since the input of certificate matching is only a list of domain names itself, and no other selection features are used.

In the next subsections, the selected methods will be further expanded upon.

5.2.2 *Analysis of the website register*

Before using the website register as input data, its content will be analyzed and validated. To this extent, information will be gathered that can aid in this validation. For starters, DNS Resource Records give information about the configuration of domain names and can be used to compare different domain names with each other. For instance, the name servers of a domain name can tell more about who is managing the domain name. To collect the DNS Resource Records, the tool 'dig' will be used[1]. This tool can be utilized to send DNS requests of various types, and in the scope of this research will be used to gather the most common records like A, AAAA, MX, NS and TXT records. DNS Resource Records will be retrieved for all domain and subdomain names that are part of the website register.

Another source of information is the WHOIS database for domain names. This contains details about the domain name registration, including the registration status, the registrant and the registrar of that domain name. This information can be partly retrieved using the 'whois' tool[2], for which a wrapper in Python is available[3]. However, registries have reduced the amount of information that is disclosed via this tool, due to privacy laws and to prevent excessive data scraping. For instance, SIDN, the registry for '.nl', does not disclose any registrant information via this tool.

[1] https://github.com/tigeli/bind-utils

[2] https://github.com/rfc1036/whois

[3] https://github.com/DannyCork/python-whois

For registrant information, permission was requested and granted by SIDN to use the WHOIS tool available on their website[4]. This information will also be collected for all domains in the website register.

Within the collected information, there will be looked for patterns and deviations manually. If governmental domain names show specific patterns, those patterns can be used in the recognition of other governmental domain names that are currently not part of the data set. On the other hand, deviations and anomalies in the data set may reveal configuration mistakes, or domain names that should not be part of the data set.

5.2.3 *Web crawling discovery*

Based on the domain names retrieved from the website register, it is possible to extract hyperlinks from their websites. These hyperlinks could point toward external domain names, which are potentially governmental. Since the website register is only a subset of the governmental domain name register, it is expected that a large number of domain names can still be discovered.

For this research, a web crawler will be developed that can be deployed on governmental websites in the website register[5]. Its purpose will be the detection of hyperlinks to new (sub)domain names, that are not part of the website register. It is expected that part of the new domain names found on governmental websites, will be owned by the government as well. However, it is also expected that governmental websites contain links to non-governmental domain names. This means that the results have to be validated afterward.

Before deploying the web crawler, the current data set of domain names will be enriched in two ways. First of all, the website register contains either the apex domain name or the domain's 'www' subdomain. Since the 'www' subdomain is technically a different domain name than the apex domain, a web crawler might identify either as a new domain. To prevent this, the data set is enriched to include both the apex domain name and the 'www' subdomain of that domain, if the 'www' domain name is in use. In cases where the website register only contains a multi-order domain name that is not the 'www' subdomain, neither the apex domain nor the 'www' subdomain is considered to be a governmental domain, before this has been validated manually.

Secondly, the website register only contains websites for the central government. It is expected that there will be a significant number of outgoing links toward decentral governments. Therefore, the data set will be expanded with known domain names from decentral governments and decentral intergovernmental collaborations. By including these domain names in the inventory, the number of new domain names is expected to be lower, when analyzing hyperlinks on governmental websites. The domain names of decentral governments will be retrieved from the Dutch governmental almanac [69].

[4] https://www.sidn.nl/whois-direct

[5] https://github.com/WKobes/GovtScraper

5.2.4 *Web crawling results validation*

It is expected that the web crawler produces a large number of new domain names. Not all discovered domain names will be (central) governmental domain names, since it is expected hyperlinks will be present to non-governmental organizations or foreign governments as well. Therefore, the web crawling results will be validated.

The first step is to check the registration status. This way, it is possible to determine whether the domain name is currently registered, or (soon to be) free for registration. If the domain name is registered, the verification is continued. If not, the hyperlink may pose a risk of domain takeover, since the domain can be registered by anyone.

The second step is to check whether the domain name is a subdomain of an already-known governmental domain name. If this is the case, it is assumed that the subdomain is also governmental-owned, and thus validated.

If not, validation is attempted based on the fingerprint of the domain name's registration information and DNS data. This information is compared to those of known governmental domain names. Some fields in the registration information are a stronger implication of being governmental than others. For instance, if the registrar of a domain name is indicated to be 'Rijksoverheid', this is a very strong indication of being governmental. The registrant information, on the other hand, could be more easily manipulated and thus is a less reliable method. On the other hand, these fingerprints can also be used to determine that a domain name is not governmental-owned, for instance when the registrant is a known external company or organization. Another example of a strong indicator is when a domain name operates on a governmental name server, as this indicates the domain name is configured by them. In case the fingerprint is a match with those of governmental domains, the domain is validated.

For the domain names that remain, the researcher applies a manual validation step. This is done by regarding both the domain name registration information as well as the website content, if present. In case the domain name is deemed to be indeed governmental-owned, the domain is included in the inventory, and if possible, the validation fingerprints are improved to accept other domain names sharing these same properties. Domain names for which this can not be determined with absolute certainty will be excluded from the data set.

5.2.5 *DNS and Certificate Transparency analysis*

The preceding methods will result in an enriched and validated version of the governmental website register. However, the methodology is limited in the discovery of domain names to those domains that are hyperlinked from another governmental domain. To expand the data set further, the data set of the OpenINTEL project [78] is used. OpenINTEL is a data project that performs active DNS measurements for domain names. For this purpose, several registries provide access to their TLD zone file, which is used for the active measurements. The domain names in these zone files are queried daily through the DNS, and their Resource Records are recorded. This is done for the apex domain name, as well as the 'www' subdomain. The data is retained indefinitely, and therefore does not only allow for research on current DNS records but also on historical DNS data.

Most zone files are not publicly accessible. Therefore, access to the OpenINTEL database is restricted and access is only granted upon request and may not be used for commercial purposes. Data derived from the OpenINTEL dataset may also be subject to publishing restrictions. For the scope of this research, permission was granted to access the dataset as long as obtained data did not leave the controlled environment. Only aggregated results derived from the data will be included in this thesis, meaning the number of domain names that have been discovered. This also means that validation of individual results is limited, as the individual results will not become available for further analysis.

The data set will be used as follows. In the preceding domain name validation step, domain name fingerprints will have been constructed that commonly match governmental domain names. The fingerprints based on DNS data can be reused to match domain names that share these same records. A program will be developed in Python with Jupyter Notebook to conduct this analysis[6].

The OpenINTEL infrastructure also hosts a mirrored data set of CT logs. These public logs contain all issued TLS certificates, as this is a current requirement for a certificate to be trusted. Certificates are assigned to domain names to support encrypted connections, most frequently over the HTTPS protocol. A single certificate may be valid for multiple domain names in two ways. Firstly, the certificate can contain a wildcard domain name, for instance '*.example.nl', which is valid for any second-level subdomain under 'example.nl'. Secondly, multiple domain names may be specified in the Subject Alternative Name (SAN) field on the certificate, making the use of a single certificate for multiple domain names possible.

The SAN field can be used for the discovery of new domain names. Based on the set of governmental domain names, it is possible to select all valid certificates that are issued for these domains. If a governmental domain name shares a certificate with other domain names, it is expected these other domain names are governmental as well. This selection will be done based on the governmental domain names that have been retrieved using the OpenINTEL DNS data. CT logs are often temporal sharded, mostly in time frames of one year based on the expiration date of the certificate. In light of this research, only CT logs for 2022 and 2023 will be considered, since expired certificates may contain domain names that are no longer used.

[6] https://github.com/WKobes/openintel-analysis

The OpenINTEL data set is based on, as mentioned, the zone files that are provided by registries. The top-level domain '.politie' (police) is currently the only TLD that is in the possession of the Dutch government. Access to the zone files of gTLDs can be requested through ICANN's Centralized Zone Data Service [39]. Access to the zone of '.politie' will be requested. Since the TLD is not open for public registration, all domain names in the zone file can be considered governmental.

5.3 RESULTS

5.3.1 *Technical results*

In this subsection, the technical results of the conducted review are provided. The intermediate results are presented for every execution step of the methodology. The most notable findings are highlighted, followed by an identification of common risks within those findings.

5.3.1.1 *Analysis of the website register*

The website register of the Dutch central government is updated monthly, to represent a current overview of domain names. A screenshot of the register is shown in Figure 4.2. Newly registered domain names with a website are added to the overview, while expiring domain names are removed. Furthermore, the statistics regarding the standards compliance scans and numbers of monthly visitors are updated. Over the course of this research, the website register has been regularly collected and stored. In Table 5.1, the number of domain names in the register over time is listed. As can be observed, a steep decrease in the number of domain names has occurred in the version of June 2021. It is expected that a more thorough analysis had been conducted at that time, after which domain names that no longer served a governmental website were removed. A second observation is that the number of domain names serving a website is, apart from this single correction, constantly increasing.

In the analysis, the register versions of July 2021 to November 2021 were used. Initially, as much information as possible was gathered about these domain names. As described, both WHOIS and DNS information was gathered. On manual inspection of the retrieved data, several inconsistencies were found concerning the registration of domain names. First of all, while requesting the WHOIS information, three domain names were found to be free for registration and could be taken over by the researcher. Domain names that are free for registration do not return any WHOIS information and are recorded with the status 'free'. After a manual inspection of the gathered WHOIS data, one other domain name was identified, which was assumed to be taken over by an external party. This assumption was based on the use of a personal e-mail address as a point of contact. This finding and its potential impact are further described in the next section, under Finding 1.

Month	#	**Month**	#
Oct/20	1608	Sep/21	1678
Nov/20	1628	Oct/21	1688
Dec/20	1651	Nov/21	1718
Jan/21	-	Dec/21	1733
Feb/21	1708	Jan/22	1751
Mar/21	1742	Feb/22	-
Apr/21	1762	Mar/22	1758
May/21	1764	Apr/22	-
Jun/21	1624	May/22	1786
Jul/21	1653	Jun/22	1790
Aug/21	1674	Jul/22	1797

Table 5.1: Number of domain names in the governmental website register over the period October 2020 until July 2022

Then, 16 subdomain names were identified in the website register, of which is disputed that their apex domain names are governmental. These domain names, while the subdomain was serving a governmental website, were suspected to be owned by an external party. These cases were found to be genuine, for instance, external service providers. This is further elaborated upon in the next section, in Finding 5.

The remainder of the analysis on the website register did not unveil any other anomalies. The remaining gathered information was used for creating a fingerprint for governmental domains. This was done for both DNS records as well as the registrant and registrar information. For example, it was identified that the central government is in the possession of three different registrar contracts: one for the central government, one for the Dutch Tax Administration and one for the Netherlands Vehicle Authority (RDW).

5.3.1.2 *Web crawling discovery*

The initial starting point of the web crawler discovery phase is at 1718 domain names from the website register of November 2021, a version in which the earlier found irregularities have been corrected. By adding the apex domains or 'www' subdomains to the inventory, 1253 new domain names are added to the inventory bringing the total to 2971. As explained, the 'www' or apex version of a domain is only added to the inventory if its counterpart is already part of the website register. Furthermore, the 'www' subdomain is not added if it has no Resource Records defined, since this indicates that the subdomain is non-existent.

To retrieve the domain names of other governments, a simple web scraping tool was developed in Python that retrieves the unique domain names from the governmental almanac. This resulted in 780 new apex domain names, and approximately the same number of 'www' subdomain versions.

The number of domain names retrieved from the governmental almanac was slightly higher, however, there was a small overlap between the almanac and the website register.

To discover entirely new domain names, not listed in one of the available resources, a web crawler[7] was developed that crawls the known governmental websites. This tool was developed in Python using the scrapy package[8]. The tool starts on the root of the website and retrieves all hyperlinks on that page using an XPath matcher. These hyperlinks include internal links, meaning they point to other pages within the same domain name, and external links to other domain names. Internal links are added to a list and will be automatically visited by the web crawler, retrieving all hyperlinks on that page as well, until all linked pages of the website have been visited. External domain names are checked against the current inventory and if the domain name is not known yet, it is logged for further inspection.

Due to time limitations, this crawler has only been applied fully on the top three domains of the website register, based on visitor numbers in November 2021 [76]. This were at the time `rijksoverheid.nl`, `coronadashboard.rijksoverheid.nl` and `rivm.nl`. Although the method has only been applied on these websites, this method still provided 3047 new (sub)domain names unknown to the inventory. These domain names were taken into the validation step, to determine how many of these were indeed governmental.

5.3.1.3 *Web crawling results validation*

The described validation methods have been applied to the 3047 new domain names that were found using the web crawler. Based on the fingerprints, it was possible to validate that 186 out of the 3047 new domain names were indeed governmental-owned. The remainder of the domain names were either domain names of external organizations, foreign governments, or otherwise. It should be noted that governmental domain names may have been excluded, if their WHOIS information or website did not indicate governmental involvement. Combined with their respective 'www' subdomains where applicable, the inventory is now expanded to a total of 4963 unique domain names, consisting of 2201 apex domain names and 2762 subdomains. Since the website register is only a small subset of the domain name register, it was the expectation that a large portion of the 186 newly verified domain names is already part of this domain name register. However, after verification with the DPC, it was found that out of the 186 domain names, 9 were incorrectly excluded from the website register. This notable finding is expanded upon in Finding 2 of the next section.

During the validation phase, the information on new governmental domain names was used to extend the fingerprints used for automatic validation. This resulted in identifying that 10 governmental institutions were in the possession of a registrar contract in total. Furthermore, 406 unique reg-

[7] `https://github.com/WKobes/GovtScraper`
[8] `https://scrapy.org/`

istrant organizations were identified that are used as registrants for domain name registrations.

5.3.1.4 *DNS and Certificate Transparency analysis*

The set of 4963 governmental domains that have been collected and validated up until now, was used as input for the data analysis on the OpenINTEL dataset. The code written in Jupyter Notebook[9], was able to load the data of the OpenINTEL dataset and conduct the analysis. The infrastructure is based on Apache Spark[10], an engine for large-scale data analytics and data files can be exported in AVRO files. The code was run on a Virtual Machine with access to the Spark infrastructure.

As mentioned earlier, access to the OpenINTEL data set is limited due to contractual constraints with the registries that provide their zone information. This means that for this research, the results obtained from the OpenINTEL analysis could not be subjected to the same validation steps as conducted in the previous section. These validation methods would require the use of external sources, like the DNS, which could indirectly lead to data leakage through passive DNS measurements. The implications of this for the validity of these research steps are described in the discussion. In addition to the number of newly identified domain names, noteworthy observations made during the manual inspection of the results in the controlled environment will also be included.

For the OpenINTEL DNS data, matching will be done based on DNS Resource Records. It was chosen only to select based on A, CAA, MX and NS records. This choice was made, because the other Resource Records in the data set either were unique per domain name and thus would not yield any results, or were too generic and would match a large number of non-governmental domain names. To illustrate, DNS Records related to DNSSEC signing would contain unique signatures per domain name, and therefore cannot be used. On the other hand, TXT records that contain a DMARC configuration, for example, had shown overlap with non-governmental domain names. One exception is the AAAA record, for IPv6 addresses, since these would provide unique matches on input data. The problem here is that the IPv6 ranges allocated to the government contain an incomprehensible number of possible IPv6 addresses, making comparisons on all possible records impossible. At the same time, all governmental domain names that contained an AAAA record also contained an A (IPv4) record. Thus, the use of IPv4 addresses in this comparison would provide the same results.

Based on the domain names found so far, a list of each of the four DNS Resource Records was created. Only values were included to which it was sure they were governmental, for instance, only NS records that contained governmental domain names, and only A records with IPv4 addresses assigned to the government. This resulted in the number of records as shown in Table 5.2. The number of A records highly outnumbers the other records,

[9] https://github.com/WKobes/openintel-analysis
[10] https://spark.apache.org/

because the known IPv4 ranges assigned to the Dutch government were fully included. These ranges were obtained by cross-referencing the IP addresses of known governmental domain names with the RIPE database[11]. At the same time, almost all governmental domain names registered at the DPC used the same name servers, explaining the low number of unique NS records.

Resource Record Type	#
A	790808
CAA	27
MX	21
NS	41

Table 5.2: The number of Resource Records selected for every RR type

The comparison of the Resource Records on the OpenINTEL dataset was executed. The date of reference of the data set is 24 July 2022. The full results are shown in Table 5.3. This shows that the number of identified governmental domain names has increased by 8655 to 13618 unique domains and subdomains.

RR Selector	# new domains
A	3996
CAA	2507
MX	352
NS	6006
Total (unique)	8655

Table 5.3: Number of governmental domain names identified in OpenINTEL using various Resource Records and the total of newly discovered domain names

Now, the 13618 known domain names are used in the SAN field matching using the CT logs data set. The results are shown in Table 5.4. As can be seen, 592 domain names using a wildcard on their certificate were identified. Since wildcard domains do not represent actual domain names, these will not be added to the inventory. This means that the CT data set is bringing the total to 18613 domain and subdomain names. Manual inspection of the matched domain names showed that there was at least one case where a governmental domain name shared a TLS certificate with a non-governmental domain of a commercial party. This is further elaborated upon in Finding 6 of the next section and the discussion of this chapter, as this may indicate the number of found domain names is an overestimation of the number of actual governmental domain names.

[11] https://www.ripe.net/manage-ips-and-asns/db

Type	**# new domains**
Domain names	4995
Wildcards (*)	592
Total (unique)	5587

Table 5.4: Number of governmental domain names identified using Certificate Transparency logs

In the last step, the zone file of the '.politie' TLD was downloaded from ICANN's CZDS, and the domain names were counted. The results confirm the earlier observation that the top-level domain is not in active use yet, since the zone file only contained 10 unique domain names.

This concludes the analysis. The final number of governmental domain and subdomain names identified is now 18623, compared to the 1718 domain names at the start.

5.3.2 *Notable findings*

The technical results of the review have been discussed in the previous section. In the following paragraphs, the most notable findings of this review will be further discussed, in terms of possible cyber risks. Where possible, the findings are mapped towards their corresponding risk category identified in chapter 3.

All findings have been reported to the responsible departments. Inconsistencies have since then been corrected in the registers by DPC. In January 2022, the results were presented in an expert session[12] concerning domain name management for governments. About 30 different governmental entities participated in this session of different layers of government.

FINDING 1 - FREE AND TAKEN-OVER DOMAIN NAMES IN MANAGEMENT SYSTEM During an initial analysis of the website register, two listed domain names were found to be unregistered. In a later phase of the research, a third domain name was put into quarantine and became available for registration after the grace period. The researcher registered these domain names in their personal capacity, to prevent any takeover by potentially malicious actors. It was observed that the domain names did not get automatically removed from the register in newer revisions. From this, it was concluded that there is no monitoring for domain name ownership changes, nor monitoring for whether domain names in the website register move into a quarantine period or become available for registration.

While investigating the registration information of the domain names in the website register, one domain name was found to be registered by a personal e-mail address. Upon inspection of the website, it was determined

[12] Beheersbaarheid Internetdomeinen (BID), 11 January 2022, Online meeting

that this was not a governmental website. Instead, the website listed advertisements of questionable legitimacy. Notable is that, while the finding was done in 2021, this domain name was last registered in 2015. From this is derived that the domain name has become available for registration during that time, and has since the takeover remained part of the domain name management systems.

Free domain names can be freely registered by anyone on the internet, meaning that these domain names could be abused to host malicious content or to mimic governments. These domains are susceptible to takeover attacks, as described in the first category of chapter 3.

Finding 2 - Missing domain names from management system

Since the public website register is only a subset of the larger domain name register of the government, it was not possible to automatically verify whether found domain names were already part of this register. However, it was possible to check if newly found domain names served a governmental website. In this case, it would be expected to be part of the website register and these cases could be reported to the DPC.

This resulted in nine discovered domain names that served websites yet were not listed on the website register. After reporting, it turned out that five of these were part of the domain name register, but not recorded as a domain name that is serving a website.

The other four domain names that served a website were unknown to both the website and the domain name register. This means that they fell outside of the existing domain name management system and thus were not included in analyses. The DPC has started an investigation into whether these domain names are indeed governmental-owned. Two of these missing domains were registered under the .nl TLD and two were registered under .com.

Two other domain names that were identified to be serving a website, were excluded from the website register due to other reasons. These two domain names, even though the registration was managed by the Dutch central government, served websites that were not considered part of the central government.

Governmental domain names missing from the register may remain out of sight in regular security assessments and not meet the required security standards set by the government, for instance, non-compliance with the security standards described in category 3 of chapter 3.

Finding 3 - Personal governmental domain name registrations

During the investigation, the WHOIS details of all domain names were retrieved to look for possible anomalies. In one of the domain names that were found to be missing from the register (see Finding 2), it was discovered that the registration was done in the personal capacity of a civil servant. This was derived from the fact that a personal (non-governmental) e-mail address was used as the administrative contact person. This means that the domain name asset that is serving a governmental website, is in fact not in the possession of the government, but rather of an individual.

This can have consequences whenever the individual civil servant stops working for the current organization. The government may lose access to the domain name configuration in this case. Secondly, personal registration details make it unclear whether the domain name is owned by the government. In case the government loses access to the domain name, it can be technically seen as a domain name takeover as described in category 1 of chapter 3.

FINDING 4 - DOMAIN NAME USAGE AFTER RELEASE Throughout the research period, various versions of the website register were analyzed and monitored whether domain names were added or removed. Table 5.1 showed the exact time periods for which the register was retrieved. In contrast to the domain names that became free for registration while they were still part of the domain name register (Finding 2), it was also observed that multiple domain names were released according to policy. This means that the domain name was first removed from the website register, after which it was set into quarantine and became available for registration.

The proper release of domain names does not lead to the risks discussed in Finding 2, however, for some of these domain names, it was discovered that they were still used in hyperlinks on other governmental websites. This means that these hyperlinks stopped working and that a new registrant could present any content to the users of the governmental websites that pressed the outdated hyperlink. Having hyperlinks to external domains may indicate a relation between the two domain names, effectively legitimizing the linked domain name as a trusted source.

Therefore, even though the policy on domain name release was followed accordingly, there remained a possibility for impersonation attacks using these domain names. Since the domain name release was intentional, this risk should not be considered a domain name takeover, but rather an impersonation attack using left-over hyperlinks. Thus, the risk can be placed under category 2 of chapter 3.

FINDING 5 - GOVERNMENTAL SUBDOMAIN-ONLY DOMAIN NAMES The fifth notable finding is that some domain names that are part of the website register are served under a higher-order domain name, while the apex domain name was not governmental-owned. This was observed primarily in cases where the government leveraged an external service and the service is hosted by the server provider, rather than the government itself. These domains could be in a form of 'government.<service>.nl', for instance.

While the service may be legitimate, end-users of said service could be confused by the usage of higher-order domain names. For example, end-users could falsely assume that the apex domain name is also controlled by the government. Another risk is that the operator of the root domain name can control and monitor all activities on their subdomains. This means that more trust has to be laid in the service provider, than when the service would be provided on a governmental domain name. A rogue service provider would effectively execute a subdomain takeover in these cases, meaning this risk falls under category 1 of chapter 3.

finding 6 - shared tls certificates with non-governmental domain names The sixth and final notable finding is derived from the CT logs data, namely that TLS certificates exist that are shared by both governmental and non-governmental domain names. This means that both domain names were included in the SAN field of such certificates. Since certificates are used to set up the confidential connection between the client and the server, anyone with access to the private key of the certificate could compromise this secure connection.

It is questionable whether governmental websites should be protected with a certificate shared with non-governmental parties. Especially when websites offer sensitive services involving personal information, these connections should only be between the user and the government directly. Using a shared certificate means that all domains that serve this certificate, could potentially eavesdrop on this connection. In terms of the risks discussed in chapter 3, the consequences of such an attack would be comparable to a domain name takeover of category 1. However, eavesdropping on secure connections is not trivial, nor is the domain name compromised in the long term. Using shared TLS certificates is, as long as it is done intentionally, more an operational risk than a technical one.

5.3.3 *Identified common risks*

Considering the notable findings done in the review, it is concluded that the current implementation of domain name management by the Dutch central government still incurs security risks from the various risk categories identified in chapter 3. The current policies and policy enforcement do not fully cover these risks. The root causes of these findings are described in terms of risks.

Findings 1, 2 and 3 all relate to domain names that were not registered according to the current domain name policy. According to the central policy [75], all domain registrations must be done at DPC directly. However, it is identified that the domain names that involved these findings are or were all registered at an external registrar. Therefore, the coordinating party did not have access to the necessary information and they were not notified by the governmental organization on (de)registration of the domain names. If these domain names were directly registered at DPC, there would be no need for additional information provision, since the technical status of the domain name could be read directly from the zone file managed by DPC. Non-compliance with the domain name policy indicates a lack of sufficient policy enforcement.

Findings 4, 5 and 6, in turn, are caused by hiatuses in the current domain name policy. While the current policy includes methods for domain name release, there is no check on whether the domain name is still used in other places on governmental websites. This is a direct cause for Finding 4. Finding 5 is partly covered as it could be argued that domain names have to be centrally registered, yet it is unclear whether this applies to subdomains as well. Finding 6 could occur as there was no policy identified that restricts the sharing of TLS certificates in any way. It now has been shown that the absence of certain policies, as well as the lack of policy enforcement, causes security issues in the current implementation of domain name management in the Dutch central government.

5.3.4 *Policy versus Practice*

This evaluation shows that in terms of cyber security, the policy outcomes of domain name management leave room for improvement. Looking back to the policy theory that was reconstructed in chapter 4, these findings are coherent with the main conclusions of the policy theory.

Cyber security being absent among the goals of the domain name policy showed that a focus was laid on the other policy fields: communication and archiving. The identified policy hiatuses all relate to situations that affect cyber security, but are less relevant in terms of communication and archiving. It could be argued that the use of outdated hyperlinks and subdomain-only domain names is also unfeasible in terms of usability and recognizability of citizens, yet their impact is larger on the security aspect. In a domain name policy with more focus on security, it could be expected that these topics would be addressed.

The identified deviations from existing policy documents are indicative of a lack of policy enforcement. Whereas found policy documents did not specify which organization would be responsible for enforcement, it may now be assumed this role was not specifically assigned. Therefore, policy deviations are not detected and thus not acted upon. Even with the aforementioned improvements in policy incorporated, the policy system will remain vulnerable to security risks while policies are not being fully adhered to.

Lastly, chapter 4 provided an outlook on ongoing policy research, like the consideration of a public suffix for all governmental domain names. Once a move would be made toward a public suffix, several of the findings made in this chapter would be less likely to occur or have their impact reduced in terms of security risk. Only Finding 6 could remain to pose a risk, if a TLS certificate is shared between governmental domains with the public suffix and external domains. The other findings could still occur, but would pose no security risk since domain names within the public suffix are not open for registration by the general public. Additionally, the risks remain relevant at least until the transition toward a public suffix would be fully completed. Even after completion, it could be argued that current regular domain names should be upheld, meaning that they also have to keep being managed.

5.4 CONCLUSION AND DISCUSSION

In this chapter, the approach to domain name management of the Dutch central government has been evaluated. Various methods obtained from related work have been applied to measure the level of correctness of this approach. The review resulted in six different notable findings on the current implementation, for which two common risks were identified: (1) a lack of information due to external domain name registration, and (2) hiatuses in the current domain name policy. While not all domain name management resources of the government are publicly available, it was possible to gather these results with the usage of various public resources as well as the OpenINTEL data set.

The observations of both the technical results and the policy theory in chapter 4 have shown that the current policy outcomes are not fully in line with the policy goals. A primary challenge for the government is that individual departments do not always comply with the set policies, and there is currently no department in charge of policy enforcement. Secondly, the set of policies shows certain hiatuses, for instance on the points of TLS certificate sharing and the prevention of remaining hyperlinks towards expired domain names. With this, the third sub research question in this thesis has been addressed: *SRQ 3: How is domain name management executed within the Dutch central government in practice?*.

A point for discussion is the limited access to the governmental domain name register. Having only had access to the website register means that the starting point was only about 10 percent of the total number of known governmental domain names. Using the described novel combination of methods, it was possible to increase this number to 18623 unique domain names, surpassing the approximate number of 16000 domain names in the governmental domain name register. However, it can be asserted how many of these domain names are part of the domain name register, and how many are to be considered new and unique finds. Furthermore, the results derived from the OpenINTEL database could not be thoroughly validated due to restrictions on access to this data. Therefore, it is possible false positives are included in this number, when domain names that are not owned by the government use the same DNS Resource Records as governmental domain names. Similarly, false positives may have been given by the Certificate Transparency data, when certificates were shared between governmental and non-governmental domain names. Future research with unrestricted access to these data sources is required to be able to fully validate the found domain names.

Secondly, time limitations made it possible to apply the web crawler discovery methodology to the three most visited governmental domain names only. It is expected that significantly more domain names would have been identified, if this methodology was conducted on more or all known governmental domain names. However, the results of the discovery clearly show that this method is a feasible way to discover new domain names, and thus in theory could be applied still if given enough time.

Lastly, the methods used in this chapter all relied on secondary sources of domain name discovery. The success of the methodological approach is highly dependent on the number of governmental domain names that are supplied at the start. Several primary sources, unavailable to the researcher, could be used in practice to improve the results. For example, if the government were to execute this methodology themselves, they could use the entire domain name register, as well as all domain and subdomain names included in the zone files of the governmental name servers. Nonetheless, the methods described in this chapter have shown to be effective in the detection of domain names that were even unknown to the governmental internal registers, and thus are still relevant for use in practice.

The current analysis relies directly on domain names and their technical properties. In future studies, it could be considered to expand this analysis using more properties that indicate governmental domain names. For instance, websites could be checked for the presence of governmental logos, or a textual analysis could be conducted on the content to determine if a website is governmental based on context. This way, governmental domain names could be detected that do not share any technical properties with the currently identified domain names.

6 CONCLUSIONS AND DISCUSSION

6.1 INTRODUCTION

In the previous chapters, the subquestions have been addressed. Based on these outcomes, an answer will be formulated for the main research question: **RQ: How should the Dutch government implement domain name management?**. After that, the main recommendations are listed for the improvement of the Dutch government's approach to domain name management.

The results are then put up for discussion, by laying out the limitations of this research. This chapter ends with the research paths that remain open for future work.

6.2 CONCLUSION

In this research, the domain name management practices of the Dutch central government have been studied. To understand what domain name management entails, a background study has been conducted on domain names and DNS. This has shown that DNS is a complex concept, which is implemented based on a large number of technical documents.

Specific attention was given to the cyber security risks that are relevant for domain name owners. These risks were abstracted from the literature and it was determined these could be divided into three different risk categories. The literature on risk in DNS commonly did not distinguish between risks that are relevant for DNS operators and those that are relevant for domain name owners. This thesis argued that for domain name management, only the latter risks are relevant to take into account. Therefore, several risks were excluded from the review. Policies that address domain name management should incorporate mitigative measures for all three risk categories. The description of these three risk categories is the first contribution of this thesis to the current literature.

As a second academic contribution, the policy theory of domain name management at the Dutch government was constructed, complemented with a visualization in the form of a goal-means tree. When investigating this policy theory, it was identified that the policies were focused on three different policy areas: communication, archiving and compliance & security. It was found that individual policies did not necessarily cover all three aspects, where most notably security & compliance were missing from a general domain name policy. This could also be observed from the composition of the goal-means tree that has been constructed in Figure b.1. The Dutch government should ensure security is included as a core goal in domain name management, specifically to mitigate the identified relevant risks. Furthermore, based on documents describing the ongoing developments in domain name management, it was concluded that no attention was given to the fact that current domain names will have to be managed, even after transitioning to a new domain name approach. The government should take this into account when deciding on future domain name management policies.

In an evaluation of the policy outcomes, several shortcomings were identified. Some of these shortcomings were caused by non-compliance with the existing policies. These should be addressed by applying active policy enforcement. Other shortcomings were caused by a gap in policy, since several potentially unwanted practices concerning domain names were not addressed in policies. By defining policies for these situations, domain name operators know how to configure their domain names better. The evaluation applied several novel approaches to the discovery and validation of domain names, which could be applied in practice to detect domain names that are missing from domain name management systems. The methodology could also be applied to detect deviations from current policies, for instance for policy enforcement. This novel methodology represents the third and last academic contribution done by this thesis.

With this, an answer has been formulated to the main research question.

6.3 RECOMMENDATIONS

This thesis brings the following five concrete recommendations for the Dutch central government to improve its approaches to domain name management:

- Make improving and maintaining cyber security one of the explicit goals of the domain name policy. Currently, the central policy is focused mostly on communication and archiving. By incorporating relevant cyber security risks, the policy could be designed better to mitigate these risks. This could improve the general awareness of cyber security risks related to domain names as well.

- Ensure that policy is enforced. This thesis has shown that current policies are not always adhered to. There is no central department in charge of policy enforcement, nor are there any consequences for not adhering to the policies. If policies were to be actively enforced, policy compliance would most likely increase.

- Do not consider a move towards a uniform domain extension as a solution to all current problems. The research conducted in a move toward such domain extension, laid focus on the costs and efforts for transitioning the current domain names. No attention is given, however, to the fact that current domain names cannot simply be canceled. These domain names have been used as the main communication channels with the government in the last twenty years. Canceling their registrations will improve the chance of impersonation attacks.

- Draft policy concerning shared TLS certificates, the use of domain names of external service providers and prevention of domain name usage after release. This research has shown that cases exist where (1) a TLS certificate is shared between governmental and non-governmental domain names, (2) some governmental websites are hosted on a subdomain of a non-governmental domain, and (3) hyperlinks toward expired governmental domain names may be left behind on other governmental websites. Current policies on domain name management did not provide directions in these situations. It could be argued that all three cases could lead to potentially unwanted situations, and thus should be prevented. If so, this should be added to the policies, including who is responsible for preventing these cases from occurring.

- Apply automated tooling to detect deviations from policy. Tools are currently already applied to measure compliance with internet standards. Methodologies in this thesis could be used to measure compliance with the domain name policies as well. This could also aid in policy enforcement.

6.4 DISCUSSION

6.4.1 *Validation*

In terms of validation, this research was based on reputable sources from literature. The selection of sources has been extensively described, and no works have been excluded on other grounds than those explicitly mentioned. Since no other works have been identified that address domain name management specifically, there are no works to which the results can be compared directly.

The policy theory in chapter 4, as well as the corresponding goal-means tree, was constructed based on public policy documents. The goal-means tree depicts from which documents individual goals and means were obtained. This makes the work verifiable for future research on domain name management policies.

The data gathering in chapter 5 has been done systematically. The fingerprints used for data validation were carefully constructed, only containing information that was confirmed to match governmental domain names. Every step of the data-gathering procedure has been explicitly mentioned including the intermediate number of results. The used code for data gathering and analysis has been publicly published, making the entire methodology verifiable and reusable.

6.4.2 *Limitations*

This research was based on the resources that were made available publicly, and the restricted database OpenINTEL. This means that internal policy documents and data sets could not be incorporated. Internal documents may provide a more complete insight into the policy theory behind domain name management. If more time would have been available, it could have been considered to request these internal documents, for instance using a freedom of information (Woo) request.

Another limitation of this research is that the results in terms of identified domain names will rapidly become out of date. The cause is that continuously new governmental domain names are registered, while others are canceled. Due to this reason, it was decided not to publish the full list of governmental domain names. Doing so would only degrade the quality of domain name sources online, since the list would not be kept up to date.

Lastly, the results of the analysis on the OpenINTEL database could not be individually verified, due to restrictions on data access. This means that possibly not all identified domain names were governmentally owned. Future research should find other means to validate these results, for instance by comparing the results with the government's internal domain name register. Also, other techniques can be considered for the analysis, like content and textual comparisons.

6.4.3 *Future Work*

This thesis has introduced several novel approaches for analyzing and improving domain name management. Multiple opportunities remain for future studies, to expand upon the work done here.

For starters, this research focused on domain name management within the government. Larger cooperations may also have an interest in improving their domain name management approaches. Unlike the government, cooperations tend to be less country-dependent and thus an analysis in domain name management would entail incorporating more TLDs and WHOIS fingerprints. It should be determined if the present methodology can be directly applied.

Secondly, this thesis has argued the need for proper domain name management based on different cyber security risks. Domain names are communication channels that allow interaction between the government and its citizens. This role is also observed for governmental accounts on social media, like Facebook, Twitter and Instagram. The Dutch government operates a large number of accounts on these services. Citizens may not always be able to verify whether a social media account is governmental, which means that social media accounts could be abused for impersonation as well. Research could be conducted on the best approaches to social media account management on a large scale.

Lastly, this study focused on domain name management at the level of the Dutch central government. It could be argued that the inclusion of decentral government in domain name management further improves the security of the government in total. Additional research should be conducted on how decentral governments can be best included in the system, as these governments will be more heterogeneous than departments on the central level of government.

Part II

APPENDIX

a

ADDITIONAL TABLES

RFC	Title	Status
1034	DOMAIN NAMES - CONCEPTS AND FACILITIES	Internet Standard
1035	DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION	Internet Standard
1101	DNS Encoding of Network Names and Other Types	Unknown
1995	Incremental Zone Transfer in DNS	Proposed Standard
1996	A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)	Proposed Standard
2136	Dynamic Updates in the Domain Name System (DNS UPDATE)	Proposed Standard
2181	Clarifications to the DNS Specification	Proposed Standard
2606	Reserved Top Level DNS Names	Best Current Practice
3007	Secure Domain Name System (DNS) Dynamic Update	Proposed Standard
3225	Indicating Resolver Support of DNSSEC	Proposed Standard
3226	DNSSEC and IPv6 A6 aware server/resolver message size requirements	Proposed Standard
3492	Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)	Proposed Standard
3596	DNS Extensions to Support IP Version 6	Internet Standard
3597	Handling of Unknown DNS Resource Record (RR) Types	Proposed Standard
4033	DNS Security Introduction and Requirements	Proposed Standard
4034	Resource Records for the DNS Security Extensions	Proposed Standard
4035	Protocol Modifications for the DNS Security Extensions	Proposed Standard
4343	Domain Name System (DNS) Case Insensitivity Clarification	Proposed Standard
4501	Domain Name System Uniform Resource Identifiers	Proposed Standard
4592	The Role of Wildcards in the Domain Name System	Proposed Standard
4955	DNS Security (DNSSEC) Experiments	Proposed Standard
5011	Automated Updates of DNS Security (DNSSEC) Trust Anchors	Internet Standard
5452	Measures for Making DNS More Resilient against Forged Answers	Proposed Standard
5731	Extensible Provisioning Protocol (EPP) Domain Name Mapping	Internet Standard
5890	Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework	Proposed Standard
5891	Internationalized Domain Names in Applications (IDNA): Protocol	Proposed Standard
5892	The Unicode Code Points and Internationalized Domain Names for Applications (IDNA)	Proposed Standard

5893	Right-to-Left Scripts for Internationalized Domain Names for Applications (IDNA)	Proposed Standard
5910	Domain Name System (DNS) Security Extensions Mapping for the Extensible Provisioning Protocol (EPP)	Proposed Standard
5936	DNS Zone Transfer Protocol (AXFR)	Proposed Standard
6014	Cryptographic Algorithm Identifier Allocation for DNSSEC	Proposed Standard
6147	DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers	Proposed Standard
6376	DomainKeys Identified Mail (DKIM) Signatures	Internet Standard
6452	The Unicode Code Points and Internationalized Domain Names for Applications (IDNA) - Unicode 6.0	Proposed Standard
6698	The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA	Proposed Standard
6725	DNS Security (DNSSEC) DNSKEY Algorithm IANA Registry Updates	Proposed Standard
6761	Special-Use Domain Names	Proposed Standard
6762	Multicast DNS	Proposed Standard
6763	DNS-Based Service Discovery	Proposed Standard
6840	Clarifications and Implementation Notes for DNS Security (DNSSEC)	Proposed Standard
6891	Extension Mechanisms for DNS (EDNS(0))	Internet Standard
7208	Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1	Proposed Standard
7489	Domain-based Message Authentication, Reporting, and Conformance (DMARC)	Informational*
7671	The DNS-Based Authentication of Named Entities (DANE) Protocol: Updates and Operational Guidance	Proposed Standard
7686	The ".onion" Special-Use Domain Name	Proposed Standard
7720	DNS Root Name Service Protocol and Deployment Requirements	Best Current Practice
7766	DNS Transport over TCP - Implementation Requirements	Proposed Standard
7858	Specification for DNS over Transport Layer Security (TLS)	Proposed Standard
7873	Domain Name System (DNS) Cookies	Proposed Standard
8310	Usage Profiles for DNS over TLS and DNS over DTLS	Proposed Standard
8484	DNS Queries over HTTPS (DoH)	Proposed Standard
8490	DNS Stateful Operations	Proposed Standard
8659	DNS Certification Authority Authorization (CAA) Resource Record	Proposed Standard

8753	Internationalized Domain Names for Applications (IDNA) Review for New Unicode Versions	Proposed Standard
8880	Special Use Domain Name ‘ipv4only.arpa’	Proposed Standard
8945	Secret Key Transaction Authentication for DNS (TSIG)	Internet Standard
9018	Interoperable Domain Name System (DNS) Server Cookies	Proposed Standard
9103	DNS Zone Transfer over TLS	Proposed Standard
9156	DNS Query Name Minimisation to Improve Privacy	Proposed Standard
9157	Revised IANA Considerations for DNSSEC	Proposed Standard
9022	Domain Name Registration Data (DNRD) Objects Mapping	Proposed Standard
9076	DNS Privacy Considerations	Informational*
9233	Internationalized Domain Names for Applications 2008 (IDNA2008) and Unicode 12.0.0	Proposed Standard
9250	DNS over Dedicated QUIC Connections	Proposed Standard

Table a.1: List of current RFC documents that contain the DNS specification with its corresponding status

b

GOAL-MEANS TREE

The goal-means tree is included on the next page.

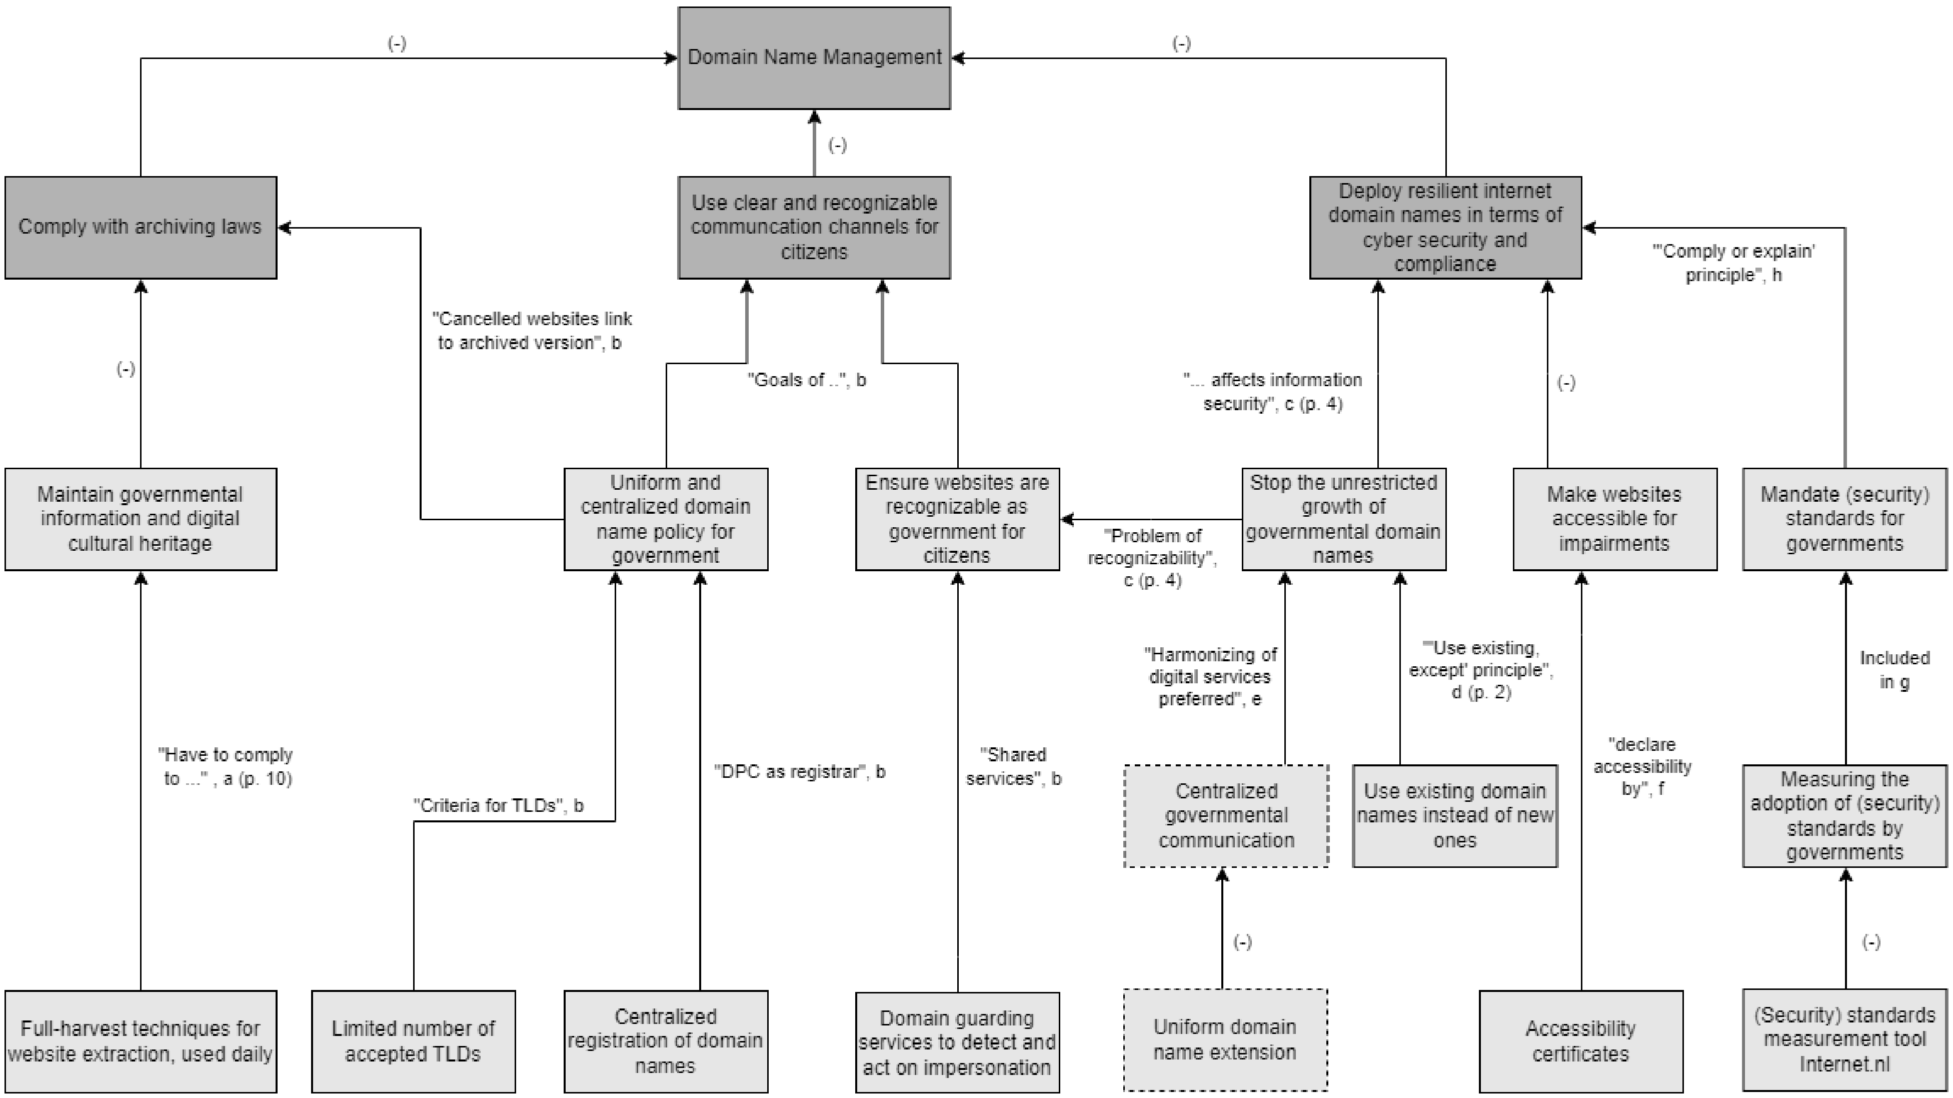
Domain Name Management
(-)
(-)
(-)
Comply with archiving laws
Use clear and recognizable communcation channels for citizens
Deploy resilient internet domain names in terms of cyber security and compliance
"'Comply or explain' principle", h
"Cancelled websites link to archived version", b
(-)
"Goals of ..", b
"... affects information security", c (p. 4)
(-)
Maintain governmental information and digital cultural heritage
Uniform and centralized domain name policy for government
Ensure websites are recognizable as government for citizens
"Problem of recognizability", c (p. 4)
Stop the unrestricted growth of governmental domain names
Make websites accessible for impairments
Mandate (security) standards for governments
"Harmonizing of digital services preferred", e
"'Use existing, except' principle", d (p. 2)
Included in g
"Have to comply to ..." , a (p. 10)
"Criteria for TLDs", b
"DPC as registrar", b
"Shared services", b
"declare accessibility by", f
Centralized governmental communication
Use existing domain names instead of new ones
Measuring the adoption of (security) standards by governments
(-)
(-)
Full-harvest techniques for website extraction, used daily
Limited number of accepted TLDs
Centralized registration of domain names
Domain guarding services to detect and act on impersonation
Uniform domain name extension
Accessibility certificates
(Security) standards measurement tool Internet.nl

Figure b.1: Goal-Means tree of domain name management at the Dutch government. Sources: a [65], b [75], c [16], d [77], e [71], f [52], g [76], h [24]

BIBLIOGRAPHY

[1] Pieter Agten, Wouter Joosen, Frank Piessens, and Nick Nikiforakis. "Seven Months' Worth of Mistakes: A Longitudinal Study of Typosquatting Abuse." In: *Proceedings 2015 Network and Distributed System Security Symposium*. February. Reston, VA: Internet Society, 2015, pp. 8–11. ISBN: 1-891562-38-X. DOI: 10.14722/ndss.2015.23058.

[2] Fatemah Mordhi Alharbi, Yuchen Zhou, Feng Qian, Zhiyun Qian, and Nael Abu Ghazaleh. "DNS Poisoning of Operating System Caches: Attacks and Mitigations." In: *IEEE Transactions on Dependable and Secure Computing* 19.4 (2022), pp. 2851–2863. ISSN: 1545-5971. DOI: 10.1109/tdsc.2022.3142331.

[3] Marios Anagnostopoulos, Georgios Kambourakis, Panagiotis Kopanos, Georgios Louloudakis, and Stefanos Gritzalis. "DNS amplification attack revisited." In: *Computers & Security* 39 (Nov. 2013), pp. 475–485. ISSN: 01674048. DOI: 10.1016/j.cose.2013.10.001.

[4] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *DNS Security Introduction and Requirements*. Tech. rep. IETF, Mar. 2005. DOI: 10.17487/rfc4033.

[5] Giuseppe Ateniese and Stefan Mangard. "A new approach to DNS security (DNSSEC)." In: *Proceedings of the ACM Conference on Computer and Communications Security* (2001), pp. 86–95. ISSN: 15437221. DOI: 10.1145/501983.501996.

[6] Zhouyu Bao, Wenbo Wang, and Yuqing Lan. "Using Passive DNS to Detect Malicious Domain Name." In: *ACM International Conference Proceeding Series* (2019). DOI: 10.1145/3387168.3387236.

[7] Steven M Bellovin. "Using the Domain Name System for System Break-ins." In: *UNIX Security Symposium*. Fifth. USENIX Association, 1995. URL: https://www.usenix.org/legacy/publications/library/proceedings/security95/full_papers/bellovin.pdf.

[8] Kevin Borgolte, Tithi Chattopadhyay, Nick Feamster, Mihir Kshirsagar, Jordan Holland, Austin Hounsel, and Paul Schmitt. "How DNS over HTTPS is Reshaping Privacy, Performance, and Policy in the Internet Ecosystem." In: *SSRN Electronic Journal* (2019). ISSN: 1556-5068. DOI: 10.2139/ssrn.3427563.

[9] Reg Brennenraedts, Melvin Hanswijk, Roos Jansen, Jessica Kats, Wazir Sahebali, and Leonie Hermanussen. *Evaluatie van de opbouw en meetbaarheid van de Nederlandse Cybersecurity Agenda*. Tech. rep. Utrecht: WODC, 2021. DOI: 20.500.12832/3065.

[10] Daiki Chiba, Takeshi Yagi, Mitsuaki Akiyama, Toshiki Shibahara, Takeshi Yada, Tatsuya Mori, and Shigeki Goto. "DomainProfiler: Discovering Domain Names Abused in Future." In: *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, June 2016, pp. 491–502. ISBN: 978-1-4673-8891-7. DOI: 10.1109/DSN.2016.51.

[11] A. Costello. *Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)*. Tech. rep. IETF, Mar. 2003. DOI: 10.17487/rfc3492.

[12] *DomainKeys Identified Mail (DKIM) Signatures*. Tech. rep. IETF, Sept. 2011. DOI: 10.17487/rfc6376.

[13] DPC. *Merkbewaking*. URL: https://www.communicatierijk.nl/vakkennis/rijkswebsites/aanbevolen-richtlijnen/domeinnaambewakingsservice-dbs (visited on 08/12/2022).

[14] S. Dickinson, D. Gillmor, and T. Reddy. *Usage Profiles for DNS over TLS and DNS over DTLS*. Tech. rep. IETF, Mar. 2018. DOI: 10.17487/RFC8310.

[15] Dienst Publiek en Communicatie. *Informatie Rijkswebsites*. URL: https://www.communicatierijk.nl/vakkennis/rijkswebsites.

[16] Directie Digitale Overheid. *Kamerbrief Voortgang informatieveiligheid bij de overheid*. 2021. URL: https://www.rijksoverheid.nl/documenten/kamerstukken/2021/03/18/kamerbrief-voortgang-informatieveiligheid-overheid.

[17] Domain Research Group. *Emoji Domain Registration*. URL: https://xn--i-7iq.ws/ (visited on 08/10/2022).

[18] Arthur Drichel, Vincent Drury, Justus Von Brandt, and Ulrike Meyer. *Finding Phish in a Haystack: A Pipeline for Phishing Classification on Certificate Transparency Logs*. Vol. 1. 1. Association for Computing Machinery, 2021. ISBN: 9781450390514. DOI: 10.1145/3465481.3470111. arXiv: 2106.12343.

[19] Dutch Internet Standards Platform. *Internet.nl - Is your internet up to date?* URL: https://internet.nl.

[20] Wolfgang Ebbers, Bob Hulsebosch, and Martijn Oostdijk. *Een Top Level Domein voor betrouwbare overheidscommunicatie*. Tech. rep. Novay, 2013. URL: https://www.kennisopenbaarbestuur.nl/documenten/rapporten/2013/02/26/een-top-level-domein-voor-betrouwbare-overheidscommunicatie.

[21] *The Unicode Code Points and Internationalized Domain Names for Applications (IDNA)*. Tech. rep. IETF, Aug. 2010. DOI: 10.17487/rfc5892.

[22] Edona Fasllija, Hasan Ferit Enişer, and Bernd Prünster. "Phish-Hook: Detecting Phishing Certificates Using Certificate Transparency Logs." In: *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, LNICST* 305 LNICST (2019), pp. 320–334. ISSN: 18678211. DOI: 10.1007/978-3-030-37231-6_18.

[23] *HTTP Semantics*. Tech. rep. IETF, June 2022. DOI: 10.17487/RFC9110.

[24] Forum Standaardisatie. *Lijst Open Standaarden*. URL: https://www.forumstandaardisatie.nl/open-standaarden.

[25] Forum Standaardisatie. *Regie op Internetdomeinen*. URL: https://magazine.forumstandaardisatie.nl/regie-op-internetdomeinen/regie-op-internetdomeinen.

[26] Forum Standaardisatie. *Meting Informatieveiligheidstandaarden overheid maart 2020*. Tech. rep. 2020. URL: https://zoek.officielebekendmakingen.nl/blg-942649.pdf.

[27] J. Gori Mohamed and J. Visumathi. "A predictive model of machine learning against phishing attacks and effective defense mechanisms." In: *Materials Today: Proceedings* (Oct. 2020). ISSN: 22147853. DOI: 10.1016/j.matpr.2020.09.612.

[28] M van de Graaf. *Domeinnaambeleid Rijksoverheid*. 2014. URL: https://www.communicatierijk.nl/binaries/communicatierijk/documenten/publicaties/2015/01/08/domeinnaambeleid-rijksoverheid/domeinnaambeleid-rijksoverheid-06022014.pdf.

[29] Valentijn Grapperhaus. *Handreiking Beheer Internetdomeinen Rijksoverheid*. Tech. rep. Rijksprogramma voor Duurzame Digitale Informatiehuishouding, 2021. URL: https://www.informatiehuishouding.nl/projecten/opschonen-websites/Producten+%26+publicaties/publicaties/2021/07/01/handreiking-beheer-internetdomeinen-rijksoverheid.

[30] P. Hallam-Baker, R. Stradling, and J. Hoffman-Andrews. *DNS Certification Authority Authorization (CAA) Resource Record*. Tech. rep. IETF, Nov. 2019. DOI: 10.17487/RFC8659.

[31] Michiel Henneke. *Toekomst .tv-domein op lange termijn onzeker*. 2021. URL: https://www.sidn.nl/nieuws-en-blogs/toekomst-tv-domein-op-lange-termijn-onzeker.

[32] P. Hoffman and P. McManus. *DNS Queries over HTTPS (DoH)*. Tech. rep. IETF, Oct. 2018. DOI: 10.17487/RFC8484.

[33] Tobias Holgers, David E. Watson, and Steven D. Gribble. "Cutting through the confusion: A measurement study of homograph attacks." In: *USENIX 2006 Annual Technical Conference*. 2006, pp. 261–266. URL: https://www.usenix.org/legacy/events/usenix06/tech/full_papers/holgers/holgers.pdf.

[34] Andries Hoogerwerf and Michiel Herweijer. *Overheidsbeleid: een inleiding in de beleidswetenschappen*. Alphen aan den Rijn: Kluwer, 1989. ISBN: 9060923057.

[35] Z. Hu, L. Zhu, J. Heidemann, A. Mankin, D. Wessels, and P. Hoffman. *Specification for DNS over Transport Layer Security (TLS)*. Tech. rep. IETF, May 2016. DOI: 10.17487/RFC7858.

[36] Walter Hulsker. *Een herkenbare en betrouwbare digitale overheid*. Tech. rep. ECORYS, 2020. URL: https://www.rijksoverheid.nl/documenten/rapporten/2020/07/08/een-herkenbare-en-betrouwbare-digitale-overheid.

[37] Matthew Humphries. *Here's Why Porn Is Appearing on News Websites Across the Web*. July 2021. URL: https://www.pcmag.com/news/heres-why-porn-is-appearing-on-news-websites-across-the-web.

[38] IANA. *Root Zone Management*. URL: https://www.iana.org/domains/root (visited on 07/19/2022).

[39] ICANN. *Centralized Zone Data Service (CZDS)*. URL: https://czds.icann.org/home.

[40] ICANN. *IDN Top-level domains*. URL: https://newgtlds.icann.org/en/about/idns (visited on 08/10/2022).

[41] ICANN. *Fees and Timelines*. 2012. URL: https://newgtlds.icann.org/en/applicants/global-support/faqs/faqs-en#fees (visited on 08/10/2022).

[42] ICANN. *Emojis in Domain Names: A Security Risk for Everyone*. Tech. rep. 2019. URL: https://www.icann.org/en/system/files/files/idn-emojis-domain-names-13feb19-en.pdf.

[43] ICANN. *TLD DNSSEC Report (2022-06-22 00:12:12)*. 2022. URL: http://stats.research.icann.org/dns/tld_report/.

[44] IETF. *The Internet Engineering Task Force*. URL: https://www.ietf.org/.

[45] International Telecommunication Union. *Individuals using the Internet (percentage of population)*. 2020. URL: https://data.worldbank.org/indicator/IT.NET.USER.ZS.

[46] Joost Schellevis. *Overheid vraagt twee eigen top-level-domeinen aan - update*. 2012. URL: https://tweakers.net/nieuws/82552/overheid-vraagt-twee-eigen-top-level-domeinen-aan.html.

[47] Panagiotis Kintis, Najmeh Miramirkhani, Charles Lever, Yizheng Chen, Rosa Romero-Gómez, Nikolaos Pitropakis, Nick Nikiforakis, and Manos Antonakakis. "Hiding in Plain Sight: A Longitudinal Study of Combosquatting Abuse." In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. August. New York, NY, USA: ACM, Oct. 2017, pp. 569–586. ISBN: 9781450349468. DOI: 10.1145/3133956.3134002. arXiv: 1708.08519.

[48] S. Kitterman. *Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1*. Tech. rep. IETF, Apr. 2014. DOI: 10.17487/rfc7208.

[49] *Domain-based Message Authentication, Reporting, and Conformance (DMARC)*. Tech. rep. Independent, Mar. 2015. DOI: 10.17487/rfc7489.

[50] Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff. "A brief history of the internet." In: *ACM SIGCOMM Computer Communication Review* 39.5 (Oct. 2009), pp. 22–31. ISSN: 0146-4833. DOI: 10.1145/1629607.1629613.

[51] Wenfeng Liu, Yu Zhang, Lu Liu, Shuyan Liu, Hongli Zhang, and Binxing Fang. "A secure domain name resolution and management architecture based on blockchain." In: *Proceedings - IEEE Symposium on Computers and Communications* (2020). ISSN: 15301346. DOI: 10.1109/ISCC50000.2020.9219632.

[52] Logius. *Register Toegankelijksverklaringen*. URL: https://toegankelijkheidsverklaring.nl/ (visited on 08/12/2022).

[53] Chaoyi Lu et al. "From WHOIS to WHOWAS: A Large-Scale Measurement Study of Domain Registration Privacy under the GDPR." In: *Proceedings 2021 Network and Distributed System Security Symposium*. February. Reston, VA: Internet Society, 2021. ISBN: 1-891562-66-5. DOI: 10.14722/ndss.2021.23134.

[54] Pin Lv, Lingling Bai, Tingwen Liu, Zhenhu Ning, Jinqiao Shi, and Binxing Fang. "Detection of Malicious Domain Names Based on Hidden Markov Model." In: *2018 IEEE Third International Conference on Data Science in Cyberspace (DSC)*. IEEE, June 2018, pp. 659–664. ISBN: 978-1-5386-4210-8. DOI: 10.1109/DSC.2018.00105.

[55] Douglas C. MacFarland, Craig A. Shue, and Andrew J. Kalafut. "Characterizing Optimal DNS Amplification Attacks and Effective Mitigation." In: *Lecture Notes in Computer Science*. Vol. 8995. 2015, pp. 15–27. ISBN: 9783319155081. DOI: 10.1007/978-3-319-15509-8_2.

[56] Douglas C. MacFarland, Craig A. Shue, and Andrew J. Kalafut. "The best bang for the byte: Characterizing the potential of DNS amplification attacks." In: *Computer Networks* 116 (Apr. 2017), pp. 12–21. ISSN: 13891286. DOI: 10.1016/j.comnet.2017.02.007.

[57] Cleborne D. Maddux and D. LaMont Johns. "The World Wide Web: History, Cultural Context, and a Manual for Developers of Educational Information-Based Web Sites." In: *Educational Technology* 37.5 (1997), pp. 5–12. URL: https://www.jstor.org/stable/44428413.

[58] Samuel Marchal, Jérôme François, Radu State, and Thomas Engel. "Proactive discovery of phishing related domain names." In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 7462 LNCS. 2012, pp. 190–209. ISBN: 9783642333378. DOI: 10.1007/978-3-642-33338-5_10.

[59] John Markwell and David W. Brooks. ""Link rot" limits the usefulness of web-based educational materials in biochemistry and molecular biology." In: *Biochemistry and Molecular Biology Education* 31.1 (2003), pp. 69–72. ISSN: 14708175. DOI: 10.1002/bmb.2003.494031010165.

[60] Saif Al Mashhadi and Selvakumar Manickam. "A brief review of blockchain-based DNS systems." In: *International Journal of Internet Technology and Secured Transactions* 10.4 (2020), p. 420. ISSN: 1748-569X. DOI: 10.1504/IJITST.2020.108134.

[61] Nicole van der Meulen. "DigiNotar: Dissecting the First Dutch Digital Disaster." In: *Journal of Strategic Security* 6.2 (June 2013), pp. 46–58. ISSN: 1944-0464. DOI: 10.5038/1944-0472.6.2.4.

[62] P.V. Mockapetris. *Domain names - concepts and facilities*. Tech. rep. IETF, Nov. 1987. DOI: 10.17487/rfc1034.

[63] P.V. Mockapetris. *Domain names - implementation and specification*. Tech. rep. IETF, Nov. 1987, pp. 1–55. DOI: 10.17487/rfc1035. URL: http://tools.ietf.org/html/rfc1035.

[64] Tyler Moore and Benjamin Edelman. "Measuring the Perpetrators and Funders of Typosquatting." In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 6052 LNCS. January. 2010, pp. 175–191. ISBN: 3642145760. DOI: 10.1007/978-3-642-14577-3_15.

[65] Nationaal Archief. *Richtlijn archiveren overheidswebsites*. Tech. rep. 2018. URL: https://www.nationaalarchief.nl/archiveren/kennisbank/Richtlijn-Archiveren-Overheidswebsites.

[66] Stephen J Nightingale. *Email authentication mechanisms: DMARC, SPF and DKIM*. Tech. rep. Gaithersburg, MD: National Institute of Standards and Technology, 2017. DOI: 10.6028/NIST.TN.1945.

[67] Nick Nikiforakis, Steven Van Acker, Wannes Meert, Lieven Desmet, Frank Piessens, and Wouter Joosen. "Bitsquatting: Exploiting bit-flips for fun, or profit?" In: *WWW 2013 - Proceedings of the 22nd International Conference on World Wide Web* (2013), pp. 989–998. DOI: 10.1145/2488388.2488474.

[68] Nick Nikiforakis, Marco Balduzzi, Lieven Desmet, Frank Piessens, and Wouter Joosen. "Soundsquatting: Uncovering the Use of Homophones in Domain Squatting." In: 2014, pp. 291–308. DOI: 10.1007/978-3-319-13257-0_17.

[69] Overheid.nl. *Register van Overheidsorganisaties*. URL: https://almanak.overheid.nl/.

[70] PBLQ. *Buitenlandonderzoek Domeinnaambeleid*. Tech. rep. 2019. URL: https://www.digitaleoverheid.nl/document/buitenlandonderzoek-domeinbeleid/.

[71] Willem Pieterson. *Oordeel Burgers en Ondernemers over Overheidsdienstverlening*. Tech. rep. Pieterson, Ltd, Kantar, 2020. URL: https://www.rijksoverheid.nl/documenten/rapporten/2020/12/07/eindrapport-onderzoek-overheidsdienstverlening-2020.

[72] Joost Prins. *Proactive Recognition of Domain Abuse*. 2021. URL: https://purl.utwente.nl/essays/84073.

[73] RDDI. *Modeltoets voldoen aan overige eisen Rijkswebsites*. 2020. URL: https://www.informatiehuishouding.nl/Producten+%26+publicaties/instrumenten/2020/02/06/modeltoets-voldoen-aan-overige-eisen-websitearchivering.

[74] S. M.Zia Ur Rashid, Md Imtiaz Kamrul, and Asraful Islam. "Understanding the Security Threats of Esoteric Subdomain Takeover and Prevention Scheme." In: *2nd International Conference on Electrical, Computer and Communication Engineering, ECCE 2019* (2019), pp. 7–9. DOI: 10.1109/ECACE.2019.8679122.

[75] Rijksoverheid. *Domeinnaambeleid*. URL: https://domeinnaambeleid.nl (visited on 01/26/2022).

[76] Rijksoverheid. *Websiteregister*. 2021. URL: https://websiteregisterrijksoverheid.nl.

[77] Rijksvoorlichtingsdienst. *Online beleid: rijksbreed afwegingskader online middelen*. Tech. rep. september. 2020. URL: https://www.communicatierijk.nl/vakkennis/rijkswebsites/documenten/publicaties/2018/9/19/rijksbreed-afwegingskader-online-middelen.

[78] Roland van Rijswijk-Deij, Mattijs Jonker, Anna Sperotto, and Aiko Pras. "A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements." In: *IEEE Journal on Selected Areas in Communications* 34.6 (June 2016), pp. 1877–1888. ISSN: 0733-8716. DOI: 10.1109/JSAC.2016.2558918.

[79] SIDN Labs. *.nl statistieken DNSSEC*. URL: https://stats.sidnlabs.nl/nl/dnssec.html (visited on 02/12/2023).

[80] SIDN. *Nieuwe SAD DNS 'cache poisoning'-aanval op Domain Name System gepubliceerd*. 2021. URL: https://www.sidn.nl/nieuws-en-blogs/nieuwe-sad-dns-cache-poisoning-aanval-op-domain-name-system-gepubliceerd (visited on 12/13/2021).

[81] SIDN. *Verdere adoptie IPv6 in Nederland nagenoeg tot stilstand gekomen*. 2021. URL: https://www.sidn.nl/nieuws-en-blogs/verdere-adoptie-ipv6-in-nederland-nagenoeg-tot-stilstand-gekomen (visited on 11/30/2021).

[82] P. Saint-Andre and J. Hodges. *Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)*. Tech. rep. Mar. 2011. DOI: 10.17487/rfc6125.

[83] Giovanni Schmid. "Thirty Years of DNS Insecurity: Current Issues and Perspectives." In: *IEEE Communications Surveys and Tutorials* 23.4 (2021), pp. 2429–2459. ISSN: 1553-877X. DOI: 10.1109/COMST.2021.3105741.

[84] Hossein Shirazi, Bruhadeshwar Bezawada, and Indrakshi Ray. “"Know Thy Doma1n Name".” In: *Proceedings of the 23nd ACM on Symposium on Access Control Models and Technologies*. New York, NY, USA: ACM, June 2018, pp. 69–75. ISBN: 9781450356664. DOI: 10.1145/3205977.3205992.

[85] Raffaele Sommese, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, K.C. Claffy, and Anna Sperotto. "The Forgotten Side of DNS: Orphan and Abandoned Records." In: *2020 IEEE European Symposium on Security and Privacy Workshops*. IEEE, Sept. 2020, pp. 538–543. ISBN: 978-1-7281-8597-2. DOI: 10.1109/EuroSPW51379.2020.00079.

[86] Janos Szurdi, Balazs Kocso, and Gabor Cseh. "The Long "Taile" of Typosquatting Domain Names." In: *Usenix.Org* (2014), p. 191. URL: https://www.usenix.org/system/files/conference/usenixsecurity14/sec14-paper-szurdi.pdf.

[87] Kristel van Teeffelen. *Politie is slordig met het afdekken van oude websites*. Jan. 2017. URL: https://www.trouw.nl/nieuws/politie-is-slordig-met-het-afdekken-van-oude-websites~b238a9f1/.

[88] 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." In: *Computer Science Review* 45 (2022), p. 100469. ISSN: 15740137. DOI: 10.1016/j.cosrev.2022.100469.

[89] Sadegh Torabi, Amine Boukhtouta, Chadi Assi, and Mourad Debbabi. "Detecting internet abuse by analyzing passive DNS traffic: A survey of implemented systems." In: *IEEE Communications Surveys and Tutorials* 20.4 (2018), pp. 3389–3415. ISSN: 1553877X. DOI: 10.1109/COMST.2018.2849614.

[90] Tweede Kamer der Staten-Generaal. *Wetgeving voor de elektronische snelweg*. 2000. URL: https://zoek.officielebekendmakingen.nl/kst-25880-11.pdf.

[91] Henk Van de Graaf and Rob Hoppe. *Een inleiding tot de beleidswetenschap en de beleidskunde*. Bussum: Couthino, 1992. ISBN: 9789062837588.

[92] Anne-Sofie Vanhaeght. "Assessing Policy IV: Goal-Means Tree Analysis." In: *The Palgrave Handbook of Methods for Media Policy Research*. Cham: Springer International Publishing, 2019, pp. 595–608. ISBN: 9783030160647. DOI: 10.1007/978-3-030-16065-4_34.

[93] Daniel Verlaan. *Groot datalek bij Jeugdriagg: medische dossiers kwetsbare kinderen gelekt*. Oct. 2020. URL: https://www.rtlnieuws.nl/nieuws/nederland/artikel/5187220/jeugdriagg-kenter-jeugdhulp-datalek-dossiers.

[94] David L. Weimer and Aidan R. Vining. *Policy Analysis*. 6th ed. New York: Routledge, Mar. 2017. ISBN: 9781315442129. DOI: 10.4324/9781315442129.

[95] *DNS Privacy Considerations*. Tech. rep. IETF, July 2021. DOI: 10.17487/RFC9076.

[96] Hongtao Yi and John T. Scholz. "Policy Networks in Complex Governance Subsystems: Observing and Comparing Hyperlink, Media, and Partnership Networks." In: *Policy Studies Journal* 44.3 (2016), pp. 248–279. ISSN: 15410072. DOI: 10.1111/psj.12141.

[97] Johannes Zirngibl, Steffen Deusch, Patrick Sattler, Juliane Aulbach, Georg Carle, and Mattijs Jonker. "Domain Parking: Largely Present, Rarely Considered!" In: *Network Traffic Measurement and Analysis Conference*. IFIP, 2022. ISBN: 9783903176478.

[98] Auke Zwaan. *Detecting malicious domain names using spatial co-occurrence in DNS traffic*. 2016. URL: https://www.sidnlabs.nl/downloads/Zn4ziXWIT8ibSaRlhMqj_Q/816c6d7597dd8ac99d16765f3fb545e4/RP65-AukeZwaan-MaliciousDomainNameDetectionSystem.pdf.

[99] Root zone changes. *Deletion of an*. 2015. URL: https://web.archive.org/web/20220409170835/https://twitter.com/diffroot/status/628272650108932096.