Forgotten & Reclaimed: Detecting and Preventing Subdomain Takeover in the Italian Medical Landscape

L. Bracciale[1,2], M. Coni[1], P. Loreti[1,2], E. Raso[1,*] and G. Bianchi[1,2]

[1]Department of Electronic Engineering, University of Rome "Tor Vergata", Rome, Italy
[2]CNIT Network Assurance & Monitoring National Lab, Rome, Italy

Abstract

This work addresses the detection and mitigation of subdomain takeover attacks, with a specific focus on the Italian medical sector. Subdomain takeover occurs when dangling DNS records remain active after a cloud resource is released, enabling malicious actors to assume control of these resources. After designing a novel custom toolchain embedding all the steps necessary to identify, enumerate, and assess subdomains, we conduct a nation-wide scan focused on the Italian medical landscape, analyzing 3,219 domains and over 60,000 subdomains to identify potential risks. A further contribution of this study lies in the use and comparison of different subdomain enumeration techniques, including proprietary and open-source tools, as well as data sources such as Passive DNS and Certificate Transparency logs. Our findings reveal a significant number of exploitable subdomains, offer actionable insights to improve the prevention and mitigation of subdomain takeover attacks, and highlight the opportunity for national-scale initiatives to proactively safeguard critical sectors.

Keywords

Subdomain takeover, dangling resources, vulnerability scanning, nation-wide analysis, lawful scans

1. Introduction

Even today, despite the proliferation of Italian (such as the recent Cybersecurity Law, L. 90/2024) and European regulations (DORA, CRA, Cybersecurity Act, etc.), and the widespread reports of cyberattacks targeting companies and states, the level of awareness among businesses about cyber risks remains insufficient. Many organizations lag behind in understanding the scope of these threats or in adopting and implementing effective mitigation strategies, leaving critical sectors vulnerable to exploitation. According to a recent report by the Bank of Italy on "Cybersecurity in Italian Companies: Risk Perception and Mitigation Practices" [1], "the awareness expressed by companies, however, *is struggling to translate into concrete actions*."

A concrete, relatively cost-effective, and scalable action is the use of nation/world-wide scanning tools such as Shodan, BinaryEdge, Onyphe, LeakIX, or Censys. For example, the United States Cybersecurity and Infrastructure Security Agency (CISA) has repeatedly reported using Shodan to monitor the spread of vulnerabilities [2]. These tools operate without explicit

Joint National Conference on Cybersecurity (ITASEC & SERICS 2025), February 03-8, 2025, Bologna, IT

[*] Corresponding author.

lorenzo.bracciale@uniroma2.it (L. Bracciale); matteo.coni@alumni.uniroma2.eu (M. Coni); pierpaolo.loreti@uniroma2.it (P. Loreti); emanuele.raso@uniroma2.it (E. Raso); giuseppe.bianchi@uniroma2.it (G. Bianchi)

authorization, systematically scanning the entire Internet multiple times a day, mapping attack surfaces, and tagging vulnerabilities.

The threat: subdomain takeover.
In this work, we do not use any of the previously mentioned systems; instead, we develop our own custom toolchain for lightweight, least-invasive, nation-wide scanning, specifically designed to detect a type of attack targeting cloud resources, known as ***subdomain takeover***. This attack is grounded on the fact that cloud resources are often associated with subdomains, serving as unique identifiers that link to specific services or applications. This practice is particularly useful for temporary, short-lived events, such as marketing campaigns or seasonal activities. A typical example is a domain like `event.mycompany.com`, which serves as an alias for `event-mycompany.cloud.com`. Once the event concludes, the cloud resources are released, especially when utilizing pay-per-use pricing models, but the DNS record may be forgotten and left active. When this happens, The subdomain then becomes **dangling**, i.e. it remains active but no longer points to a valid resource. As shown in past research works [3, 4, 5], attackers can exploit dangling subdomains in various ways, using them for a variety of questionable or even harmful purposes. Despite being recognized for nearly a decade, with early research dating back to 2016 [3], many dangling DNS records are still overlooked, leaving their associated subdomains vulnerable to takeover attacks. A striking example of this was presented in a recent DefCon talk [6], which revealed over 66,000 affected top-level domains, including thousands from prominent organizations (Google, Amazon, New York Times, Harvard, MIT, Samsung, Qualys, Hewlett-Packard),

The target: medical domain.
Recognizing the importance of producing results that can be effectively utilized by our national institutions, in this work we aim to understand what is the current posture of Italian domains against this threat. We specifically focus on the Italian medical landscape for several reasons: it is well-defined and clearly delineated (see Section 3), it aligns with our team's existing medical-domain-specific expertise in open data crawling [7, 8], and, most importantly, it is a highly critical sector. Medical domains handle sensitive data, making them highly vulnerable and requiring enhanced security measures. As largely expected, our scans revealed several dangling subdomains. However, beyond identification, our work further aims to develop methodologies and tools to assist national institutions in detecting such vulnerabilities more effectively. To achieve this, our scans employed subdomain information extraction from open, transparent resources (including Certificate Transparency Logs), and scanning techniques similar to web crawling, limited to sending HTTP/GET requests on root paths and performing DNS queries. This approach ensured that the process remained ethical, compliant with existing regulations, and demonstrated that valuable insights can be obtained without resorting to proprietary or intrusive methods.

Contribution. In summary, our contributions are threefold:

- We develop a custom toolchain capable of periodically monitoring a significant portion of resources within the Italian medical sector.
- Using this toolchain, we report results from scanning 3,219 medical domains and analyzing 59,655 subdomains for dangling resources.

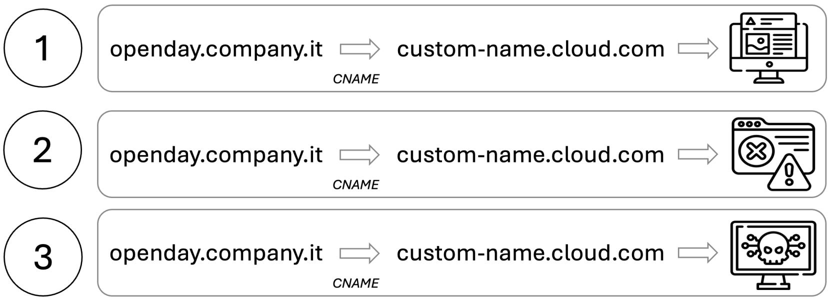

Figure 1: Subdomain takeover in action - A company creates a domain alias pointing to a cloud resource (Step 1). After decommissioning the resource (Step 2), a malicious actor reclaims it (Step 3), gaining control over a resource still linked to the company's domain, enabling misuse.

- As a side contribution, we provide several comparative insights into different subdomain enumeration methods, with a specific focus on the Italian context.

The rest of the paper is organized as follows: Section 2 provides background on the threat; Section 3 describes our methodology and our toolchain; Section 4 presents the results of the scanning activity and assesses the performance of different subdomain enumeration sources; Section 5 highlights lessons learned; finally, conclusions are drawn.

2. Subdomain takeover: basics

To provide a basic understanding of the vulnerability, we illustrate a practical example shown in Figure 1, based on a real-world case identified during our experimental scans. Imagine a company preparing for an Open Day event by creating a new website using a cloud service. This might involve uploading a static site to a cloud storage bucket, such as AWS S3, or using a web-building platform like Wix.com.

When the website goes live, the cloud provider assigns it a URL such as `custom-name.cloud.com`. Typically, users can choose the first part of the URL (e.g., `custom-name`), while the second-level domain (e.g., `cloud.com`) is predefined by the provider.

However, to promote the event more effectively, the company might prefer not to directly use the assigned default URL, and rather create a custom subdomain, such as `openday.company.it`, which points to `custom-name.cloud.com` (**Step 1** in Figure 1).

After the event, the company deactivates the service to cut costs but fails to remove the associated DNS record. Consequently, accessing `openday.company.it` results in an error message indicating that the resource no longer exists (**Step 2**). The domain `openday.company.it` now becomes *dangling*. At this point, an attacker could exploit this situation by identifying the unused domain and registering a new resource with the same `custom-name` on the same cloud provider. This allows the attacker to take control of `custom-name.cloud.com` and, by extension, `openday.company.it` (**Step 3**).

Once the attacker has assumed ownership of the dangling domain, the primary exploitation of this vulnerability lies in leveraging its established reputation for questionable or even malicious purposes, as discussed in the next examples.

Blackhat SEO. According to a recent large-scale study on dangling resources [9], the most common form of exploitation is the so-called Blackhat SEO (Search Engine Optimization), found in 75% of the results therein provided. The attacker's goal is to exploit the reputation of the original domain for improving visibility of dubious or even harmful sites. This involves a range of techniques such as cloaking, where content shown to search engines differs from what users see, or doorway pages designed to rank highly but redirect users to monetized sites. For instance, in the example above, the promotional efforts for the Open Day could inadvertently direct users to a gambling or scam site. Click-jacking is another tactic, manipulating user clicks to trigger malicious actions or generate traffic for targeted websites.

Phishing, Scams, Malware Distribution. By exploiting the user's trust in the subdomain—perceived as associated to a legitimate domain—, Attackers can steal credentials or sensitive information by hosting phishing pages, distribute malware, or impersonate legitimate websites. For instance, under certain circumstances, it is possible to take over a dangling e-commerce website created using the popular Shopify platform [10], which currently supports over 2 million merchants and generates more than $7 billion in annual revenue [11].

Spread of Misinformation. Compromised subdomains can disseminate false information and harm public opinion, leveraging the trust in hijacked entities, potentially including governments or media organizations. For instance, a satirical article falsely claiming war tensions was published on a New York Times production domain [6], demonstrating the potential for public confusion and reputational harm.

Identity Theft. Hijacked subdomains can exploit misconfigured Cross-Origin Resource Sharing (CORS) or Content Security Policies (CSP) to bypass same-origin protections. Attackers can inject scripts to steal cookies or hijack sessions, gaining access to personal accounts. Additionally, compromising MX (Mail Exchange) records enables attackers to redirect email traffic, intercepting password resets, two-factor authentication (2FA) codes, etc.

3. Methodology

3.1. Dataset preparation

Public Healthcare Facilities: We started with the open dataset *"Index of Digital Addresses of Public Administrations and Public Service Managers"*, managed by the Agency for Digital Italy (AgID)[1]. Italian public administrations are required to update this dataset promptly and at least semiannually, as mandated by Article 6-ter of the Digital Administration Code (Legislative Decree March 7, 2005, No. 82). The dataset contains 23,590 entries, 96% of which include a website domain. To focus on the healthcare sector, we filtered the dataset using the following criteria: *ATECO codes*[2], medical/hospital-related keywords, and legal nature codes. This filtering process resulted in 404 entries, representing public healthcare facilities, including major hospitals and inpatient care centers across Italy.

Private Healthcare Facilities: We expanded our analysis to include private healthcare facilities. According to Article 41 of Legislative Decree 33-2013, titled *"Transparency of the National Health*

[1] https://indicepa.gov.it/ipa-dati/dataset/enti

[2] The ATECO code in Italy refers to the national classification system used to categorize economic activities

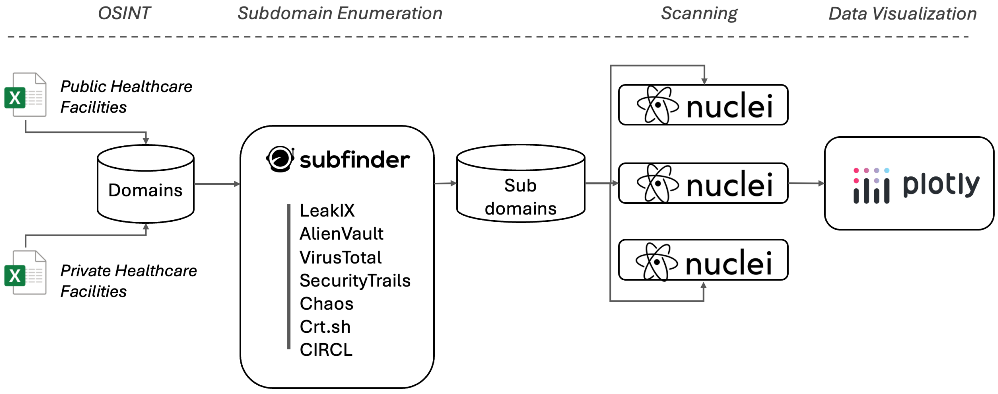

Figure 2: Description of the adopted toolchain

Service", each Italian region is required to publish a list of accredited private healthcare facilities on its official website. However, these regional datasets often lack specific domain information for the facilities. To overcome this limitation, we employed data mining techniques to identify the corresponding domain names for each private healthcare facility. This process allowed us to compile a comprehensive dataset of 3,219 Italian medical domains.

3.2. Subdomain enumeration

For each domain, we identified all its associated subdomains. We utilized the tool "subfinder" [12], which aggregates multiple subdomain enumeration services. Specifically, we leveraged Leakix, Subdomaincenter, Alienvault, VirusTotal, SecurityTrails, Digitorus, Chaos, and Crtsh. Due to API restrictions, payment requirements, and trial limitations, only the last two services were available for free use. For each subdomain identified, we verified its existence by querying the DNS for the corresponding name and discarded those that returned an NXDOMAIN response. Ultimately, out of 90,624 *total* subdomains extracted, we obtained a list of 59,655 *valid* subdomains.

3.3. Dangling resource scanning

We examined each (valid) subdomain for dangling resources by sending simple HTTP/GET requests and analyzing the responses. These responses were matched against specific fingerprints associated with different cloud providers. For example, a dangling resource on Microsoft Azure might return an NXDOMAIN response, while an AWS S3 bucket might generate an error such as "The specified bucket does not exist." This analysis was tailored to each cloud provider. We conducted these checks across 73 different providers using the Nuclei "templates" [13].

3.4. Toolchain description

We integrated all the above steps into a toolchain highlighted in Figure 2. In the initial *OSINT phase*, crawlers collect domain data from public and private healthcare facilities, which is then stored in a MongoDB database. During the subsequent *subdomain enumeration phase*, the

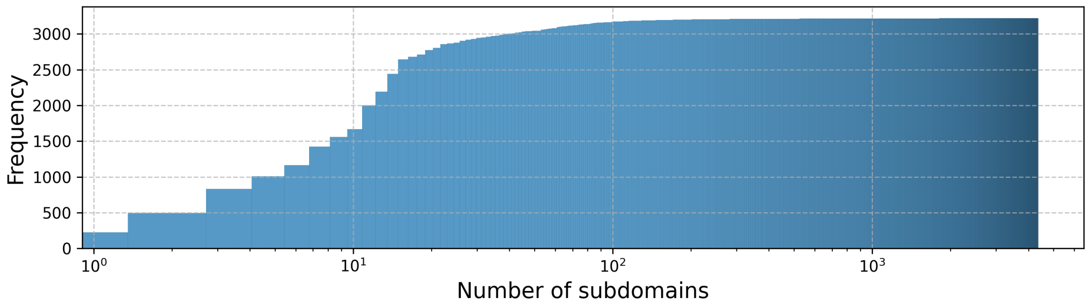

Figure 3: Distribution of the number of subdomains for Italian medical domains

tool `subfinder` queries seven distinct sources for each domain, uncovering all associated subdomains and adding them to the database. In the *scanning phase*, multiple instances of Nuclei perform HTTP/GET requests to each subdomain, comparing the responses against predefined fingerprints to identify potential subdomain takeover vulnerabilities. Finally, Plotly is used in the *data visualization phase* for analyzing results.

4. Results

4.1. Subdomain statistics

Starting with 3,219 domains, we identified a total of 59,655 valid subdomains, whose distribution is illustrated in Figure 3. Among the 3,219 domains examined, 50% have fewer than 11 subdomains, 80% have fewer than 17, and a mere 5 domains collectively account for 10,145 subdomains. Almost 1% of these valid subdomains, specifically 566, were found to be dangling, i.e., characterized by CNAME records resolving to `NXDOMAIN` responses. Further analysis revealed that 4 of these subdomains had exploitable vulnerabilities. Notably, during the investigation, one of these subdomains—associated with a prominent Italian medical institution—was actively taken over by an Indonesian e-commerce website, as illustrated in Figure 7 (appendix 1).

To better understand dependencies on cloud infrastructure, we analyzed CNAME records to identify subdomains pointing to cloud-based domains. We found 704 subdomains pointing to cloud resources: 168 to Amazon Web Services, 418 to Microsoft Azure, and 107 to Italian national cloud providers. It is important to note that AWS and Azure are known to be susceptible to subdomain takeover attacks [10] (Italian cloud providers were not addressed in that study).

4.2. Subdomain Enumeration methods: comparison

A comprehensive list of subdomains serves as the baseline for our analysis. Given that the effectiveness of discovery tools apparently varies across countries and domains, we focused on identifying the sources best suited for the Italian healthcare domain. To ensure reproducibility, we restricted our analysis to the 404 officially provided domains representing public healthcare facilities, as detailed in Section 3. These entries provide a reliable baseline for our analysis, as

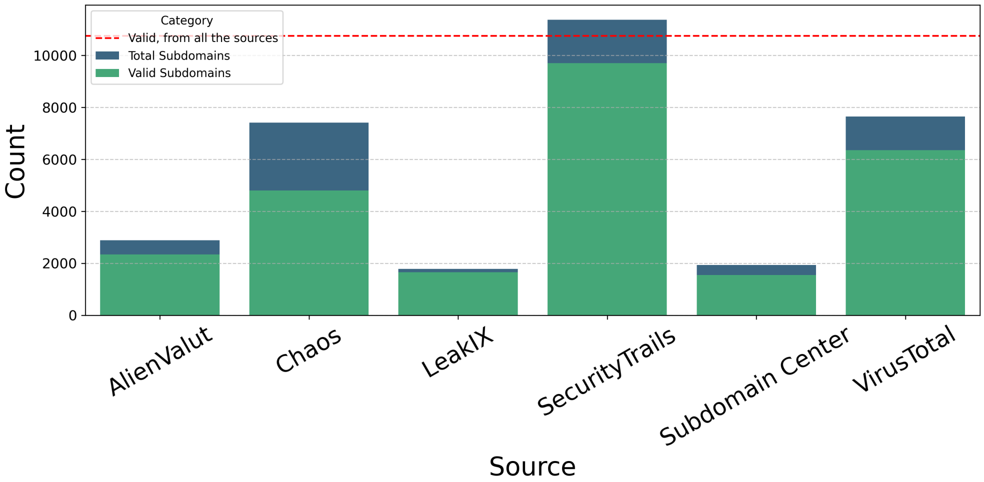

Figure 4: Subdomains returned by each source: total vs valid entries

the dataset is legally required to be updated at least semiannually and includes domain names. We then conducted separate analyses using the following sources:

- **LeakIX**: a platform designed to discover exposed services and leaks in public-facing infrastructures, helping identify misconfigurations and security risks;
- **Subdomain Center**: a service offering subdomain enumeration and monitoring, assisting in tracking DNS records and identifying related assets;
- **AlienVault**: a threat intelligence platform that provides the Open Threat Exchange (OTX), sharing indicators of compromise (IoCs) and information about malicious activity;
- **VirusTotal**: a widely-used online service that aggregates results from multiple antivirus engines, tools, and intelligence sources to analyze files and URLs for malicious content;
- **Security Trails**: a comprehensive database offering DNS records, domain history, and associated infrastructure for cyber threat intelligence and asset management;
- **Chaos**: a public dataset from Project Sonar that includes subdomains and DNS data, supporting research on internet-facing assets;

Figure 4 illustrates the total number of subdomains discovered for each source. A notable disparity is observed, with SecurityTrails identifying more subdomains, followed by VirusTotal and Chaos.

However, not all the domains provided by the sources were valid. Specifically, querying some domains resulted in a DNS non-existent domain (NXDOMAIN) response. Figure 4 illustrates the proportion of "Valid Subdomains" (i.e., domains that do not return an NXDOMAIN response) relative to the total number of subdomains listed by each source. As shown, some sources return a significant number of non-existent domains, which are excluded before proceeding to the scanning phase. The aggregation of all sources resulted in the identification of 10,755 unique and valid subdomains.

The graph in Figure 6 (Appendix) shows the number of subdomains discovered for each of the 404 medical domains analyzed. In particular, some domains have more than 500 subdomains. For example, there is one health company that has 517 verified subdomains identified by Chaos. Investigating such domains, it results that many of those are linked to activities such as testing, open days, monitoring, vaccinations, or past events (e.g., Christmas 2023).

4.3. Certificate Transparency: a free source for subdomain enumeration?

We compared our analysis against a "special" source, CRT.sh,a specialized tool for searching Certificate Transparency (CT) logs. As well known, Certificate Transparency (CT) is *not* meant to help enumeration, but is a security framework that enhances the Public Key Infrastructure (PKI) by providing publicly accessible, cryptographically secure, tamper-evident logs of issued SSL/TLS certificates [14]. Its primary goal is to mitigate risks from mis-issued or malicious certificates, enabling domain owners and the public to audit certificates in real time. Web browsers and TLS clients enforce CT compliance by verifying Signed Certificate Timestamps (SCTs) in the certificate chain, rejecting certificates without valid SCTs.

CT has gained significant traction across the web. Its enforcement has been mandatory in Google Chrome since 2018, by 2019, over 60% of HTTPS traffic supported CT [15], and its adoption continues to grow. Given the widespread use of HTTPS [16], with over 99% of Chrome browsing time now occurring over HTTPS [17], we can ask whether scraping CT logs offers a viable, free solution for obtaining an maintaining an up-to-date list of newly issued subdomains without relying on feeds from cybersecurity companies.

In our analysis, illustrated in Figure 5, we compared the total number of subdomains identified across all sources with those discovered exclusively through CT logs retrieved through CRT.sh. In the set of tested domains, CT represented approximately 28% of the subdomains, identifying 2,970 out of the total 10,755 subdomains.

We examined which subdomains were excluded and found that they were primarily non-web domains, such as those related to mail, VoIP, and WebSockets. In some cases, wildcard domains were registered, making CT logs unsuitable for enumerating the subdomains. This suggests that while CT logs are useful, they cannot be solely relied upon for subdomain enumeration, and integrating multiple sources is essential for comprehensive results.

Furthermore, we examined the presence of dangling domains within the CT logs and identified 57 entries. Notably, all four exploitable domains were included in this source.

4.4. Comparison with Passive DNS Replica

We compared the subdomains identified through various sources with those retrieved from a Passive DNS replica. Passive DNS (pDNS) is a technique that logs DNS responses in a way that minimizes privacy concerns for users, introduced by Florian Weimer in 2005 [18]. It is widely used by security researchers to investigate malware, identify command and control servers, detect phishing domains, and explore subdomains [19].

Unlike traditional DNS, pDNS captures only "cache fill" responses from authoritative servers, which are triggered by recursive resolvers. This approach does not expose user device IPs or ports, making it a privacy-friendly method to identify connections to known malicious domains.

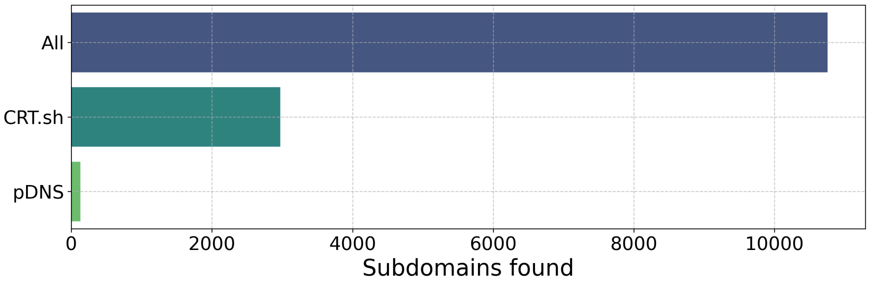

Figure 5: Subdomain enumeration: Certificate Transparency vs CIRCL pDNS vs other sources

pDNS records the DNS responses that were ultimately served to users, which may differ from those provided to researchers or malicious actors. For example, authoritative servers may offer different responses based on the request source to optimize performance [20] or to hide malicious activity when queried from known security research IPs.

While many sources incorporate pDNS data, we specifically compared it with the CIRCL pDNS service [21], which is part of a government-driven initiative. CIRCL aggregates historical DNS records from various sources, including malware analysis, and adheres to the Passive DNS - Common Output Format [22].

We requested subdomain enumeration from the dataset of 404 Italian medical domains, but only 128 subdomains were identified through the CIRCL pDNS service. This represents just 1% of the total number of subdomains discovered using other sources, as shown in Figure 5.

This finding highlights a significant gap: many Italian subdomains are not recognized by the Luxembourg-based pDNS replica. It underscores the need to replicate successful open initiatives like CIRCL's pDNS in Italy to improve the coverage and accuracy of subdomain discovery.

5. Discussion

From this analysis, we identify three key emerging needs.

Large-scale scanners for cybersecurity hygiene: There is a need for light-weight large-scale scanning tools integrated with efficient reporting systems, to promptly address cybersecurity issues and improve national cybersecurity hygiene. These tools should operate (semi)automatically across multiple domains without requiring authorization to remain practical.

Italian Passive DNS: National-scale tools and initiatives are needed to create specific assets, such as passive DNS replica systems [21], populated with local data—current foreign-based systems apparently fail to capture the nuances of Italy's cybersecurity landscape.

National Clouds: Although such attacks are highly impactful, they can be effectively detected and mitigated with appropriate systems in place. These threats highlight the critical need for proactive measures to prevent potential takeovers. Notably, Italian cloud infrastructures, despite not being included in any current Nuclei templates, host numerous critical strategic services and therefore deserve greater attention.

6. Related Work

Dangling DNS records and subdomain takeover vulnerabilities are extensively studied in the literature, with their associated risks and impacts analyzed in numerous works. In [9], the authors developed a methodology to detect and analyze real-world abuse of dangling DNS records. Over a three-year period, they collected a dataset of (sub)domains pointing to deallocated cloud assets, starting with 1,508,273 records and expanding to 3,101,992 by the end of the study. They identified 20,904 cases of abused dangling records, with the majority (75%) used for SEO purposes. Attackers leveraged domains with established reputations to boost the ranking of malicious content in search engines. The remaining cases involved cookie theft, fraudulent certificate issuance, and malware distribution. In a similar study [23], using passive DNS traffic, the authors analyzed over 130 million domains pointing to cloud network IP addresses. They discovered more than 700,000 domains that were unclaimed and vulnerable to takeover attacks.

A comprehensive analysis of subdomain takeover vulnerabilities is presented in [24], where the authors evaluated the causes of such vulnerabilities and described various attack scenarios. They also discussed the operational aspects of these attacks, including subdomain enumeration and takeover execution, and proposed a model for prevention. An attacker type known as a related-domain attacker is examined in [25]. Such attackers can control a sibling domain of a target web application, often through subdomain takeover. The authors conducted an in-depth analysis of the attacker's capabilities, demonstrating how these can be exploited to compromise the security of the target application.

Solutions to detect and mitigate subdomain takeover threats are proposed in [26] and [27]. In [26], the authors introduced a system to identify subdomain takeover threats and alert the relevant organizational authorities. Their model evaluates CNAME records and HTTP responses, triggering alerts on communication platforms like Slack, Discord, and Telegram. In [27], two distinct techniques are presented: a query-based approach using machine learning to verify existing subdomains and a method for identifying potentially risky subdomains.

7. Conclusion

This study highlights the growing risk of subdomain takeover attacks, especially within Italian medical institutions, and demonstrates the effectiveness of a custom toolchain for large-scale scanning. By leveraging multiple subdomain enumeration sources and passive DNS logs, we identified thousands of potential vulnerabilities in critical sectors. Our findings emphasize the need for proactive monitoring and fast response mechanisms to mitigate the impact of these attacks. Despite increasing awareness and the establishment of regulatory frameworks, there remains a significant gap in the adoption of effective mitigation strategies. This research suggests that national-scale initiatives, such as the development of localized passive DNS replicas, are crucial for enhancing the cybersecurity posture of institutions and industries vulnerable to subdomain hijacking. Future work should focus on improving detection capabilities, expanding the toolchain to address additional attack vectors, and fostering collaboration between governmental and private sectors to strengthen defenses against these emerging threats.

Acknowledgments

This work has been partially supported by the PRECISION project (CUP H73C22001600002) funded by Sardinia Region, Italy, and by project “SEcurity and RIghts In the CyberSpace – SERICS, SPOKE 9” (PE00000014 - CUP B53C22003990006) under the National Recovery and Resilience Plan (NRRP) funded by the European Union – NextGenerationEU.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to: grammar and spelling check, paraphrase and reword. After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

References

[1] B. d’Italia, Qef 2024-0852, https://www.bancaditalia.it/pubblicazioni/qef/2024-0852/index.html?dotcache=refresh, 2024. Accessed: 2024-12-02.

[2] Cy2021 cy2021 administrative subpoena for vulnerability notification year in review, 2024. URL: https://www.cisa.gov/sites/default/files/2023-01/CY2021_Admin_Subpoena_Summary_Factsheet_FINAL.pdf.

[3] D. Liu, S. Hao, H. Wang, All your dns records point to us: Understanding the security threats of dangling dns records, in: Proc. of the 2016 ACM SIGSAC Conf. on Computer and Commun. Security, CCS ’16, 2016.

[4] E. Alowaisheq, S. Tang, Z. Wang, F. Alharbi, X. Liao, X. Wang, Zombie awakening: Stealthy hijacking of active domains through dns hosting referral, in: Proc. of the 2020 ACM SIGSAC Conf. on Computer and Commun. Security, CCS ’20, 2020.

[5] M. Zhang, X. Li, B. Liu, J. Lu, Y. Zhang, J. Chen, H. Duan, S. Hao, X. Zheng, Detecting and measuring security risks of hosting-based dangling domains, Proc. ACM Meas. Anal. Comput. Syst. 7 (2023).

[6] B. Demirkapi, Secrets & shadows: Leveraging big data for vulnerability discovery, Proceedings of DEFCON 32, Las Vegas, talk avaliable at https://media.defcon.org/DEF CON 32/DEF CON32video and slides/ (2024).

[7] L. Bracciale, P. Loreti, G. Bianchi, Cybersecurity vulnerability analysis of medical devices purchased by national health services, Scientific Reports 13 (2023). Nature Publishing Group UK London.

[8] L. Bracciale, P. Loreti, E. Raso, G. Bianchi, et al., In plain sight: A pragmatic exploration of the italian medical landscape (in) security, in: ITASEC 2024 - CEUR WORKSHOP PROCEEDINGS, volume 3731, 2024.

[9] J. Frieß, T. Gattermayer, N. Gelernter, H. Schulmann, M. Waidner, Cloudy with a chance of cyberattacks: Dangling resources abuse on cloud platforms, in: 21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24), 2024, pp. 1977–1994.

[10] EdOverflow, Can i take over xyz?, https://github.com/EdOverflow/can-i-take-over-xyz, 2024. Accessed: 2024-12-02.

[11] T. S. Shepherd, Shopify statistics: Key Data, Facts, & Trends, https://thesocialshepherd.com/blog/shopify-statistics, 2024. Accessed: 2024-12-02.

[12] ProjectDiscovery, Subfinder: A subdomain discovery tool, https://github.com/projectdiscovery/subfinder, 2024. Accessed: 2024-12-02.

[13] ProjectDiscovery, Nuclei: Fast, flexible, and customizable vulnerability scanning, https://github.com/projectdiscovery/nuclei, 2024. Accessed: 2024-12-02.

[14] B. Laurie, Certificate transparency, Communications of the ACM 57 (2014) 40–46.

[15] E. Stark, R. Sleevi, R. Muminovic, D. O'Brien, E. Messeri, A. P. Felt, B. McMillion, P. Tabriz, Does certificate transparency break the web? measuring adoption and error rate, in: 2019 IEEE Symposium on Security and Privacy (SP), 2019, pp. 211–226. doi:10.1109/SP.2019.00027.

[16] A. P. Felt, R. Barnes, A. King, C. Palmer, C. Bentzel, P. Tabriz, Measuring {HTTPS} adoption on the web, in: 26th USENIX security symposium (USENIX security 17), 2017, pp. 1323–1338.

[17] Google, Https transparency report, https://transparencyreport.google.com/https/overview, 2024. Accessed: 2024-12-02.

[18] Enyo, Dnslogger: Passive dns logging and analysis, https://www.enyo.de/fw/software/dnslogger/first2005-paper.pdf, 2005. Accessed: 2024-12-02.

[19] SecurityTrails, Understanding passive dns: How it works and why it matters, https://securitytrails.com/blog/passive-dns, 2024. Accessed: 2024-12-02.

[20] IETF, Rfc 7871 - client subnet in dns queries, https://tools.ietf.org/html/rfc7871, 2016. Accessed: 2024-12-02.

[21] CIRCL, Passive dns service, https://circl.lu/services/passive-dns/, 2024. Accessed: 2024-12-02.

[22] Y. Dulaunoy, Passive dns format, https://datatracker.ietf.org/doc/html/draft-dulaunoy-dnsop-passive-dns-cof, 2024. Accessed: 2024-12-02.

[23] K. Borgolte, T. Fiebig, S. Hao, C. Kruegel, G. Vigna, Cloud strife: Mitigating the security risks of domain-validated certificates, in: Proceedings of the 2018 Applied Networking Research Workshop, ANRW '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 4. URL: https://doi.org/10.1145/3232755.3232859. doi:10.1145/3232755.3232859.

[24] S. Z. U. Rashid, M. I. Kamrul, A. Islam, Understanding the security threats of esoteric subdomain takeover and prevention scheme, in: 2019 International Conference on Electrical, Computer and Communication Engineering (ECCE), IEEE, 2019, pp. 1–4.

[25] M. Squarcina, M. Tempesta, L. Veronese, S. Calzavara, M. Maffei, Can i take your subdomain? exploring same-site attacks in the modern web, in: 30th USENIX Security Symposium (USENIX Security 21), 2021, pp. 2917–2934.

[26] M. Biswas, S. P. Singh, S. K. Shaha, Detecting subdomain takeover threats and real-time alerting for rapid response, in: 2023 26th International Conference on Computer and Information Technology (ICCIT), IEEE, 2023, pp. 1–6.

[27] Y. Wang, Z. Li, T. Wu, I. Duncan, Q. Lyu, An empirical study: automated subdomain takeover threat detection, in: 2021 International Conference on Cyber Situational Awareness, Data Analytics and Assessment (CyberSA), IEEE, 2021, pp. 1–10.

Appendix 1

Figure 6 presents a scatterplot illustrating the distribution of subdomains identified across the 404 targeted Italian healthcare domains.

Figure 7 illustrates an example of a real subdomain takeover that we observed, involving a subdomain belonging to a major Italian medical institution. In details, during scans performed until Friday December 6, 2024, we identified such subdomain as dangling. A further scan performed the day after, on Saturday December 7, revealed that such subdomain had been taken over by an unauthorized party. We promptly notified the national CSIRT, and the incident was swiftly resolved. Without proactive nation-wide scanning activities such as the ones performed in this work, we suspect this compromised subdomain might have remained active and unnoticed by Italian stakeholders for a significantly longer period.

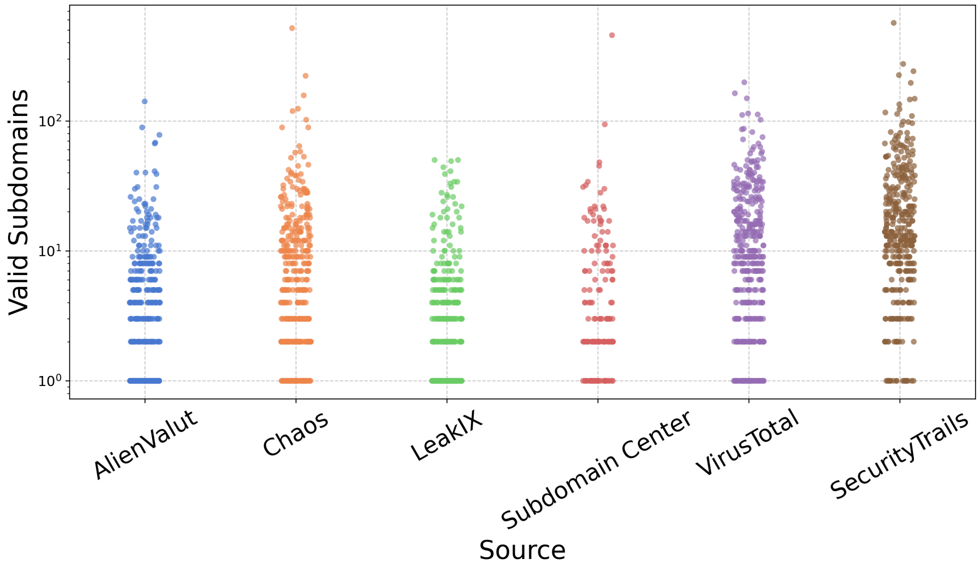

Figure 6: Number of subdomains discovered for each analyzed domain

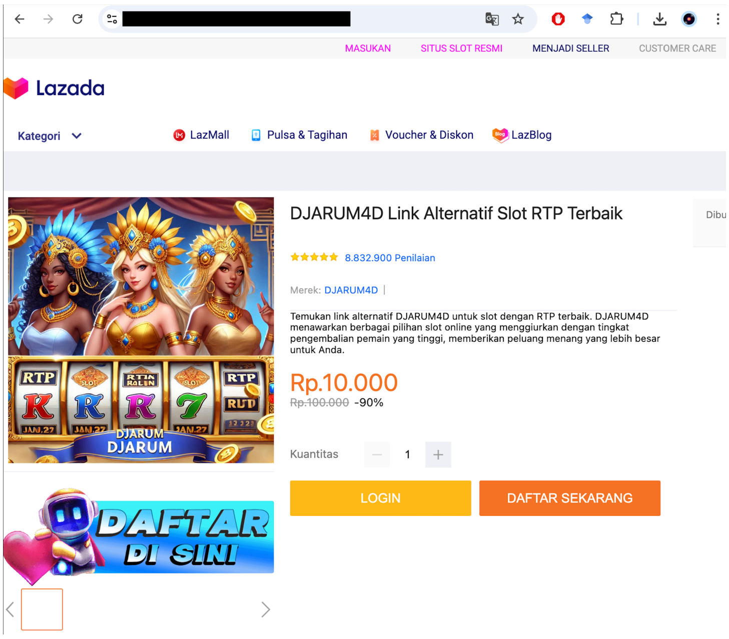

DJARUM4D Link Alternatif Slot RTP Terbaik

★★★★★ 8.832.900 Penilaian

Merek: DJARUM4D |

Temukan link alternatif DJARUM4D untuk slot dengan RTP terbaik. DJARUM4D menawarkan berbagai pilihan slot online yang menggiurkan dengan tingkat pengembalian pemain yang tinggi, memberikan peluang menang yang lebih besar untuk Anda.

Rp.10.000

~~Rp.100.000~~ -90%

Kuantitas − 1 +

LOGIN

DAFTAR SEKARANG

Figure 7: Subdomain of a medical institution taken and exploited