Challenges in Decentralized Name Management: The Case of ENS

Pengcheng Xia
Beijing University of Posts and Telecommunications
China

Haoyu Wang*
School of Cyber Science and Engineering, Huazhong University of Science and Technology
China

Zhou Yu
Xinyu Liu
Beijing University of Posts and Telecommunications
China

Xiapu Luo
The Hong Kong Polytechnic University
China

Guoai Xu
Beijing University of Posts and Telecommunications
China

Gareth Tyson
The Hong Kong University of Science and Technology (GZ)
China

ABSTRACT

DNS has often been criticized for inherent design flaws, which make the system vulnerable to attack. Further, domain names are not fully controlled by users, meaning that they can easily be taken down by authorities and registrars. Due to this, there have been efforts to build a decentralized name service that gives greater control to domain owners. The Ethereum Name Service (ENS) is a major example. Yet, no existing work has systematically studied this emerging system, particularly regarding security and misbehavior. To address this gap, we present the first large-scale measurement study of ENS. Our findings suggest that ENS has shown growth during its four years' evolution. We identify several security issues, including traditional name system problems, as well as new issues introduced by the unique properties of ENS. We find that attackers are abusing the system with thousands of squatting ENS names, a number of scam blockchain addresses and indexing of malicious websites. We further develop a new record persistence attack, to find that 22,716 .eth names (3.7% of all names) are vulnerable to name hijacking. Our exploration suggests that our community should invest more effort into the detection and mitigation of issues in decentralized name services.

CCS CONCEPTS

• **Security and privacy** → **Intrusion/anomaly detection and malware mitigation**; **Web application security**; • **Information systems** → **Web mining**.

KEYWORDS

naming system; blockchain; decentralization; security

*Corresponding author: Haoyu Wang (haoyuwang@hust.edu.cn).

IMC '22, October 25–27, 2022, Nice, France

ACM ISBN 978-1-4503-9259-4/22/10…$15.00
https://doi.org/10.1145/3517745.3561469

ACM Reference Format:
Pengcheng Xia, Haoyu Wang, Zhou Yu, Xinyu Liu, Xiapu Luo, Guoai Xu, and Gareth Tyson. 2022. Challenges in Decentralized Name Management: The Case of ENS. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22), October 25–27, 2022, Nice, France.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3517745.3561469

1 INTRODUCTION

The domain name system (DNS) is an indispensable component of the Internet, translating human-readable domain names to numerical IP addresses. DNS is built in a hierarchical and distributed way. However, it has often been criticized for inherent design flaws that make the system vulnerable to various attacks [106]. For instance, a recent report suggests that a group of hackers have launched DNS hijacking attacks against at least 30 organizations, including government ministries, embassies and security services as well as companies and other groups in Europe and the Middle East [22].

Thus, many research efforts have attempted to make DNS more secure and reliable [66, 68, 69]. For example, the Domain Name System Security Extensions (DNSSEC) [68] is a set of extensions to support cryptographic authentication of DNS data, authenticated denial of existence, and data integrity. However, DNSSEC does not provide confidentiality, which means that DNSSEC responses are authenticated but not encrypted. Besides, the signing and checking of digital signature increases query latency and affects user experience. Moreover, the complexity of implementing and maintaining DNSSEC has impeded its adoption [63].

With these limitations in mind, some have proposed ways to address issues of traditional DNS by combining blockchain concepts with DNS: so-called Blockchain Naming Systems (BNS). Since blockchain has unique properties such as immutability and decentralization, it is argued that such approaches can improve both resilience and security. For example, Namecoin [38] claims to be the first blockchain-based DNS solution and it is a fork of the Bitcoin network that offers a new `.bit` top-level domain (TLD) for its names. Similar to Namecoin, UnstoppableDomains [44] and EmerDNS [16] propose new TLDs like `.zil` , `.crypto` and `.emc`, etc with guaranteed ownership. Handshake [26] takes another approach, in which it attempts to replace the DNS root with a more decentralized, secure system. The names on these BNS services are exclusively managed by their owners, and they cannot be taken

down by authorities (e.g., the government). Despite this, the advantages of BNS can be exploited for malicious purposes. According to a recent study [80], blockchain domain names have been exploited as command and control (C&C) channels by attackers. Furthermore, domain squatting issues [99] will be more severe than the normal DNS, due to the difficulty of shutting down malicious domains.

One particularly prominent BNS is the Ethereum Name Service (ENS) [17]. This is a decentralized naming service built atop of Ethereum [28], one of the most popular blockchain technologies that allows users to create Decentralized Applications (dApps) by developing smart contracts. In contrast to the aforementioned BNS solutions, instead of taking the place of DNS, it focuses foremost on resolving its blockchain domain names (short for "ENS names") to `web3` resources like blockchain addresses and Decentralized Websites (dWebs, web pages hosted on decentrlized servers.). Its service takes advantage of smart contracts on Ethereum to manage the registration and resolution of blockchain domain names. In other words, blockchain users can use their blockchain addresses to interact with ENS contracts. They must pay a certain amount of Ether for using an ENS domain name over a period of time, which can be used to resolve designated blockchain addresses and other resources. According to the official announcement [17], ENS has been integrated into more than 500 popular services including blockchain wallets, dApps and browsers. For example, browsers like Chrome and Firefox have extensions to resolve ENS domain names when users type them into browsers directly. Therefore, ENS has become one of the most popular blockchain name systems in-the-wild [50].

Although ENS has been deployed for over 4 years, to the best of our knowledge, it has not yet been systematically studied. Consequently, we lack empirical insight into the operations and challenges of ENS, particularly related to security. For instance, there remains a number of unexplored questions such as *How many domain names are registered in ENS? How do people use ENS? And whether security issues and abuse is prevalent in ENS?*

This work. In this paper, we take the first steps to systematically characterize ENS. To fully understand the registration and resolution process, we fetch and analyze all event logs (7.7 million) of ENS-related smart contracts and third party resolver contracts (**§4**). By decoding these logs, we harvest 617, 250 registered names and 184, 490 Ethereum addresses. Based on this dataset, we perform a detailed temporal analysis (**§5**). We observe that 55.6% of all ENS domain names are active (i.e. unexpired second-level domain names, or unexpired 2LD names for short, and other subdomain names) and 83.4% of ENS users are active by the time of our study (i.e., Ethereum addresses that have ever had an ENS name still have at least one name). We then explore the use of ENS names and find that over 85.8% of record settings are related to blockchain addresses (**§6**). To further examine the security of ENS, we adopt a series of measurement techniques (**§7**) to investigate both traditional security issues (i.e., domain name squatting, malicious domain indexing and scam addresses) and then ENS specific security issues (i.e., record persistence attack). We obtain the following key findings:

- **ENS has experienced notable growth during its four years' evolution (§5).** Over 617K ENS names have been registered, and 341K of them are active by the time of this study. We find a number of users willing to pay high prices for rare ENS names, as well as others who obtain large numbers of names.
- **ENS is an open system where domain names index a wide range of records (§6).** The most common use of an ENS name is to link to blockchain addresses (85.8% of the record settings). Further, it is common to use ENS names for decentralized websites. In addition to blockchain addresses and decentralized websites, many other kinds of records like public key records and text records are found in ENS resolvers, highlighting the diversity of ENS. Curiously, we observe people exploring new ways to interact using ENS text records, e.g., using them as accounts for a P2P database or for decentralized voting.
- **The open nature of ENS makes it easy to be abused by attackers (§7).** We identify several security issues and misbehaviors, including traditional DNS security issues and new ones introduced by ENS smart contracts. A few squatters are found to be hoarding famous brand names and their variants (which could be used for malicious purposes). Some malicious decentralized websites and scam addresses are also found in ENS name records. Beyond these traditional name system attacks, we develop a new record persistence attack and find that 22, 716 `.eth` names (3.7% of all names) are vulnerable to hijacking.

To the best of our knowledge, this is the first comprehensive study of ENS. Our results motivate the need for more research to illuminate widely unexplored BNS systems. We believe that our efforts can attract the focus of the research community and promote best operational practices. We will release our dataset, along with the experimental results: https://ensnames.github.io/ensnames/.

2 BACKGROUND

2.1 DNS and Zooko's Triangle

2.1.1 Domain Name System (DNS). DNS is a network protocol that associates domain names with various information. The DNS is distributed throughout the Internet in a hierarchical authority. A typical DNS resolution process is shown in Figure 1. When a client wants to query the IP of a domain name, it first queries the recursive resolver. If the resolver previously resolved the record and has it in its cache, it will return the result immediately. Otherwise, it will query the root servers which has the information of top-level domains (TLDs), and get the information of the relevant TLD server. Then, a query will be made to the TLD server for the second-level

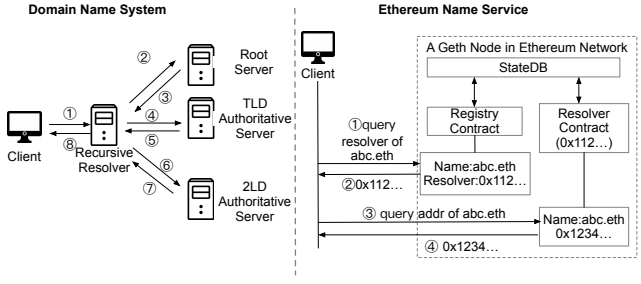

Figure 1: An overview of DNS and ENS operations.

domain (2LD) authoritative server. After receiving the request, the 2LD server responds the corresponding IP address to the resolver.

DNS has been criticized, due to various vulnerabilities and security issues. For example, due to the lack of authentication and integrity checks, attacks like cache poisoning and DNS tunneling were prevalent [106]. Further, DNS domain names are not fully controlled by users and can be taken down by authorities. Some efforts like DNS over HTTPS (DoH), DNS over TLS (DoT) and DNSSEC have been designed to alleviate some of these issues [66–69], although they are yet to reach the ubiquity of DNS [72, 87, 90]. People also use other solutions like Tor or the Invisible Internet Project (I2P) for privacy [34, 43]. Nevertheless, DNS and the aforementioned solutions still cannot achieve *human-readability*, *security* and *decentralization* simultaneously, i.e, Zooko's Triangle.

2.1.2 Zooko's Triangle. Zooko's Triangle [56] defines three properties that an ideal name system should possess: 1) *human-meaningful*, i.e., the names should be readable and memorable by humans; 2) *secure*, i.e., names should be translated correctly even when the system is attacked; 3) *decentralised*, i.e., names should be translated without a central authority. Zooko Wilcox-O'Hearn, the creator of Zcash [59], speculated that any name system could only achieve two of these three properties at most. This triangle has been used to evaluate name system performance in many other literature [71, 82, 97]. For example, DNS names are human-readable but the DNS system is not secure and decentralized. Tor is secure and decentralized but its addresses are not human-readable. I2P addresses are human-readable and secure, but I2P is not decentralized. With the rise of blockchain technologies, some blockchain-based name systems have been proposed, which claim to fulfill all three properties in Zooko's Triangle [16, 26, 38, 44]. Whether ENS truly solves the paradox constitutes one of the motivations for our study.

2.2 Blockchain Name Service and Ethereum Name Service

2.2.1 Blockchain Name Service (BNS). Since blockchain offers immutability and decentralization, some blockchain-based name services have been proposed in recent years. The most common purpose of BNS is to replace the traditional DNS with blockchain-based alternatives. They usually propose new TLDs that are incompatible with the traditional DNS and the Internet Corporation for Assigned Names and Numbers (ICANN). For example, Namecoin [38] is the first blockchain-based DNS, which claims to be the solution of Zooko's Triangle. It proposes the `.bit` TLD for users and has the ability to attach identity information or human-meaningful Tor domains. Similarly, UnstoppableDomains [44] and EmerDNS [16] also propose some new TLDs. Handshake [26] is another kind of BNS, which seeks to replace the DNS root with its decentralized, secure system.

2.2.2 Ethereum Name Service (ENS). Unlike the above BNS attempts, ENS focus foremost on translating its blockchain domain names to on-chain resources. It is built on top of Ethereum, a shared and distributed ledger that enables recording transactions and tracking assets in a P2P network. Based on a cryptographic design, each transaction in the block is verified by the confirmation of most participants in the system. Ethereum also supports smart contracts — a computer program that runs atop of the Ethereum blockchain [51]. It contains a set of code (functions) and data (state). When a smart contract's predetermined conditions are met, it runs automatically to change the states of involved participants. A user who controls an Ethereum address can interact with a smart contract by sending transactions that execute the functions defined within the contract. Invoking a function (i.e., performing a transaction) will cost a fee (termed "gas"), as a reward for the miner who deals with the transaction. One exception is where the functions does not change the state of involved participants (like read-only functions). These are termed "external view" functions, which do not make transactions and do not need gas fee.

ENS is controlled by several such smart contracts, and users can call the interfaces provided by these contracts to register and manage names. Among all the contracts, the multi-signature wallet contract [29] controlled by ENS core members can make changes to the whole system when all members agree. ENS mainly consists of three kinds of contracts: the registry, the registrars and the resolvers [33]. We detail these below.

(1) The *Registry* stores the mapping of ENS names (of any level) to owners, resolvers and the caching time-to-live (TTL) for ENS name records. In order to avoid trivial enumerations of names during the initial auction (see §3.1) and to map names to a fixed length identifier, ENS stores names in the form of hashes, which are generated through the process named "namehash". The namehash can be calculated by combining the hash of the highest-level part of ENS domain names (called "labelhash") with the namehash of the other part, and then performing a hash again on it.[1] This algorithm preserves the hierarchical properties of ENS names. We describe how we restore ENS names from hash values in §4.2.

(2) *Registrar* is a kind of smart contract that owns a name, and can automatically assign subdomain names to users based on some rules (e.g., payment). The ENS development team has used various registrar contracts for `.eth` name registrations, including the Vickrey auction registrar and the permanent registrar (see §3). Along with the permanent registrar, the concept of the registrar controller was introduced to delegate the name management of name owners. We detail the registration process in §3. In particular, users whose DNS names are supported by ENS can claim their DNS names in ENS by proving the ownership through DNSSEC and setting the TXT records containing their Ethereum addresses [57]. TLDs like `.kred` and `.luxe` can be linked with owners' Ethereum accounts directly in their DNS registrars [32].

(3) The *Resolver* stores the mapping of names to records. ENS can store arbitrary records while the "public resolvers" implemented by the ENS team have eight predefined types of records (see Table 1).

As shown in Figure 1, the ENS name resolution is a two-step process. The user who wants to resolve the name needs to query the registry to find the correct resolver and then get the resolution results from the resolver. Note that these queries are processed by external view functions, which do not cost gas and are not in the blockchain transaction list.

[1] namehash(test.eth) = keccak256(keccak256(test) + namehash(eth))

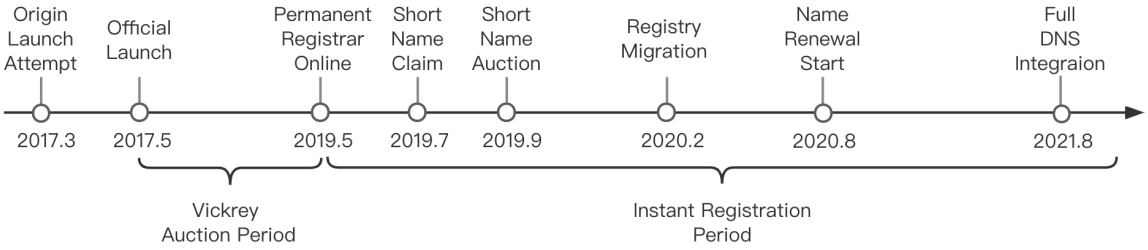

Figure 2: The timeline of major ENS events.

Table 1: The 8 record types in the public resolvers.

Record Type	Description
Address	Can be ETH address or other blockchain address
Name	Used for reverse resolution, i.e., mapping wallet addresses to ENS names
Content Hash	IPFS hash, Swarm hash for dWebs and Tor .onion address hash
Text	Key-value text record, and key can be "email", "URL", "vnd.twitter", etc.
DNS Record	DNS record in wire-format
Pubkey	ECDSA SECP256k1 public key
ABI	Application Binary Interface for interacting with contracts
Authorisation	Granting one address full access to one name except authorisations

3 THE EVOLUTION OF ENS

As ENS has gone through multiple development cycles, we start by detailing the timeline of ENS development. Figure 2 summarizes the timeline of ENS evolution, according to the official ENS blogs [19]. The ENS service was initially launched in March 2017. It encountered two severe bugs and the service went offline soon after the launch [20]. We next describe how, following this, ENS has managed its name registration process.

3.1 The Initial Auction (Vickrey Auction)

When ENS formally launched on May 4th 2017, the ENS team deployed a smart contract[2] implementing a Vickrey auction for registering names that have a length of more than 6. A *Vickrey auction* [64] is a type of sealed-bid auction where bidders submit their bids without knowing how much others have bid. The winner of the auction is the highest bidder, while they only need to pay the second-highest price. In an ENS Vickrey auction, `.eth` names are transferred into hashes (as depicted in §2.2.2) to avoid trivial enumeration; further, names are gradually released during an 8-week period. These schemes, to some extent, help people have a greater probability of obtaining the names they want. The Ether paid by a name's bidders will be deposited into a smart contract called a "deed" and all the losers of the auction will get a refund, less 0.5.%.[3] The winner of the name only needs to pay the second-highest price, and they could give up the ownership to withdraw all the Ether they paid after registration for one year. We perform a detailed analysis on the behavior of this auction period in § 5.2.

3.2 Permanent Registrar, Short Name Claim and Short Name Auction

3.2.1 Permanent Registrar. After two years of auction, the ENS team launched the "Permanent Registrar" for registering names over 6 characters in length instead of the auction registrar on May 4th 2019. This Permanent Registrar aims to run continuously (until any high-risk vulnerabilities are identified in the registrar contract). The charging method of `.eth` names follows an annual rental model, in which each name needs to be charged $5 per year based on the real-time exchange rate when the registration transaction occurs. Along with the permanent registrar, the concept of a "Registrar Controller" was introduced to delegate the name management of name owners. Thus, a name registered by the registrar controller can set the resolver and name records within the single registration transaction. This simplifies the registration process.

3.2.2 The Short Name Claim & Auction. In July 2019, the ENS development team opened the reservation of short `.eth` names (names with a length of 3-6 characters). This means that owners of eligible traditional TLD names can request corresponding `.eth` names and pay the rent in advance to obtain access to their corresponding `.eth` names for one year ($640 in ETH for a 3 character name, $160 for a 4 character name, $5 for a 5-6 character name). An owner of a short second-level traditional name registered on or before May 4th 2019 can claim one of the following names: 1) An exact match of the original name (e.g., `foo.com` to `foo.eth`). 2) Removing the `eth` suffix of original name (e.g., `fooeth.com` to `foo.eth`). 3) Combining the 2LD and TLD of the original name (`foo.com` to `foocom.eth`). Upon application, the ENS team will review the request for validity.

In September 2019, the auction for short names (with a length of 3-6) started. The ENS team chose OpenSea [39], a well-known crypto assets marketplace, as the auction platform, and used the *English auction* [102] as the sales method. In an English auction, bids are public and bidders can bid multiple times. The bidder who submits the highest price will win the name and the payment deposited will be the registration fee for the first year (which is quite different from the Vickrey auction period). After the short name auction, the remaining short names will be open for registration at a price based on their length. According to the analysis in §5.3, there were few DNS name owners claiming corresponding ENS names, while many famous brand names are selling for high price in the short name auctions. This could be related to squatting behaviors and we investigate this in §7.1.1.

3.3 Name Renewal Start

Since ENS introduced the "Permanent Registrar", expiration and renewal mechanisms were also introduced. Currently in ENS, all `.eth` names are charged annually based on their name length and anyone can renew no matter whether they own the name or not. Old names registered through the Vickrey auction, expired on May

[2] Address:0x6090a6e47849629b7245dfa1ca21d94cd15878ef (Etherscan label: "ENS: Old Registrar")

[3] The deed contract would burn 0.5% of the paid Ether in order to deter large numbers of registrations for capturing valuable names.

4th 2020 if not renewed. Besides, all `.eth` names have a 90-day grace period after expiration where payment can be made to retain ownership.

Since a large number of names were registered in the Vickrey auction period, most names would expire on May 4th 2020 (actually, August 2nd due to the 90-day grace period). To avoid squatting behaviors and gas competition for registration priority, the ENS team implemented a "decaying price premium" [10]. This was where the price of an expired name would start at $2,000 (besides normal annual rent), but would then decrease linearly to the normal annual rent after 28 days.

3.4 Full DNS Integration

As stated in §2.2.2, instead of creating more TLDs that are incompatible with DNS, ENS aims to integrate more DNS TLDs through DNSSEC. Over the 4 year-old history of ENS, it gradually supported 6 existing DNS TLDs and on August 26th, ENS launched the full DNS integration, which means that all 2LD domain name owners can import their DNS names to ENS. The DNS names have no protocol fee (e.g., the $5 annual fee on normal `.eth` names) and have the same records as `.eth` names. Specifically, DNS 2LD domain owners can claim their DNS names in ENS by proving the ownership through DNSSEC and setting the TXT records containing their Ethereum addresses [57]. However, since the ownership of a DNS name on ENS is based on DNSSEC, the security of DNS names on ENS depends on the security of these names on DNS.

4 STUDY DESIGN

We present the details of our measurement study on ENS in this section. We first describe the research questions, and then present how we collect the ENS data used for our study.

4.1 Research Questions

Our study aims to understand the status quo of the ENS ecosystem, and investigate its security issues. To this end, our study is driven by the following research questions (RQs):

- RQ1 **Popularity of ENS.** Considering ENS launched over 4 years ago and its purpose is to map on-chain resources, it is necessary to investigate its popularity, i.e., *how many domain names are registered, and how many addresses are involved in ENS?* The results shed light on the adoption level of ENS in the community. This RQ is answered in §5.
- RQ2 **Usage of ENS.** Considering the unique features of ENS, it is unknown *how people use ENS in-the-wild?* Therefore, we analyze the record types and information associated with ENS names. This RQ is answered in §6.
- RQ3 **Security Issues of ENS.** Since no prior work has analyzed security issues in the ENS ecosystem, it is important to understand *whether security issues (both traditional DNS and new emerging issues) are present in ENS*? The observations have implications for the design of future decentralized name services. This RQ is answered in §7.

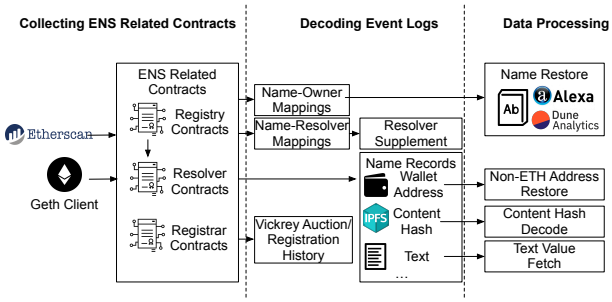

Figure 3: The workflow of our data collection.

4.2 Dataset Collection

It is non-trivial to harvest information related to ENS. First, ENS stores ENS domain names (ENS names for short) in the form of hash values, so we cannot get their human-readable names directly. Second, ENS has multiple resolvers (including third-party resolvers), which often use different types of protocols to encode records. Thus, we follow a hybrid workflow to extract comprehensive information on ENS, shown in Figure 3. Our dataset collection contains three major steps.

4.2.1 Step 1: Collecting ENS-related smart contracts. The first step is to collect all ENS official smart contracts, which are related to core functions of ENS (e.g., name registration and name renewal). We exploit Etherscan [21], a commonly used Ethereum explorer, to search for related contracts. Etherscan has labeled 28 ENS official smart contracts with human-meaningful names. For example, the contract used for the Vickrey auction[4] is labeled as the "Old Registrar". Note that some smart contracts are not related to the core functionalities of ENS, e.g., the multi-signature contract is used to effect administrative changes. Thus, we only focus on the aforementioned three types of smart contracts that are related to the resolution of ENS, including registry contracts, resolver contracts, registrar contracts (including registrar controller contracts and a short name claim contract). We manually analyze all these contracts and label 13 of them.

4.2.2 Step 2: Decoding Event Logs. After collecting related contracts, we take advantage of Geth [24], a well-known Ethereum client to synchronize the ledger of Ethereum. Specifically, to get the state changes of each contract, we extract event logs from the ledger. Event logs record the major activities of smart contracts and thus help track smart contracts' behaviors. An event log contains the event name and its related parameters, and will emit when predefined conditions meet. The conditions of these events are defined in contracts' codes and the format of event logs can be found in the Application Binary Interfaces (ABIs)[5] of smart contracts. For example, when a name is registered during the Vickrey Auction through the "Old Registrar" contract, a "HashRegistered" event will emit and record the name registered (in hashed format), the name owner, auction value and registration time in its event log. We summarize the details of all the events we fetched in Table 10 of the Appendix. Since ENS official contracts are open-sourced on Etherscan, we

[4] Address: 0x6090A6e47849629b7245Dfa1Ca21D94cd15878Ef
[5] ABI encodes how to interact with the functions of a smart contract and can be used to fetch event logs or function details from transaction receipt.

fetch the ABIs of each contract and decode event logs based on their ABIs. Thus, we get name-owner mappings and name-resolver mappings through registry contracts. We obtain a name record history from resolver contracts, and collect auction/registration history from registrar contracts. Furthermore, in name-resolver mappings, we find that many names point to additional resolvers. Thus, we further include 13 open-source extra resolvers that have more than 150 event logs. We then fetch their event logs, and decode them based on their ABIs. These additional resolvers are shown in Table 6 of the Appendix.

4.2.3 Step 3: Data Processing. Some information, like the Vickrey auction history, can be directly extracted from the event logs. However, some additional data processing tasks are still needed due to the design of these contracts. Specifically, we need to get the unhashed ENS names from name-owner mappings, restore non-ETH wallet addresses and content hashes based on their encoding rules and fetch values of text records from corresponding transactions.

As stated in §2.2.2, ENS smart contracts store hash values of ENS names instead of the names themselves. Thus, we take efforts to restore these hash values to readable names using three techniques. First, the ENS developers have uploaded their name-hash dictionary in the Vickrey auction to Dune Analytics [12], a platform for extracting Ethereum data, for sharing data to developers and researchers. We fetch this dictionary and update our database based on it. Second, we manually generate the labelhash based on a list of over 460K English words and 2LD of the Alexa [2] top-100K name list (downloaded on Sept. 14th 2021). We then match these predicted names with the hashes in the registry event logs (this data is also used in §7.1.1 for identifying squatting names). Third, the "NameRegistered" and "NameRenew" events of new ENS registrar controllers contain the plain text of newly registered names; we simply add these to our database.

For the address records, since non-ETH addresses have been processed for uniformity, we restore them based on the rules in EIP-2304 [14]. For example, ENS resolvers store Pay-to-Public-Key-Hash (P2PKH) Bitcoin addresses [6] in the form of scriptPubkey [94]. We restore the BTC addresses by extracting public key hashes and encoding them based on Base58Check [4]. For content hash records, based on EIP-1577 [13], the IPFS hash strings [101] are encoded by Base58 and Swarm hash strings are hex encoded. Thus, we decode them accordingly[6]. For text records, as stated in EIP-634 [15] and ENS documentation [33], the event logs only contain the keys (but not the values). Thus, we use the transaction data related to these event logs and decode them based on ABIs to get the text values.

4.3 Dataset Overview

Using the above approach, we obtain all ledger information until block 13,170,000 (i.e., 2021-09-06 04:14:27 UTC) on Ethereum. The overall statistics are shown in Table 2. In total, we gather over 2.7 million registry logs, 4.4 million registrar logs, and 635 thousand resolver logs. In addition to event logs, we fetch and decode over 13,000 transactions related to text records. We find 617,250 ENS names in the registry records.[7] Further, we restore 514,567 names (including 447,116 `.eth` names, which accounts for 90.1% of all `.eth` names) in total. To the best of our knowledge, this is the largest ENS name dataset to date (even larger than the dataset provided by the official ENS development team). Note that, the data gathered in this paper is public registration data, akin to DNS records. Thus, our analysis is limited to information that registrants have published. Indeed, the purpose of publishing ENS records is such that third parties can view them and this does not raise any ethical concerns.

5 OVERVIEW OF ENS

5.1 Overall Statistics

5.1.1 Overview. Table 3 shows an overview of the 617,250 ENS names. 184,490 addresses have participated in the registration of ENS `.eth` names. There are over 343K active names (include names in the grace period), related to 83.4% of addresses (153,553). Compared with two well-known blockchain-based naming systems published by Patsakis et al. [92] in 2020, where there are over 140K names in both Namecoin and Emercoin, ENS has a relatively high number of registrations. Besides `.eth` names, there are 2,434 DNS names involved in the registration of ENS, showing that some DNS domain name owners are also interested in this naming service on blockchain [18].

5.1.2 The evolution of ENS names. Figure 4 shows the number of monthly registrations. For each name, we use the first block time of the "NewOwner" event to reflect the registration time. The figure shows the trend of ENS names (all ENS names and `.eth` names) registered for the first time each month. Various key occurrences can be observed. The ENS team started the service in March 2017, but encountered two severe bugs and the service went offline [20]. Thus, only basic ENS names like `.eth` or `addr.reverse` were registered in March 2017 and were de-registered soon after. On May 4th 2017, ENS relaunched and the first name registered (after a 5-day auction period) is `rilxxlir.eth`. The first 7 months after the launch witnessed enthusiasm for holding ENS `.eth` names: 192,471 names (51.6% of all `.eth` names) were registered. There is a peak in November 2018, when 43,832 were registered. In this month, 4 addresses registered a large number of Chinese pinyin names (e.g., `tianxian.eth`) and names composed of dates or numbers (e.g., `20140409.eth`), which resulted in them ranking 2-5 in the registration numbers of the auction period. On May 4th 2019, the ENS team launched a new registration registrar (instead of the old auction). The number of registrations increased slightly until the short name auction started from September to November. Note, the short name auction also affected the registration of other names in October and November. In February 2020, Decentraland, a decentralized virtual reality platform, created over 12K subdomain names for their own naming system [30], which led to a rise in the numbers of names. Since June 2021, the number of creations rose sharply partly due to the drop in gas prices [48], and over 30K names were created in June.

[6] Since hashes in the "ContentChanged" events of "OldPublicResolvers" do not have a uniform format and users may upload variable content, their protocol cannot be easily detected and they are treated as Swarm hashes.

[7] We exclude ENS TLDs records and reverse resolution names because our study is focused on second and higher level ENS names.

Table 2: An overview of the ENS event logs we collect.

Contract Type	Etherscan Name Tag	Address	# of event logs
Registry	Eth Name Service	0x314159265dD8dbb310642f98f50C066173C1259b	1,101,810
	Registry with Fallback	0x00000000000C2E074eC69A0dFb2997BA6C7d2e1e	1,612,253
Registrar	Base Registrar Implementation	0x57f1887a8BF19b14fC0dF6Fd9B2acc9Af147eA85	1,757,884
	Old ENS Token	0xFaC7BEA255a6990f749363002136aF6556b31e04	337,073
	Old Registrar	0x6090A6e47849629b7245Dfa1Ca21D94cd15878Ef	1,996,549
	Short Name Claims	0xf7C83Bd0c50e7A72b55a39FE0DABF5e3A330d749	883
Registrar Controller	Old ETH Registrar Controller 1	0xF0AD5cAd05e10572EfcEB849f6Ff0c68f9700455	14,976
	Old ETH Registrar Controller 2	0xB22c1C159d12461EA124b0deb4b5b93020E6Ad16	20,827
	ETHRegistrarController	0x283Af0B28c62C092C9727F1Ee09c02CA627EB7F5	276,954
Resolver	OldPublicResolver1	0x1da022710dF5002339274AaDEe8D58218e9D6AB5	15,283
	OldPublicResolver2	0x226159d592E2b063810a10Ebf6dcbADA94Ed68b8	19,576
	PublicResolver1	0xDaaF96c344f63131acadD0Ea35170E7892d3dfBA	3,494
	PublicResolver2	0x4976fb03C32e5B8cfe2b6cCB31c09Ba78EBaBa41	460,408
	Additional Resolvers	–	136,447

Table 3: The distribution of ENS names. Note that the .eth subdomain owners of expired parent names and integrated name owners of expired DNS names still have control over their names. These are considered active ENS names.

Active ENS Names	343,492	Unexpired .eth Domains	222,456
		Subdomains	118,602
		DNS Integrated Names	2,434
Expired .eth Domains	273,758	-	-
Total	617,250	-	-

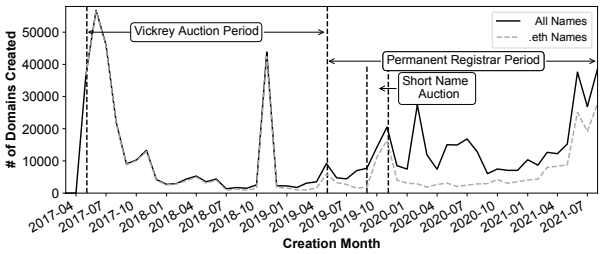

Figure 4: Timeseries of ENS name registrations.

5.1.3 Owners of ENS names. We next track every ownership change of `.eth` names from the ENS registry (i.e., "NewOwner" and "Transfer" events) and analyze the number of `.eth` names held by each address. These 496K `.eth` names have been owned by 107,895 Ethereum addresses. Over 26% of the addresses have more than one name, indicating that, although ENS now has an annual fee mechanism on `.eth` names, there is still a large amount of users holding numerous names. The address[8] that holds most names retains 1,668 names. These hoarded names cover many words in the dictionary (e.g., `pianos.eth` and `judicial.eth`), as well as famous brands (e.g., `ipods.eth`). This address and some other top holders are suspicious for their massive registrations, and we conjecture they are involved in name squatting (see §7).

5.1.4 Name Length. We further analyze the popularity of `.eth` names of different lengths. For this, we use the restored ENS names, i.e., the human-readable names we reconstructed. Figure 5 shows the distribution of `.eth` names whose length is under 20. ENS initially only accepted names longer than 6 characters. Currently, names with a length of under 5 will be charged more than $160 annually. This explains why there are only about 800 `.eth` names with a length of under 5 registered per month after the opening of short name registrations. Instead, `.eth` names that have a length of larger than 6 characters are far more popular. Names with a length ranging from 5 to 8 account for 48.7% of unexpired `.eth` names. There have been 3,531 names whose length is over 20 and the longest name has 10K characters of the "Grinning Cat" emoji [54].

5.2 The Initial Auction (Vickrey Auction)

5.2.1 Overview. Between 2017–2019, name registrations were performed using Vickrey auctions. In this period, there were 361,751 names that have been bid for. Among them, 274,052 names were registered with 338,252 valid bids by 17,625 addresses. Note that over 80K names did not finish the auction process. Figure 6 shows the distribution of the bids and the final auction price. 45.7% of the bids were 0.01 ETH while 92.8% of the names' final auction price were 0.01 ETH, suggesting that most people want to get their desired names with the minimum price, and most names got few bids due to the protection of namehash. The highest bid was 201,709 ETH for `ethfinex.eth`, while its auction price is 0.01 ETH.

5.2.2 The most valuable names. The most valuable name is `darkmarket.eth`, which cost over 20K ETH (about $5 million). The owner of it[9] also registered another three valuable names including `openmarket.eth`, `tickets.eth`, and `payment.eth`. The owner address belongs to a well-known exchange, Bitfinex [7]. Note that 7 of the top-10 valuable names had not set any records by the time of this study, indicating that they may be used for squatting purposes.

[8] 0xbcbd4885ee8b2b74249c5ad9b8b668fb256a51b1

[9] 0x8759b0b1d9cba80e3836228dfb982abaa2c48b97

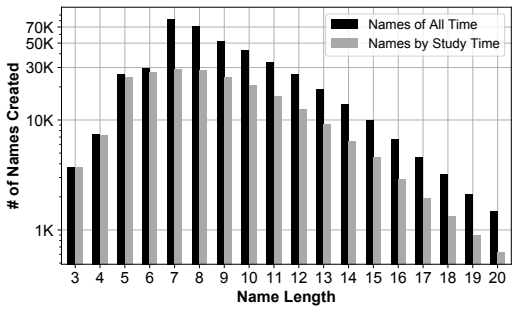

Figure 5: The distribution of `.eth` names' length.

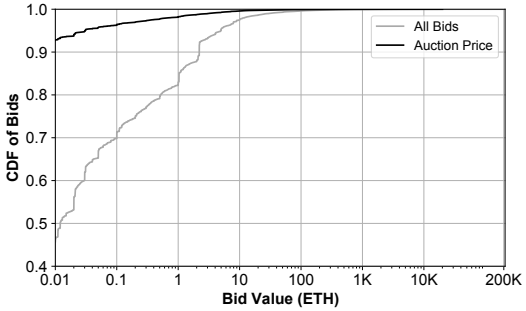

Figure 6: The distribution of bids and auction prices.

5.2.3 Holders involved in ENS Vickrey auction. We find that each of the top-10 name holders registered over 2,000 names. The last 9 in the top-10 each spent under 150 ETH in total. The top name holder[10] ranked 15th among the top spent addresses. For the top-10 spent addresses, each spent over 2,500 ETH on relatively few names. The behaviors of these top bidders suggests that there are two straightforward strategies when people bid for desired ENS names. Some people tend to register as many names as they can with low prices, while others prefer to bid for a few high-value names that are worth a lot of money.

5.3 Short Names (English Auction)

5.3.1 The Short Name Claim. During this period, 344 requests were submitted and 193 were approved. Among the applications, some traditional websites like `nba.com`, `paypal.cn`, `ebay.net` and `opera.com` applied and got corresponding `.eth` names, indicating that ENS has received attention from the wider business community.

5.3.2 The Short Name Auction. Since this auction took place in OpenSea and the details of this auction are not shown in the ENS contracts' event logs, we take advantage of the data shared by OpenSea in the ENS blog [37] to analyze the trends of the auction. In total, there are over 50K bids and 7,670 names sold, amounting to 5,697 ETH during this auction. The price distribution is shown in Figure 7. Roughly 10% of the names have a price of over 1.5 ETH (about $300 at that time) and over 22% of the names were bid for over 10 times. The top-10 popular names and expensive names are shown in Table 4. It is not surprising to see that famous companies like "apple", "google", "amazon" and terms like "sex" and "porn" are present. We also find some blockchain-related ENS names like "assets" and "dapp". Compared with the name price in the Vickrey auction period, the name price in the short name auction tends to be relatively low since users need to pay the bids (instead of depositing the payments in the deeds and having the ability to retrieve after one year). Considering that there were few brands claiming their corresponding `.eth` names in the short name claim period, it is possible that bad actors bid for famous brand names and use them for malicious purposes. We further investigate whether there are name squatters in §7.1.

[10] 0xa7f3659c53820346176f7e0e350780df304db179

Table 4: The top-10 popular names and expensive names during the short name auction.

Name	# of Bids	Price in ETH	Name	# of Bids	Price in ETH
amazon	36	100	asset	83	30
wallet	51	75	banker	78	10.5
google	47	52.9	durex	70	1.4
apple	67	51	apple	67	51
sex	44	41	lawyer	66	7.1
porn	44	40	hotel	60	20
com	16	39.8	pussy	58	8
dapp	34	38.7	kering	58	1.4
loan	30	38	foster	58	1.1
jobs	22	35.4	poker	57	33.5

5.4 Name Renewal

Since its introduction, the Permanent Registrar allows the expiration and renewal of names. The first batch of name expiration began in May 2020. Figure 8 shows the distribution of expired names and renewed names (status by the time of the study). Note that we take the 90-day grace period into consideration. We see that most of the names expired in August 2020 due to the expiration of Vickrey auction names. The renewals occurred mainly around August 2020. Another peak occurred around May 2021 possibly due to the second year's renewal of first renewed names.

Figure 9 shows the distribution of 1,859 premium name registrations. 44 premium names were registered on the first day (August 2nd), suggesting that they were registered with almost the full premium. For example, decentralised finance (DeFi) related ENS names (e.g., `opensea.eth` and `balancer.eth`) were registered almost immediately once they were released. A further examination of these premium names reveals that they were registered by many blockchain related exchanges or companies such as Opensea, Crypto Valley [49], MyCrypto [53], Synthetix [58], etc. This shows that ENS is seen as having value to these blockchain developers. By design, the first batch of names was available for registration without premium on August 30th. This was clearly noted by users and explains the spike around the end of August. Over 1300 names (72% of all premium names) were registered on August 29th. To some extent, this premium method gives people a greater chance of obtaining the names they strive and shows that financial incentives have played a major role in the early stages of ENS.

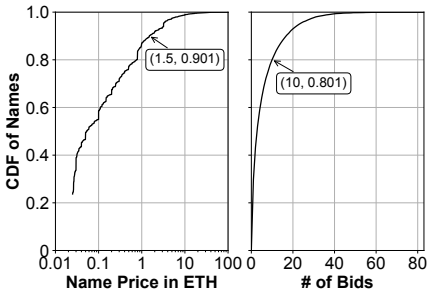

Figure 7: The distribution of short names' price and bids.

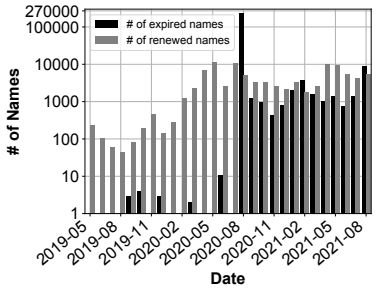

Figure 8: The distribution of expired names and renewed names.

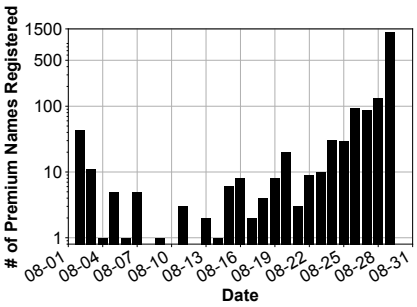

Figure 9: The distribution of premium name registration.

Table 5: The distribution of names that have records and records per name.

Names	# of Names	# of Record Types Per Name	# of Names
Name that has records	278,117	1	255,900
`.eth` name that has records	204,524	2	15,372
Unexpired `.eth` name that has records	181,468	3 - 58	6,845

Answer to RQ1: *ENS has experienced growing popularity. Over 617K names were registered and 341K of them were active by the time of this study. A number of users pay high prices for rare ENS names or try to get as many names as they can.*

6 THE RECORDS OF ENS NAMES

Next, we analyze the use of ENS names, particularly in terms of the information stored in records.

6.1 Overview of ENS records

Over 278K names have their records set over 370K times. Figure. 10(a) presents the distribution of official-defined record settings. The most widespread use of ENS is to index blockchain addresses, accounting for 85.8% of total record changes. Other records include content hash records, public key records, text records, etc. The distribution of names that have each type of record is shown in Table 5. Interestingly, only 45% of the names have ever had records. Users could set resolvers and records when they register their names in a single transaction through registrar controller contracts. However, before the existence of registrar controllers, a user needs to make additional transactions for setting records, which adds the cost of the user and may lead to a low rate of record settings. This suggests users may be purchasing names for later use.

Note that, in ENS, a name can be associated with multiple types of records. Specifically, the records could contain different blockchain addresses, text records with different keywords and other single records. Table 5 shows the distribution of the record counts per name. Most names have one record and 98.8% of records are Ethereum addresses. The name that has the most record types is `qjawe.eth`: it set 58 kinds of records including 51 blockchain addresses and 7 text records such as Twitter, Github, emails, etc.

6.2 Use of Blockchain address records

Most of the addresses contained within records pertain to Ethereum (307, 964 records) or BTC addresses (3, 980 records). The distribution of the top-5 non-ETH addresses is shown in Figure 10(b). In total, there are 82 kinds of non-ETH blockchain addresses are set over 9000 times, suggesting ENS is being used by more and more investors in blockchains besides Ethereum.

6.3 Use of content hash records

Another major use of ENS is to store content hashes. Over 9, 200 names have been set to content hash records. There are roughly 6, 000 names with non-empty values and the distribution is shown in Figure 10(c). Most of the content hashes (99.6%) are set for the InterPlanetary File System [35] (IPFS) and Swarm [41],[11] two prominent solutions for decentralized storage. Specifically, the "ipns-ns" hashes are used for the InterPlanetary Name System (IPNS), which is designed for mutable content [31]. A few names are also set for Tor `.onion` addresses, while 10 of them (e.g. `facebooktor.eth`, `protonmailtor.eth`) have been set by the ENS team [36] for guidance on Tor site resolutions. The remaining records consist of malformed IPFS hashes. For example, the nine "multicodec" hashes are generated by one user by encoding IPFS hashes twice. The results show that ENS still has some scope for dWeb resolution. We further investigate malicious dWeb uses in §7.2.

6.4 Use of text records

Finally, ENS allows users to set arbitrary text records (in the form of key-value pairs). The top-10 types of text records are shown in Figure 10(d). Most settings are for URLs, and we find that over 10% of the records are set to subdomains of OpenSea (a P2P marketplace used in short name auctions), suggesting that these names are for sale. Other URLs are used for official sites (e.g., `tokenfactory.global`), personal blog sites (e.g., `marvin-elsen.com`), etc. Besides predefined keywords like "com.twitter" (Twitter accounts, legacy "vnd.twitter") and "description" (description for names), there are also customized keywords in the text records. For example, there are 1846 "snapshot" records, which are used for Snapshot [55], a decentralised voting system. There are also records like "dnslink", which are mainly used by DAppNode [9] for their dWebs and "gundb"

[11] IPFS and Swarm are distributed data sharing networks, which are commonly used for decentralized webs (dWebs)

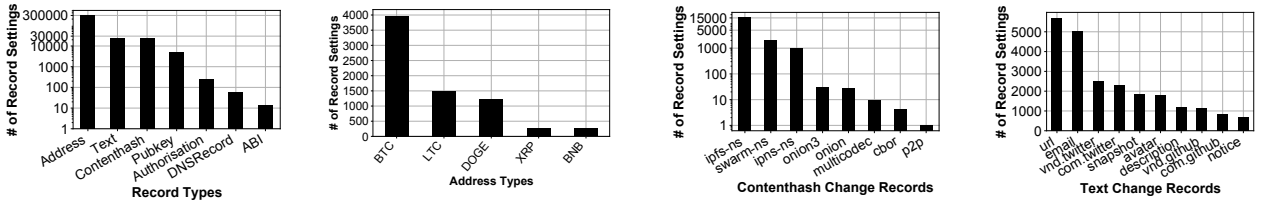

(a) The distribution of ENS names' records (b) The distribution of top-5 non-ETH addresses. (c) The distribution of content hash records. (d) The distribution of top-9 text records.

Figure 10: The distribution of all record types (a), and three major record types (b, c, d).

records used for GunDB [25], a P2P database. We identify 150 customized keywords in 2, 938 record settings in total, indicating that people are exploring new ways of using ENS. We further analyze security issues related to URL records in §7.2.

Answer to RQ2: *ENS is an open system where names can index a wide range of records. The most common is linking to blockchain addresses (85.8% of record changes). It is also common to use names for resolving other dWeb resources, websites and text. Further, developers are also exploring new ways of using ENS. This suggests that ENS has a potential to be the part of infrastructure in the blockchain world.*

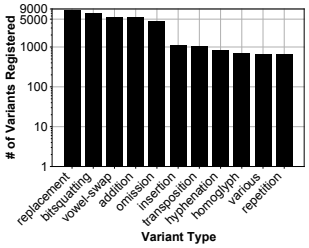

Figure 11: The distribution of squatting types.

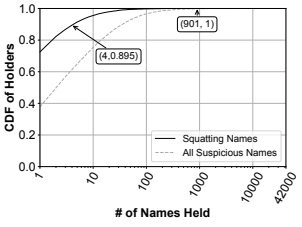

Figure 12: The distribution of ENS squatting/suspicious name holders.

7 SECURITY & MISBEHAVIOR IN ENS

We next investigate security issues on ENS, including the applicability of existing (DNS) attacks as well as new issues.

7.1 Domain Squatting

As observed in previous DNS research [100], some attackers perform typo-squatting to register domain names that are similar to well-known ones in order to hijack a website's users, in cases where typos are made. We now explore if ENS also suffers from this attack.

7.1.1 Explicit Squatting of Known Brands. As stated in §4.2, we use the Alexa top-100K domains to match hashed versions of the ENS names. Among the top-100K names in Alexa, there are 18, 984 names that could be found in ENS native 2LDs (only 32 of which were claimed in the short name claim period).

One challenge is that the above does not mean the matching ENS names are squatting, since they could be registered by the actual owner. As a heuristic, if one Ethereum address owns more than one known ENS name (e.g., both `google.eth` and `facebook.eth`) and if these domains belong to different owners (shown via Whois) in DNS, we assume this address is performing a squatting attack since it is likely to seize these names for future sale or conducting other misbehaviors based on these squatting names.

Through this, 15, 117 ENS `.eth` squatting names controlled by 2, 005 Ethereum addresses are identified. Many famous brands are registered for squatting. For example, one address[12] has registered `google.eth`, `mcdonalds.eth`, and `redbull.eth`. As these brands do not belong to the same owner, this is likely name squatting. Among these explicit squatting `.eth` names, over 64.5% are still active, suggesting that attackers hold these names for a long time. We conjecture this is a form of speculation, where owners perceive future value in name ownership.

7.1.2 Typo-Squatting ENS Names. Typo squatting occurs when a domain is registered that is similar to a real domain such as `facebok.com`. To detect typo-squatting ENS names, we use dnstwist [11], a widely used tool [70, 73, 96] to generate typo-squatting variants of domain names and it can generate 12 kinds of squatting variants. We feed all Alexa top-100K domains to dnstwist and get 764, 235, 725 variants. We then calculate the labelhash of their 2LDs to check whether these squatting names have been registered in ENS. To reduce false positives, we only keep names (and their raw names) with a length of more than 3. Note that, as owners that claim their `.eth` names in ENS (see Section 3.2.2) may continue to register variants of their names, we first check if these squatting variants are ever owned by them. We find no cases in which they own these squatting variants. The final result shows that squatting is surprisingly common: we identify 28, 189 ENS typo-squatting `.eth` names, targeting 16, 097 Alexa domain names. Figure 11 shows the distribution of variant types. There are over 6K bitsquatting variants and 683 homoglyph domains. Over 72% of the typo-squatting ENS `.eth` names are still active by the time of our study. This suggests that ENS has experienced extensive typo squatting and urgently requires mitigation strategies to protect legitimate name owners.

7.1.3 Squatting Names Analysis. In total, we collect 43, 306 unique squatting ENS `.eth` names based on the previous heuristics. We next investigate the records and holders of these squatting domains.

Records of squatting names. 23, 166 squatting ENS `.eth` names (21, 941 active ones) have records set, and most of these (86%) set only blockchain address records. Besides Ethereum addresses, some of other records are related to sales like Opensea links and IPFS

[12] 0x782cf6b6e735496f7e608489b0c57ee27f407e7d

websites posting sale information. This suggests that monetization is a key driving force behind squatting.

Relations Between Addresses. For the identified squatting ENS `.eth` names, the *name-to-holder* and *holder-to-holder* relations are further investigated. Overall, these 43,306 names have been owned by 15,726 addresses. Interestingly, some names have actually been owned by multiple addresses. Further investigation reveals that some addresses also transferred their names. We see that several addresses hold a large number of ENS `.eth` names. Figure 12 plots the number of squatting names held by each address. The top 10% of addresses hold more than 4 squatting ENS `.eth` names each. These names account for a remarkable 64% of all squatting ENS names, suggesting that a small number of attackers are highly active.

Guilt-by-Association Expansion. As our prior analysis relies on variants of the top Alexa domains, it is possible to miss other squatting names. We next inspect the behavior of users who participate in squatting at least once. The heuristic is that, if a squatter has seized a popular name or its variant, they tend to squat on other names too. We call this "guilt-by-association", as used in previous work to identify malicious domains and malware [85, 95]. We thus analyze all ENS names held by the identified squatters. Through this, we find 321,459 suspicious squatting `.eth` names. Figure 12 plots the distribution of these addresses, for all suspicious names. Over 33% of the squatters have held more than 10 ENS `.eth` names, accounting for 92% of all suspicious names in total. The top-10 holders are shown in Table 7 of the Appendix. The top address, `0xbd21109e2bdcb24c4fbcdc16a4c90f34e81228e2`, has acquired 901 ENS squatting names (confirmed using heuristics in §7.1.1 and §7.1.2), and has held more than 40K names in total (suspected to be other squatting names). Again, this confirms that just a few key players are driving the majority of squatting within ENS. Remarkably, these top-10 addresses have held around 18% of *all* `.eth` names. This indicates that squatters are likely to register more ENS names even when these names are not quite popular. Thus, blocking just a small number of addresses will allow a significant fraction of squatting attacks to be mitigated. Although such attackers could then change addresses, blocking could still complicate their activities.

The evolution of squatting names. Figure 13 shows the evolution of suspicious ENS squatting names. Multiple spikes are present. The first spike of squatting name registrations (e.g., `zhifubao.eth`, alias of Alipay) occurs around 2017 May 9th, almost at the same time as the initial auction. Beyond this, the overall squatting trends follow a similar pattern to the general names (see §5). This suggests that there has been squatting behavior across all periods of ENS. Interestingly, we notice that when the Permanent Registrar increased the cost of holding a large amount of names, most expired names were given up by the squatters. In particular, during the 2018 Nov. spike, one address[13] registered more than 40K names. Yet by the time of this study, the account does not own any names. The number of active explicit squatting names also decreased to 5,230 (2.3% of all active ENS `.eth` names). As a comparison, Patsakis et al. [92] found over 30% of active Namecoin names and 58% of Emercoin names are explicit squatting names. This suggests the mechanisms of ENS registrations mitigate the impact of explicit squatting behaviors. However, there are still some squatters holding names for potential profit: 124,253 suspicious ENS squatting names (56% of all active ENS `.eth` names) are still actively held.

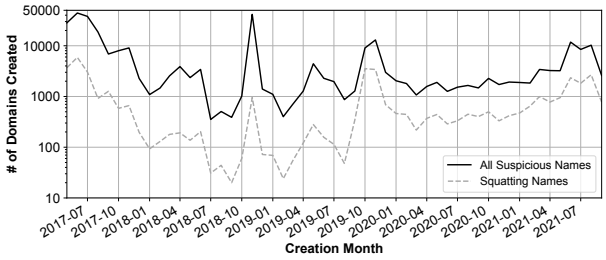

Figure 13: The evolution of squatting names.

7.2 Websites with Misbehaviors

Since ENS records can store URLs, it is a possible that bad actors index content or websites with misbehaviors. Thus, we inspect the web content that has been stored in ENS. As mentioned in §6, we get 15,320 unique dWeb hashes, 50 onion hashes and 4,644 URLs.

7.2.1 Methodology. We upload all URLs to VirusTotal [45]. Following previous studies [79, 93], if a URL is reported by 2 or more anti-virus engines, it is marked as suspicious. We then use Eyewitness [23] to get screenshots and source codes for each website, which are uploaded to Google Cloud Natural Language API and Vision API [46] to check whether they contain sensitive (e.g., adult and gambling) content. We also tag URLs that have keywords like "casino" or "generator" in their names or web contents. All suspicious URLs are manually inspected to reduce false positives.

7.2.2 Results. In total, we identify 29 dWeb URLs with misbehaviors and one phishing domain in the 28 2LD ENS names. Examples are shown in Figure 16 of the Appendix. These websites are involved in gambling (11), adult content (6) and scams (13). Note, since dWeb URLs may not store content online persistently, some content cannot be reached and the actual number of dWeb sites with misbehaviors may be higher than identified. This observation is inconsistent with prior findings identifying malware on IPFS [84, 91]. A possible reason is that malware developers choose not to bind IPFS websites to ENS names, since they can distribute malware in other ways or use IPFS as C&C servers directly.

7.3 Scam Addresses

As mentioned in §6, ENS is used mainly for storing Blockchain address records. Thus, we further seek to analyse whether any addresses are used for malicious purposes.

7.3.1 Methodology. There is no available comprehensive dataset of scam blockchain addresses. Hence, we first compile a scam address list from various sources. Etherscan and Bloxy [8] have labelled a list of "phishing" or "hacked" Ethereum addresses. BitcoinAbuse [5] and CryptoScamDB [27] host websites for tracking malicious addresses. We also fetch a scam token list from previous literature [74]. We crawl all the addresses above and obtain 90K in total. We then match the addresses stored in ENS with the scam address list.

[13] 0xbd21109e2bdcb24c4fbcdc16a4c90f34e81228e2

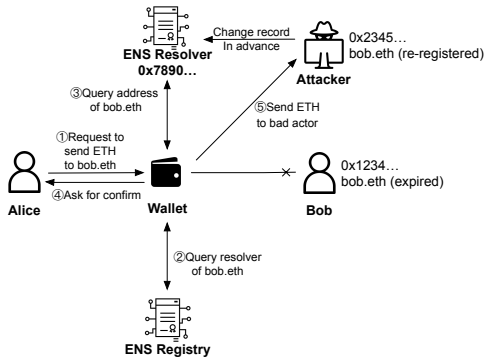

Figure 14: An Example of the attack scenario that exploits the record persistence issue.

7.3.2 Results. We find 13 scam addresses registered in ENS, as shown in Table 9 of the Appendix. They are used in airdrop scams, Ponzi schemes, ransomware and Rug Pull scams. In particular, we find 3 homoglyph ENS names that impersonate Vitalik Buterin (his own ENS name is `vitalik.eth`), one of the Ethereum co-founders, to perform giveaway scams. The few occurrences might be due to the cost of proper scam name registrations. Despite few occurrences, this confirms that ENS *can* be used as part of scams.

7.4 Record Persistence Attack

We find a novel attack that relies on the unique attributes of ENS.

7.4.1 Attack Scenario. Similar to the orphan records in DNS [81, 98], when an ENS name expires, the name and its subdomain names' records are kept. Hence, some ENS-supported wallets can still resolve them. Resolver smart contracts of ENS do not erase the old records until the new ones replace them. A standard resolution process will not check the expiration status of one name alongside its 2LD name. Further, the erase operation and status check operation could be resource-consumed when resolving high-level domains (3LDs, 4LDs, etc). In contrast to orphan records in DNS, ENS names are usually linked with blockchain addresses and we conjecture this may lead to scams or financial losses. An example of this is shown in Figure 14: When an attacker registers the expired name and changes the name's record, people who are unaware of the change and do not check the recipient addresses (this is originally one of the purposes of using ENS) would send money to the attacker.

7.4.2 Vulnerable ENS Names. We attempt to find how many ENS names are vulnerable to this attack. This involves checking how many expired ENS names still have records. In total, we find that 22,716 expired `.eth` names have records within them or their subdomains (2,318 subdomains). Examples are shown in Table 8 of the Appendix. These names contain not only blockchain addresses but also dWeb hashes, which may direct users to malicious websites if abused. In particular, the ENS name `thisisme.eth` has 706 subdomain names and all of them have Ethereum address records. It was listed in `enslisting.com` for free registration and was transferred to a smart contract to ensure that subdomain name records of this name could not be modified easily. Also, the ENSListing team claimed that they would cover the cost of annual renewals [1]. However, the name expired on May 4th 2020, thus *we registered it for protection.* This expiration meant that it could not assign any more subdomains (although the remaining records could still be resolved). Thus, our registration has no negative impact on the names nor their users. We will not change the record on these subdomain names and we will further devise ways to handle the name, like transferring the name to the ENS team or interacting with the subdomain registrar contract. Interestingly, there are still some people using subdomain names of this name. One example is an address taking part in an airdrop activity with their `thisisme.eth` names (see Figure 15 of the Appendix), which was widely publicized on Twitter. As we mentioned, attackers can identify vulnerable ENS names and register them for making a profit. Although in June 2020 ENS team has proposed email notifications to remind people to renew their names [40], this is still a severe security concern.

> **Answer to RQ3:** *ENS is vulnerable to traditional security issues, as well as the new record persistence attack. A small set of squatters hold many famous brand names and their variants, which could be used for malicious purposes. Some malicious decentralized websites and scam addresses are also found in records. Expired ENS names are vulnerable to record persistence attack, and there is currently no deployed mitigation.*

8 DISCUSSION

8.1 The status quo of ENS

It has been one year since we collected the data in this paper. To check the status quo, we therefore collect the ledger information between block 13,170,000 (2021-09-06 04:14:27 UTC) to block 15,420,000 (2022-08-27 06:23:05 UTC) and briefly summarize them. By August 27th, 2022, there are over 16 million additional event logs found in the smart contracts, including 5 million registry logs, 5 million registrar logs, 1 million registrar controller logs and 3 million resolver logs. Among all 1,678,502 newly registered names, 97% of them are `.eth` names. The majority (73%) of `.eth` names are registered after April 2022, possibly due to attention to ENS raised by the interest in short digit domain names in secondary markets like OpenSea [60]. Besides blockchain address records, we also find that over 40K names have a avatar record, which is used to link the image links of NFT tokens. This indicates that the popularity of ENS among investors and users continues to expand.

8.2 Implications

Our findings are important for stakeholders in the ecosystem. As all BNS share some common properties, the methods used in this work could be extended to study other BNS systems.

(Blockchain-based) naming system developers. Considering the number of active addresses and users on ENS, along with the integrated dApps and DNS TLDs, we find that ENS still has a relatively healthy ecosystem. According to the official team [47], they have planned to make many improvements including scaling on Layer 2 [52]. This will facilitate daily use and reduce the cost for ENS users. However, since we have found several security issues, this breaks the security requirement of Zooko's triangle. We argue the ENS team must therefore consider solutions to address these, much akin to activities within the DNS community. For other naming system developers, there are some experience that could be

learnt from ENS. First, the ENS team uses a multisig wallet contract and it can make changes on ENS core contracts like registry contracts and registrar contracts. This may diminish the decentralization claim of ENS. However, the evolution of ENS shows that this setup gives them more chance to avoid severe vulnerabilities and keep the contracts functioning properly. Thus, we argue that naming system developers need to balance these trade-offs of using fully decentralised structures. Second, during different periods of ENS registrations, the ENS team designed many mechanisms to reduce the effects of squattings such as the namehash design and the short name claim. Although we still find many suspicious squatting ENS names, compared with other well-known blockchain-based naming systems, the design has prevented many explicit squatting behaviors. It is necessary for other naming systems to design a fair registration mechanisms. Third, ENS provides a wide range of records for users to set, including wallet addresses from other blockchains and highly customizable text records. ENS also allows DNS names to integrate in ENS for resolving on-chain resources. The dApps from other blockchains can easily use ENS for resource resolution. Other naming systems can also consider these designs to build a more scalable system.

Cryptocurrency wallets, dApps, exchanges and blockchain browsers. We further argue that, considering the difficulty of making changes to deployed smart contracts, developers of blockchain wallets dApps, exchanges and blockchain browsers should take measures to detect squatting names or malicious records. This can be used to give reminders to users who are trying to interact with suspicious names. In particular, blockchain wallets should warn subdomain users of expired ENS names. They should also know the risk of the persistence record attack and take active measures.

Investors Awareness should also be raised among investors. They need to learn more about this blockchain-based naming system and the possible risks when interacting with it. Specifically, it is wise to validate the real addresses under the ENS names they resolve.

8.3 Limitations

We have only restored 90.1% of all `.eth` names to their readable names (see §4.2). This means we may have missed certain attacks, e.g., combo-squatting [86] ENS names. Nevertheless, our dataset of recovered names is the largest of its kind, and it has allowed us to study explicit name squatting and typo-squatting (since we could calculate their hash values). When identifying typo-squatting ENS names, we use dnstwist to generate variants of target Alexa top domains. Although we use name length and the legitimate owner to reduce false positives, dnstwist may still generate false positive typo-squatting names. As it is hard to distinguish whether a suspicious typo-squatting name is a squatting one, we intend to explore more accurate approaches in the future. For a naming system, it is more effective to track its real-world usage like daily resolution counts or daily users. However, since ENS uses "external view" functions for record resolution (and this does not involve blockchain transactions), it is hard for us to track the actual usage of ENS. Nevertheless, the status quo of ENS still suggests its popularity and activities to some extent. Beyond the attacks inspected, we note there are many other categories of attacks not covered (e.g., DDoS). Investigating these constitute our future work.

9 RELATED WORK

Design of BNS Systems. Many works have been studying the design of BNS systems. Hari et al. [77] is one of the first to propose the design of blockchain-based name system. They analyzed the limitations of traditional DNS and their dependencies on Public Key Infrastructures (PKIs). They also proposed a distributed, tamper-resistant DNS infrastructure. Similarly, Guan et al. [76] presented a domain authentication scheme named AuthLedger to reduce the level of trust in certificate authorities (CAs). Besides, some studies are focused on improving the security of the DNS nodes [78, 88, 103, 105]. For example, He et al. [78] proposed a trustworthy decentralised DNS root management architecture based on permissioned blockchain. Gourley et al. [75] proposed an improved DNSSEC based on blockchain, which provides the same security benefits as DNSSEC whilst addressing its drawbacks.

Analysis of BNS Systems. A few studies have analyzed different kinds of blockchain-based DNS systems [61, 62, 65, 82, 83, 89, 92, 104]. Kalonder et al. [82] performed a empirical analysis of Namecoin. They found only 278 of 120K registered domains in Namecoin are resolved to IP addresses with "regular" content (i.e., web pages without error responses or default pages of web servers). They also found that its domain market is thin-to-nonexsitent. After that, other studies have characterized the properties of different kinds of blockchain-based naming systems. Patsakis et al. [92] surveyed the threat of blockchain-based naming systems, including malware, underlying registrar mechanisms, domain markets, phishing, and immutability. Their findings revealed several potential domain extortion attempts and possible phishing schemes. Nevertheless, very few studies have inspected ENS, the leading smart contract based BNS system. Liu et al. [89] and Karaarslan et al. [83] compared designs of several blockchain-based naming systems including ENS. However, they only introduced the architecture of ENS and lack a systematic study of it, particularly related to security issues.

10 CONCLUSION

We have presented the first systematic study of the Ethereum Name Service (ENS). By collecting and analyzing millions of ENS event logs, we find a number of observations. We believe ENS is a promising system, on its way to being part of the naming service infrastructure in the blockchain world. Yet, we argue that spam and security issues may impede its progress. Our efforts in this paper can positively contribute to the BNS ecosystem and offer practical insights to support users in identifying misbehavior.

ACKNOWLEDGEMENTS

We sincerely thank our shepherd Prof. Paul Barford (University of Wisconsin–Madison) and all the anonymous reviewers for their valuable suggestions and comments to improve this paper. This work was supported in part by National Key R&D Program of China (2021YFB2701000), the National Natural Science Foundation of China (grants No.62072046, No.61873069), the Fundamental Research Funds for the Central Universities (HUST 3004129109), and Hong Kong RGC Projects (No. PolyU15219319, PolyU15222320, PolyU15224121).

REFERENCES

[1] Ensnow soft launch! get an instant ens name for your wallet for free! https://medium.com/@enslisting.com/ensnow-soft-launch-get-an-instant-ens-name-for-your-wallet-for-free-3b56ace6b60a, 2017.
[2] Alexa top 1 million sites. http://s3.amazonaws.com/alexa-static/top-1m.csv.zip, 2020.
[3] As feds fumble with bitcoin, the internet trolls the fbi's "private" wallet. https://techcrunch.com/2013/10/07/as-feds-fumble-with-bitcoin-the-internet-trolls-the-fbis-private-wallet/, 2020.
[4] Base58check encoding. https://en.bitcoin.it/wiki/Base58Check_encoding, 2020.
[5] Bitcoin abuse database. https://www.bitcoinabuse.com/, 2020.
[6] Bitcoin address formats and performance comparison. https://medium.com/fixedfloat/bitcoin-address-formats-3522cf47bdf4, 2020.
[7] Bitfinex | cryptocurrency exchange | bitcoin trading | futures ... https://www.bitfinex.com/, 2020.
[8] Bloxy. https://bloxy.info, 2020.
[9] dappnode.io. https://dappnode.io/, 2020.
[10] A decaying price premium for newly released .eth names. https://medium.com/the-ethereum-name-service/new-decaying-price-premium-for-newly-released-names-72080a650c15, 2020.
[11] Domain name permutation engine for detecting typo squatting, phishing and corporate espionage. https://github.com/elceef/dnstwist, 2020.
[12] Dune analytics. https://duneanalytics.com, 2020.
[13] Eip-1577: contenthash field for ens. https://eips.ethereum.org/EIPS/eip-1577, 2020.
[14] Eip-2304: Multichain address resolution for ens. https://geth.ethereum.org/, 2020.
[15] Eip-634: Storage of text records in ens. https://eips.ethereum.org/EIPS/eip-634, 2020.
[16] Emerdns. https://emercoin.com/en/emerdns, 2020.
[17] Ens. https://ens.domains/, 2020.
[18] Ens + .kred: Major integration of dns and ens launches. https://medium.com/the-ethereum-name-service/ens-kred-major-integration-of-dns-and-ens-launches-e7efb4dd872a, 2020.
[19] The ethereum name service. https://medium.com/the-ethereum-name-service, 2020.
[20] Ethereum name service launch postmortem. https://medium.com/the-ethereum-name-service/ethereum-name-service-launch-postmortem-a941864f4b5, 2020.
[21] Etherscan. https://etherscan.io/, 2020.
[22] Exclusive: Hackers acting in turkey's interests believed to be behind recent cyberattacks - sources. https://www.reuters.com/article/us-cyber-attack-hijack-exclusive/exclusive-hackers-acting-in-turkeys-interests-believed-to-be-behind-recent-cyberattacks-sources-idUSKBN1ZQ10X, 2020.
[23] Fortynorthsecurity/eyewitness. https://github.com/FortyNorthSecurity/EyeWitness, 2020.
[24] Go ethereum. https://geth.ethereum.org/, 2020.
[25] Gun — the database for freedom fighters. https://gun.eco/, 2020.
[26] Handshake. https://handshake.org/, 2020.
[27] Home | cryptoscamdb. https://cryptoscamdb.org/, 2020.
[28] Home | ethereum.org. https://ethereum.org/en/, 2020.
[29] https://www.innominds.com/blog/securing-ethereum-wallet-with-multisig, 2020.
[30] Id goes di. https://decentraland.org/blog/announcements/id-goes-di/, 2020.
[31] Interplanetary name system (ipns). https://docs.ipfs.io/concepts/ipns/#interplanetary-name-system-ipns, 2020.
[32] Introducing .luxe on ens. https://medium.com/@weka/introducing-luxe-on-ens-35a9ee2383ce, 2020.
[33] Introduction - ethereum name service. https://docs.ens.domains/, 2020.
[34] The invisible internet project. https://geti2p.net/en/, 2020.
[35] Ipfs powers the distributed web. https://ipfs.io/, 2020.
[36] List of ens names that resolve to tor .onion websites. https://medium.com/the-ethereum-name-service/list-of-ens-names-that-resolve-to-tor-onion-websites-99140a4c674f, 2020.
[37] The most popular .eth names in the ens short name auction (final). https://medium.com/the-ethereum-name-service/the-most-popular-eth-names-in-the-ens-short-name-auction-final-5d3466dd8837, 2020.
[38] Namecoin. https://www.namecoin.org/, 2020.
[39] Opensea: Buy crypto collectibles, cryptokitties ... https://opensea.io/, 2020.
[40] Receive email notifications to renew your .eth names with new tool from buidlhub. https://medium.com/the-ethereum-name-service/receive-email-notifications-to-renew-your-eth-names-with-new-tool-from-buidlhub-72aaba226194, 2020.
[41] Swarm. https://swarm.ethereum.org/, 2020.
[42] Tokenview. https://btc.tokenview.com/en/address/385cR5DM96n1HvBDMzLHPYcw89fZAXULJP, 2020.
[43] Tor project | anonymity online. https://www.torproject.org/, 2020.
[44] Unstoppable domains. https://unstoppabledomains.com/, 2020.
[45] Virustotal. https://www.virustotal.com/gui/, 2020.
[46] Vision ai | derive image insights via ml | cloud vision api. https://cloud.google.com/vision, 2020.
[47] 2021 ens roadmap - feedback welcome! https://discuss.ens.domains/t/2021-ens-roadmap-feedback-welcome/328, 2021.
[48] 4 metrics showcasing ens adoption. https://info.etherscan.com/4-metrics-showcasing-ens-adoption/, 2021.
[49] Crypto valley association: Home. https://cryptovalley.swiss, 2021.
[50] Ethereum name service grows by 10,700 addresses in june, now attached to $277 million. https://www.theblockcrypto.com/post/111541/ethereum-name-service-grows-by-10700-addresses-in-june-now-attached-to-277-million, 2021.
[51] Introduction to smart contracts. https://ethereum.org/en/developers/docs/smart-contracts/, 2021.
[52] Layer 2 scaling. https://ethereum.org/en/developers/docs/layer-2-scaling/, 2021.
[53] Mycrypto - ethereum wallet manager. https://mycrypto.com, 2021.
[54] Smiling cat face with open mouth emoji - emojipedia. https://emojipedia.org/grinning-cat/, 2021.
[55] Snapshot: Home. https://docs.snapshot.org/, 2021.
[56] Squaring the triangle: Secure, decentralized, human-readable names. http://www.aaronsw.com/weblog/squarezooko, 2021.
[57] Step-by-step guide to importing a dns domain name to ens. https://medium.com/the-ethereum-name-service/step-by-step-guide-to-importing-a-dns-domain-name-to-ens-d2d15feb03e8, 2021.
[58] Synthetix. https://synthetix.io, 2021.
[59] Zcash: Privacy-protecting digital currency. https://z.cash/, 2021.
[60] Record-high surge in ethereum name service domains triggers 90% rally in ens. https://cointelegraph.com/news/record-high-surge-in-ethereum-name-service-domains-triggers-90-rally-in-ens, 2022.
[61] Saif Al-Mashhadi and Selvakumar Manickam. A brief review of blockchain-based dns systems. *International Journal of Internet Technology and Secured Transactions*, 10(4):420–432, 2020.
[62] Faizan Safdar Ali and Alptekin Kupcu. Improving pki, bgp, and dns using blockchain: A systematic review. *arXiv preprint arXiv:2001.00747*, 2020.
[63] Suranjith Ariyapperuma and Chris J Mitchell. Security vulnerabilities in dns and dnssec. In *The Second International Conference on Availability, Reliability and Security (ARES'07)*, pages 335–342. IEEE, 2007.
[64] Lawrence M Ausubel et al. A generalized vickrey auction. *Econ0 metrica*, 1999.
[65] Dmitry Bagay. Blockchain-based dns building. *Procedia Computer Science*, 169:187–191, 2020.
[66] 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. *Performance, and Policy in the Internet Ecosystem (July 27, 2019)*, 2019.
[67] Timm Böttger, Felix Cuadrado, Gianni Antichi, Eder Leão Fernandes, Gareth Tyson, Ignacio Castro, and Steve Uhlig. An empirical study of the cost of dns-over-https. In *Proceedings of the Internet Measurement Conference*, 2019.
[68] Taejoong Chung, Roland van Rijswijk-Deij, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M Maggs, Alan Mislove, and Christo Wilson. A longitudinal, end-to-end view of the {DNSSEC} ecosystem. In *26th {USENIX} Security Symposium ({USENIX} Security 17)*, pages 1307–1322, 2017.
[69] Taejoong Chung, Roland van Rijswijk-Deij, David Choffnes, Dave Levin, Bruce M Maggs, Alan Mislove, and Christo Wilson. Understanding the role of registrars in dnssec deployment. In *Proceedings of the 2017 Internet Measurement Conference*, pages 369–383, 2017.
[70] Tobias Dam, Lukas Daniel Klausner, Damjan Buhov, and Sebastian Schrittwieser. Large-scale analysis of pop-up scam on typosquatting urls. In *Proceedings of the 14th International Conference on Availability, Reliability and Security*, pages 1–9, 2019.
[71] Erik Daniel and Florian Tschorsch. Ipfs and friends: A qualitative comparison of next generation peer-to-peer data networks. *IEEE Communications Surveys & Tutorials*, 24(1):31–52, 2022.
[72] Trinh Viet Doan, Irina Tsareva, and Vaibhav Bajpai. Measuring dns over tls from the edge: Adoption, reliability, and response times. In *International Conference on Passive and Active Network Measurement*, pages 192–209. Springer, 2021.
[73] Yahia Elsayed and Ahmed Shosha. Large scale detection of idn domain name masquerading. In *2018 APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–11. IEEE, 2018.
[74] Bingyu Gao, Haoyu Wang, Pengcheng Xia, Siwei Wu, Yajin Zhou, Xiapu Luo, and Gareth Tyson. Tracking counterfeit cryptocurrency end-to-end. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(3):1–28, 2020.
[75] Scarlett Gourley and Hitesh Tewari. Blockchain backed dnssec. In *International Conference on Business Information Systems*, pages 173–184. Springer, 2018.
[76] Zhi Guan, Abba Garba, Anran Li, Zhong Chen, and Nesrine Kaaniche. Authledger: A novel blockchain-based domain name authentication scheme. In *ICISSP*, pages 345–352, 2019.
[77] Adiseshu Hari and TV Lakshman. The internet blockchain: A distributed, tamper-resistant transaction framework for the internet. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*, pages 204–210, 2016.

[78] Guobiao He, Wei Su, Shuai Gao, and Jiarui Yue. Td-root: A trustworthy decentralized dns root management architecture based on permissioned blockchain. *Future Generation Computer Systems*, 102:912–924, 2020.

[79] Ren He, Haoyu Wang, Pengcheng Xia, Liu Wang, Yuanchun Li, Lei Wu, Yajin Zhou, Xiapu Luo, Yao Guo, and Guoai Xu. Beyond the virus: A first look at coronavirus-themed mobile malware, 2020.

[80] Zhangrong Huang, Ji Huang, and Tianning Zang. Leopard: Understanding the threat of blockchain domain name based malware. In *International Conference on Passive and Active Network Measurement*, pages 55–70. Springer, 2020.

[81] Andrew J Kalafut, Minaxi Gupta, Christopher A Cole, Lei Chen, and Nathan E Myers. An empirical study of orphan dns servers in the internet. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pages 308–314, 2010.

[82] Harry A Kalodner, Miles Carlsten, Paul Ellenbogen, Joseph Bonneau, and Arvind Narayanan. An empirical study of namecoin and lessons for decentralized namespace design. In *WEIS*. Citeseer, 2015.

[83] Enis Karaarslan and Eylul Adiguzel. Blockchain based dns and pki solutions. *IEEE Communications Standards Magazine*, 2(3):52–57, 2018.

[84] Christos Karapapas, Iakovos Pittaras, Nikos Fotiou, and George C Polyzos. Ransomware as a service using smart contracts and ipfs. *arXiv preprint arXiv:2003.04426*, 2020.

[85] Issa M Khalil, Bei Guan, Mohamed Nabeel, and Ting Yu. A domain is only as good as its buddies: Detecting stealthy malicious domains via graph inference. In *Proceedings of the Eighth ACM Conference on Data and Application Security and Privacy*, pages 330–341, 2018.

[86] 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*, CCS '17, pages 569–586, 2017.

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

[88] Jingqiang Liu, Bin Li, Lizhang Chen, Meng Hou, Feiran Xiang, and Peijun Wang. A data storage method based on blockchain for decentralization dns. In *2018 IEEE Third International Conference on Data Science in Cyberspace (DSC)*, pages 189–196. IEEE, 2018.

[89] Yang Liu, Yuwen Zhang, Siyu Zhu, and Cheng Chi. A comparative study of blockchain-based dns design. In *Proceedings of the 2019 2nd International Conference on Blockchain Technology and Applications*, pages 86–92, 2019.

[90] Chaoyi Lu, Baojun Liu, Zhou Li, Shuang Hao, Haixin Duan, Mingming Zhang, Chunying Leng, Ying Liu, Zaifeng Zhang, and Jianping Wu. An end-to-end, large-scale measurement of dns-over-encryption: How far have we come? In *Proceedings of the Internet Measurement Conference*, pages 22–35, 2019.

[91] Constantinos Patsakis and Fran Casino. Hydras and ipfs: a decentralised playground for malware. *International Journal of Information Security*, 18(6):787–799, 2019.

[92] Constantinos Patsakis, Fran Casino, Nikolaos Lykousas, and Vasilios Katos. Unravelling ariadne's thread: Exploring the threats of decentralised dns. *IEEE Access*, 8:118559–118571, 2020.

[93] Peng Peng, Limin Yang, Linhai Song, and Gang Wang. Opening the blackbox of virustotal: Analyzing online phishing scan engines. In *Proceedings of the Internet Measurement Conference*, pages 478–485, 2019.

[94] Mohamed Rahouti, Kaiqi Xiong, and Nasir Ghani. Bitcoin concepts, threats, and machine-learning security solutions. *IEEE Access*, PP:1–1, 11 2018.

[95] Silvia Sebastian and Juan Caballero. Towards attribution in mobile markets: Identifying developer account polymorphism. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 771–785, 2020.

[96] Lee Joon Sern and Yam Gui Peng David. Typoswype: An imaging approach to detect typo-squatting. In *2021 11th IFIP International Conference on New Technologies, Mobility and Security (NTMS)*, pages 1–5. IEEE, 2021.

[97] Reza Soltani, Uyen Trang Nguyen, and Aijun An. A survey of self-sovereign identity ecosystem. *Security and Communication Networks*, 2021, 2021.

[98] Raffaele Sommese, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, Kimberly C Claffy, and Anna Sperotto. The forgotten side of dns: Orphan and abandoned records. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 538–543. IEEE, 2020.

[99] Jeffrey Spaulding, Shambhu Upadhyaya, and Aziz Mohaisen. The landscape of domain name typosquatting: Techniques and countermeasures. In *2016 11th International Conference on Availability, Reliability and Security (ARES)*, pages 284–289. IEEE, 2016.

[100] Janos Szurdi, Balazs Kocso, Gabor Cseh, Jonathan Spring, Mark Felegyhazi, and Chris Kanich. The Long "Taile" of Typosquatting Domain Names. In *Proceedings of the 23rd USENIX Conference on Security Symposium*, SEC'14, pages 191–206, USA, 2014. USENIX Association.

[101] Dennis Trautwein, Aravindh Raman, Gareth Tyson, Ignacio Castro, Will Scott, Moritz Schubotz, Bela Gipp, and Yiannis Psaras. Design and evaluation of IPFS: a storage layer for the decentralized web. In *Proceedings of the ACM SIGCOMM 2022 Conference*, 2022.

[102] Cynthia Wall. The english auction: narratives of dismantlings. *Eighteenth-Century Studies*, 31(1):1–25, 1997.

[103] Wentong Wang, Ning Hu, and Xin Liu. Blockzone: A blockchain-based dns storage and retrieval scheme. In *International Conference on Artificial Intelligence and Security*, pages 155–166. Springer, 2019.

[104] Hu Wei-hong, A. Meng, Shi Lin, Xie Jia-gui, and L. Yang. Review of blockchain-based dns alternatives. 2017.

[105] Jichuan Zhang, Jianhong Zhai, Ru Yang, and Shuyan Liu. Research on enterprise dns security scheme based on blockchain technology. In *International Conference on Blockchain and Trustworthy Systems*, pages 690–701. Springer, 2019.

[106] Futai Zou, Siyu Zhang, Bei Pei, Li Pan, Linsen Li, and Jianhua Li. Survey on domain name system security. In *2016 IEEE First International Conference on Data Science in Cyberspace (DSC)*, pages 602–607. IEEE, 2016.

APPENDIX FOR THE WORK

Table 6: The 13 additional resolvers we add.

Resolver Name	Address	# of Event Logs
ArgentENSResolver1	0xDa1756Bb923Af5d1a05E277CB1E54f1D0A127890	70,535
OldPublicResolver3	0x5FfC014343cd971B7eb70732021E26C35B744cc4	28,790
OldPublicResolver4	0xD3ddcCDD3b25A8a7423B5bEe360a42146eb4Baf3	6,610
AuthereumEnsResolverProxy	0x4DA86a24e30a188608E1364A2D262166a87fCB7C	10,328
OpenSeaENSResolver	0x9C4e9CCE4780062942a7fe34FA2Fa7316c872956	226
ArgentENSResolver2	0xb23267e7a0DEe4DCBA80C1D2FFDb0270aF76fe80	505
PortalPublicResolver	0x0B3eBEccC00E9CEae2BF3235d558EdA7398BE91E	256
TokenResolver	0x074d58C0a0903d4C7DB9388205232602a0bF9Bf0	152
LoopringENSResolver	0xF58D55F06bB92f083E78bb5063A2DD3544f9B6a3	13,191
ChainlinkResolver	0x122eb74f9d0F1a5ed587F43D120C1c2BbDb9360B	4,539
MirrorENSResolver	0xc11796439c3202f4EF836EB126CC67eB378D52c8	624
ForwardingStealthKeyResolver	0xB37671329ABE589109b0bDD1312cc6ACcF106259	224
PublicStealthKeyResolver	0x7D6888e1a454a1fb375125a1688240e5D761fFa6	467

Table 7: The top-10 holders of ENS squatting names.

Address	Owned Squatting Names (Unexpired)	First Registraion	Owned Suspicious Names (Unexpired)
0xbd21109e2bdcb24c4fbcdc16a4c90f34e81228e2	901 (0)	2018/10/30	40937 (0)
0xa7f3659c53820346176f7e0e350780df304db179	776 (1)	2017/5/19	23537 (7)
0x94048eb0db0ccb7d38219f28ed6522937b339aaf	19 (0)	2017/5/25	5174 (0)
0xae18d320383a3598c65767dfd97c8df8ab465d26	151 (0)	2018/3/4	4469 (0)
0xf5f700e1912b93ad09597bfa22484e01c0035b04	375 (0)	2017/5/11	2866 (0)
0xbcbd4885ee8b2b74249c5ad9b8b668fb256a51b1	269 (226)	2017/5/10	2361 (1668)
0x64372db6405879214a0a76a7f1e9c013fd2fd84b	110 (0)	2018/4/30	2249 (0)
0x000fb8369677b3065de5821a86bc9551d5e5eab9	170 (0)	2017/5/9	2006 (0)
0xd8c958f774de4b671e43f78fd0a04255e2291a13	192 (67)	2017/5/28	1697 (189)
0xd2fa59b040852952bf4b4639edd4d8a718a4598a	175 (1)	2017/6/9	1665 (1)

Table 8: Some examples of expired (sub) domains with records.

.eth Name	Record Type	.eth Name with Subdomains	# of subdomains	Record Type
ammazon.eth	ETH Address	thisisme.eth	706	ETH address, Swarm hash, IPFS hash
wikipediaa.eth	ETH Address	[unknown].eth	360	Swarm hash
instabram.eth	ETH Address	unibeta.eth	154	ETH Address
valmart.eth	ETH Address	eth2phone.eth	61	ETH address
faceb00k.eth	ETH Address	smartaddress.eth	30	ETH address

Figure 15: Some people are still using `thisisme.eth` names.

Table 9: Identified suspicious scam addresses in ENS.

ENS names	Address	Description
valus.smartaddress.eth	ETH: 0x903bb9cd3a276d8f18fa6efed49b9bc52ccf06e5	An airdrop scam
four7coin.eth	BTC: 385cR5DM96n1HvBDMzLHPYcw89fZAXULJP	Reported as a Ponzi scheme by BitcoinAbuse, actually is a Bittrex cold wallet [42]
jessica.chainlinknode.eth, jessica.atethereum.eth, crunk.eth	BTC: 1F1tAaz5x1HUXrCNLbtMDqcw6o5GNn4xqX	Reported to be ransomware address, actually is a old Silkroad seized wallet [3]
okex.tokenid.eth, okb.tokenid.eth	ETH: 0x6ada340863c340cab266f4c6ef5e0067932a8bd8	Fake token of OKEx's OKB
ciaone.eth	ETH: 0x171664573e3969874dba31c35082151ea4f181f3	Uniswap scam token
lira.viewwallet.eth	ETH: 0xcf76f32ebe10139e4370127d5789cdb0750d460d	Uniswap scam token
sale.lidofi.eth	ETH: 0x4e344fa2ac01f1fb53b388fad51427de170241a4	Uniswap scam token
cndao.eth	ETH: 0xd94831a33560cd8c4fcded3e1579ab908b9bafae	Uniswap scam token
main.caketoken.eth	ETH: 0x759b0eb08ffaffef2215ac9865483b5e97a1f23c	Uniswap scam token
xn--vitli-6ve6e.eth	ETH: 0x096dc87c708d96033ab7862b14a6f23c038a9394	A scammer pretending to be Vitalik
xn--vitalik-8mj.eth	ETH: 0xda28b1eb9450978b9e3fd6a98f76a293920ce708	A scammer pretending to be Vitalik
xn--vitlik-5nf.eth	ETH: 0x12ccf4b7010f5b201c8fda0f880f0ba63b1a88f3	A scammer pretending to be Vitalik

(a) bobabet.dcl.eth (a gambling site)

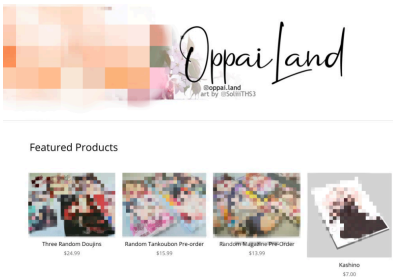

(b) oppailand.eth (an adult book shop)

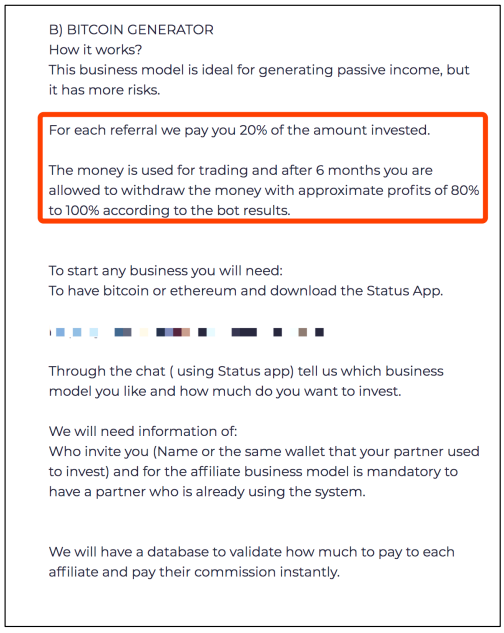

(c) bitcoingenerator.eth (a Ponzi scheme site)

Figure 16: Examples of websites with misbehaviors.

Table 10: The details of events we fetched from ENS contracts. Parameters in description are displayed with typewriter font. Note that, the events of additional resolvers have a similar schema with public resolvers and we do not list them here.

Etherscan Name Tag	Events We Fetched	Main Parameters	Description
Eth Name Service, Registry with Fallback	NewOwner	node, label, owner	A node (domain) registers a label (subdomain)
	NewResolver	node, resolver	A node is set a resolver
	Transfer	node, owner	A node is assigned to a new owner
Old Registrar	AuctionStarted	hash, registrationDate	Start an auction for a hash (.eth subdomain), and register it on registrationDate
	NewBid	hash, bidder, deposit	Bidder bids for a hash with deposit Ether (Possibly higher than the actual bid)
	BidRevealed	hash, owner, value, status	Reveal the owner's value (bid) on hash and its status (1st place, 2nd place, other place, late reveal, low bid)
	HashRegistered	hash, owner, value, registrationDate	The owner registers the hash with value at registrationDate
	HashReleased	hash, value	The owner releases the hash and get a value (refund)
	HashInvalidated	hash, name, value, registrationDate	The hash (unhashed version is name) registered with value at registrationDate is unregistered due to its short length
Base Registrar Implementation, Old ENS Token	NameRegistered	id, owner, expires	The owner registers a id (integer format of a name's labelhash) that will expire at expires
	NameRenewed	id, expires	The id is renewed and will expire on expires
	Transfer	from, to, tokenId	The tokenId (integer format of a name's labelhash) is transferred from from to to
Short Name Claims	ClaimSubmitted	claimed, dnsname, paid, claimnant, email	The dnsname owner (email) uses the claimant (Ethereum address) to claim for claimed (corresponding name on ENS) with paid Ether
	ClaimStatusChanged	claimId, status	The claim request of claimId (request id, generated by hashing claimed, dnsname, claimant, email) has a status change (pending, approved, declined, withdrawn)
Old ETH Registrar Controller 1 & 2, ETHRegistrarController	NameRegistered	name, label, owner, cost, expires	The name (label hash is label) is registered by the owner with cost and will expire at expires
	NameRenewed	name, label, cost, expires	The name (labelhash is label) is renewed with cost and will expire at expires
OldPublicResolver1	ContentChanged	node, hash	The node is set a hash record hash
OldPublicResolver1 & 2 PublicResolver1 & 2	AddrChanged	node, a	The node is set a Ethereum address a
	NameChanged	node, name	The node is set a reverse record name
	ABIChanged	node, contentType	The node is set a ABI record (its type is contentType)
	PubkeyChanged	node, x, y	The node is set a public key record (x, y)
OldPublicResolver2, PublicResolver1 & 2	AddressChanged	node, coinType, newAddress	The node is set an address record (newAdress on coinType blockchain)
	AuthorisationChanged	node, owner, target, isAuthorised	The owner of the node granting/ungranting (based on isAuthorised) target full acccess to his node
	TextChanged	node, indexedKey, key	The node is set a text record with keyword key (indexedKey for search in events)
	InterfaceChanged	node, interfaceID, implementer	The node is set an interface record (the address implementer has an InterfaceID interface on its contract)
	ContenthashChanged	node, hash	The node is set a hash record hash
PublicResolver1 & 2	DNSRecordChanged	node, name, resource, record	The node is set a DNS record for a DNS name. Its record type is resource and its content is record
	DNSRecordDeleted	node, name, resource	The DNS record (record type is resource) of name in node is deleted
	DNSZoneCleared	node	The node's DNS zone information is cleared