# Toward Detecting Stealthy DNS Tunneling Attacks: A Dataset Analysis Perspective

Naotake Ishikura*, Daishi Kondo†, and Hideki Tode*
*Graduate School of Informatics, Osaka Metropolitan University, Osaka, Japan
†Information Technology Center, The University of Tokyo, Tokyo, Japan
Email: st23937s@st.omu.ac.jp, daishi.kondo@nc.u-tokyo.ac.jp, tode@omu.ac.jp

***Abstract*—The domain name system (DNS) has increasingly been used to establish covert communication channels for malware command-and-control and data exfiltration; this is known as DNS tunneling. Numerous DNS tunneling detection methods have been proposed and shown to be effective against known tunneling tools. However, existing approaches mostly rely on features that can be easily obfuscated or manipulated by sophisticated tunneling techniques to mimic legitimate DNS traffic. As advanced persistent threat (APT) campaigns have demonstrated the practical application of stealthy DNS tunneling, we examined the validity of publicly available DNS tunneling datasets from the perspective of both legitimate and tunneling traffic. The focus was on stealthy tunneling scenarios relevant to APTs, in which tunneling traffic blends into legitimate DNS behavior through low-throughput communication patterns. The analysis revealed that existing datasets contain few samples representing practical stealthy tunneling scenarios relevant to APTs and that the statistical and temporal characteristics of legitimate DNS traffic vary substantially depending on how such traffic is defined and collected. These discrepancies between existing datasets and operational DNS environments can limit the realism of the evaluations, potentially resulting in an overestimation of detection performance. Therefore, we argue that a reliable evaluation requires datasets that inject the tunneling traffic generated under practical APT scenarios into the DNS traffic collected from real-world operational environments.**

***Index Terms*—Dataset analysis, Domain name system (DNS), Low-throughput communication, Stealthy DNS tunneling.**

## I. Introduction

The domain name system (DNS) has increasingly been exploited by attackers to establish covert communication channels between malware deployed within victim networks and authoritative external DNS servers. This technique, commonly referred to as *DNS tunneling*, enables adversaries to conceal command-and-control communications and exfiltrate data by embedding information into fully qualified domain names (FQDNs) of DNS queries or resource records of DNS responses. For example, malware acting as a DNS tunneling client may encode exfiltrated data into the prefix of an FQDN, such as *(outbound_data)*`.attacker.net`, where `attacker.net` denotes the attacker-controlled domain. The corresponding authoritative DNS server extracts the embedded data from the query and returns commands to the malware through the response. As DNS traffic is permitted in most operational environments, DNS tunneling has become attractive for malicious activities, enabling attackers to bypass conventional network security mechanisms.

Numerous countermeasures have been proposed for detecting DNS tunneling [1]–[22]. Although these methods have been demonstrated to be effective for specific DNS tunneling tools, many existing approaches rely on features that can be easily obfuscated or manipulated by advanced tunneling techniques to mimic legitimate DNS traffic. In particular, advanced persistent threat (APT) campaigns leveraging DNS tunneling have been observed [23], [24], suggesting that attackers may develop and adopt increasingly stealthy tunneling techniques to evade existing detection methods. Therefore, future research should focus on detection techniques that can effectively identify stealthy DNS tunneling attacks.

In this study, the validity of publicly available datasets used for DNS tunneling detection [20], [22], [25], [26] was investigated from the perspective of both legitimate and tunneling traffic, focusing on stealthy DNS tunneling scenarios relevant to APTs. We assumed that APT-oriented DNS tunneling attacks are designed to evade detection by blending legitimate DNS behavior with low-throughput communications. However, existing public datasets contain few samples that represent practical stealth tunneling scenarios relevant to APTs, making it difficult to evaluate detection methods against such threat models. Our analysis also showed that the statistical and temporal characteristics of legitimate DNS traffic vary substantially depending on the definition and collection method, that is, whether it is observed in real-world operational environments or generated through simulated user interactions. These discrepancies between existing datasets and operational DNS environments can limit the realism of evaluations, potentially resulting in an overestimation of detection performance, particularly for low-throughput stealthy tunneling communications. Consequently, a reliable evaluation of stealthy DNS tunneling detection requires datasets that are constructed by injecting DNS tunneling traffic generated under practical APT scenarios into real-world DNS traffic collected from operational environments.

## II. Dataset Analysis

### A. Dataset Description

Žiža *et al.* [20] noted that numerous datasets used in

This study was supported in part by JST SPRING (Grant Number JPMJSP2139) and in part by the Telecommunications Advancement Foundation.

978-3-903176-83-6 ©2026 IFIP

related studies were unavailable because of privacy concerns or broken links, making it impossible to compare the research results. Among the publicly available datasets, we analyzed (i) CIC-Bell-DNS-EXF-2021 (CIC) [25], [26], also used by Wang *et al.* [19] and Gao *et al.* [22], (ii) the public dataset by Žiža *et al.*, also used by Ozery *et al.* [21], and (iii) another by Gao *et al.*, which is the most recent. These datasets provide legitimate DNS and DNS tunneling traffic data.

The CIC dataset contains (i) legitimate DNS traffic collected by sending hypertext transfer protocol (HTTP) requests to the web servers of Alexa's top one million domains using a Python script and (ii) DNS tunneling traffic generated by exfiltrating various file sizes and types using DNSExfiltrator [27]. The periods for legitimate and tunneling traffic generation were separated, and data were collected over five days, comprising 641,642 legitimate DNS queries and 377,679 DNS tunneling queries. Identifying DNS queries for unique FQDNs was infeasible because this dataset does not include FQDN information. Therefore, we investigated only the number of registrable domains[1].

The dataset by Žiža *et al.* was built using approximately 52 million DNS query logs observed over 24 h on a primary DNS server operated by a major Internet service provider (ISP) in Serbia. This dataset comprises legitimate DNS queries from general ISP users and 89,843 intentionally generated DNS tunneling queries using iodine [29] and DNSExfiltrator with various parameters, such as the FQDN length, encoding schemes, and query transmission intervals. Furthermore, to avoid detection, Žiža *et al.* implemented a custom encoding scheme called Base32map, which intentionally reduces the character entropy in DNSExfiltrator to generate more sophisticated DNS tunneling traffic. Additionally, disposable domain names [30], [31] found in the logs that use DNS queries for purposes such as data transfer by anti-malware tools, were classified as malicious traffic because they constitute information leakage. In tunneling queries, the textual representations of 53 FQDNs exceeded the maximum length permitted by DNS specifications [32]. This was possibly caused by the encoding process; therefore, we excluded these queries from our analysis.

The multiple datasets released by Gao *et al.* were constructed for the performance evaluation of anomaly detection, classification of DNS tunneling tools, and wildcard domains. The dataset comprised 2,012,494 legitimate DNS traffic records and 962,859 DNS tunneling traffic records generated using various DNS tunneling tools. Legitimate DNS traffic was generated by invoking distributed browser instances to simulate genuine user interactions with the assigned sublists of Cloudflare's top one million domains. DNS tunneling traffic was collected in a simulated intranet environment through scenarios in which an attacker employed various information-gathering tools after establishing tunnels using ten different DNS tunneling tools, including iodine and dnscat2 [33].

[1]A registrable domain is defined as the portion of a domain name comprising a suffix from the Public Suffix List [28] maintained by the Mozilla Foundation, plus a single label that immediately precedes it (i.e., one label + an effective top-level domain).

### B. Comparative Analysis

The characteristics of legitimate DNS and DNS tunneling traffic in public datasets were analyzed from the perspectives of DNS query transmission throughput and number of registrable domains. Previous studies did not sufficiently consider the potential for obfuscation by attackers. Here, we define DNS query transmission throughput as comprising two metrics: the prefix domain length of unique FQDNs[2] and transmission interval between consecutive unique FQDNs within each registrable domain.

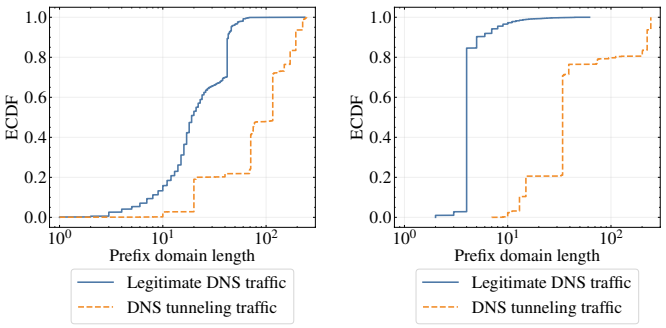


(a) Žiža *et al.* [20] dataset. (b) Gao *et al.* [22] dataset.

Fig. 1. ECDF of prefix domain lengths for unique FQDNs.

Figure 1 shows the empirical cumulative distribution function (ECDF) of the prefix domain lengths for unique FQDNs observed in legitimate DNS and DNS tunneling queries in each dataset. In the Žiža *et al.* dataset, the average prefix domain length is considerably longer for DNS tunneling traffic (103.84) than for legitimate DNS traffic (24.99). The median value for DNS tunneling traffic (116) is more than six times that of legitimate DNS traffic (19). Regarding the 90th and 99th percentiles, the prefix domain lengths for legitimate traffic are small (43 and 60, respectively), whereas those for DNS tunneling traffic are very large (194 and 234, respectively), indicating a remarkable difference. In the Gao *et al.* dataset, the average prefix domain length is considerably longer for DNS tunneling traffic (69.55) than for legitimate DNS traffic (4.63). The median for tunneling traffic (34) is more than eight times that of legitimate traffic (4), indicating a significant divergence. Furthermore, at the 90th and 99th percentiles, the prefix domain lengths are small for legitimate traffic (5 and 16, respectively) and very large for DNS tunneling traffic (224 and 243, respectively). This quantitative analysis confirms that in these public datasets, DNS tunneling traffic communicates with significantly longer prefix domain lengths than legitimate DNS traffic.

Figure 2 shows the ECDF of the transmission intervals between consecutive unique FQDNs per registrable domain for

[2]A prefix domain is the part of the FQDN that remains after removing the registrable domain.

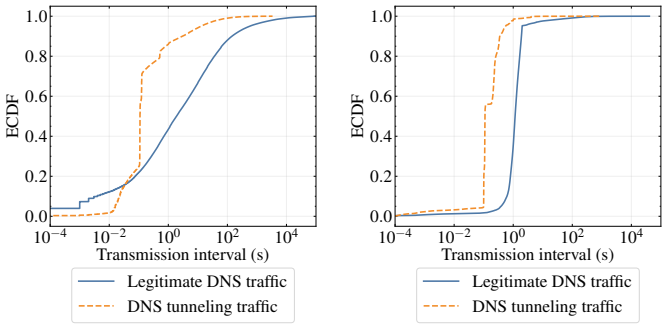


(a) Žiža *et al.* [20] dataset. (b) Gao *et al.* [22] dataset.

Fig. 2. ECDF of transmission intervals between consecutive unique FQDNs per registrable domain.

legitimate and DNS tunneling traffic. In the Žiža *et al.* dataset, the average and median values indicate that the interval between consecutive unique FQDNs is shorter for DNS tunneling traffic (5.25 and 0.11, respectively) than for legitimate traffic (430.37 and 1.89, respectively). Furthermore, the differences between legitimate and tunneling traffic are large at the 90th (142.55, 3.01) and 99th percentiles (9,003.00, 97.43), respectively. In the Gao *et al.* dataset, the average and median values indicate that the transmission interval is extremely long for legitimate traffic (10.30 and 1.16, respectively) and extremely short for DNS tunneling traffic (0.25 and 0.11, respectively). The differences between legitimate and tunneling traffic are also large at the 90th (1.90, 0.35) and 99th percentiles (82.38, 2.21), respectively. Overall, consecutive unique FQDNs appear at shorter intervals in DNS tunneling traffic than in legitimate DNS traffic across the datasets.

TABLE I
COMPARISON OF THE NUMBER OF REGISTRABLE DOMAINS.

| Class | CIC [25], [26] | Žiža *et al.* [20] | Gao *et al.* [22] |
|---|---|---|---|
| Legitimate DNS traffic | 67,625 | 12,833 | 922,475 |
| DNS tunneling traffic | 38 | 3 | 3 |

Table I lists the number of registrable domains for legitimate and DNS tunneling traffic. For DNS tunneling traffic, the CIC dataset uses 38 registrable domains, the highest number among public datasets. However, 35 registrable domains were observed in both legitimate and DNS tunneling traffic, indicating that the two cannot be strictly separated. Furthermore, the volume of exfiltrated data was substantial (168.3 MB), making it impractical to conceal the exfiltration activity by distributing the data across 38 registrable domains because this would still amount to approximately 4.4 MB per domain on average. Note that the Žiža *et al.* and Gao *et al.* datasets use only three registrable domains.

These results have two important implications in evaluating DNS tunneling detection methods. First, DNS tunneling traffic in existing public datasets exhibits conspicuous characteristics, including long prefix domain lengths, short transmission intervals, and a small number of registrable domains. Such traffic may be easier to detect using simple per-domain statistical features; however, it does not sufficiently represent low-throughput stealthy tunneling scenarios in which attackers reduce the amount of information transmitted per query, increase query intervals, and distribute queries across numerous registrable domains. Second, the characteristics of legitimate DNS traffic vary substantially across datasets, indicating that the definitions and collection methodologies are not consistent. For example, the Žiža *et al.* dataset considers ISP-observed DNS traffic as legitimate, whereas the Gao *et al.* dataset generates legitimate traffic through distributed browser-based simulations of the top-ranked domains. Therefore, both the lack of low-throughput tunneling samples and inconsistent definitions of legitimate traffic should be considered when interpreting the evaluation results obtained from public datasets.

### C. Feature Obfuscation in Low-Throughput DNS Tunneling

The characteristics identified above are not inherent to DNS tunneling and can be manipulated by attackers. In low-throughput DNS tunneling, attackers can reduce tunneling traffic visibility by tuning the amount of information embedded in each query, the query transmission interval, and the number of registrable domains used for communication. Specifically, they can take measures, such as maintaining the prefix domain length of the transmitted FQDNs within a range commonly observed in legitimate DNS traffic (e.g., six to nine characters), setting the DNS tunneling query transmission interval of malware that has infiltrated the target network to 1 h or more, preparing numerous registrable domains, and continuously updating the registrable domain after sending a few DNS tunneling queries. Low-throughput DNS tunneling enables data exfiltration while blending with legitimate DNS traffic. However, this stealthiness requires attackers to maintain DNS tunneling for a long period.

## III. CONCLUSION

In this study, we analyzed publicly available datasets for DNS tunneling detection from the perspectives of both legitimate and tunneling traffic. Our analysis showed that existing datasets contain few samples representing low-throughput stealthy tunneling scenarios relevant to APTs and that the characteristics of legitimate DNS traffic vary substantially depending on how such traffic is defined and collected. These findings suggest that evaluations using existing public datasets may be less representative of operational DNS environments and may overestimate the detection performance against stealthy low-throughput tunneling. Therefore, a reliable evaluation requires datasets that inject the tunneling traffic generated under practical APT scenarios into real-world DNS traffic collected from operational environments. In future work, we plan to construct such datasets and evaluate the detection methods under practical low-throughput tunneling scenarios.

## References


[1] K. Born and D. Gustafson, "Detecting DNS tunnels using character frequency analysis," *CoRR*, vol. abs/1004.4358, 2010. [Online]. Available: http://arxiv.org/abs/1004.4358

[2] C. Qi, X. Chen, C. Xu, J. Shi, and P. Liu, "A bigram based real time DNS tunnel detection approach," *Procedia Computer Science*, vol. 17, pp. 852–860, 2013, doi: 10.1016/j.procs.2013.05.109.

[3] M. Aiello, M. Mongelli, and G. Papaleo, "Basic classifiers for DNS tunneling detection," in *Proc. 2013 IEEE Symposium on Computers and Communications (ISCC)*, 2013, pp. 000 880–000 885, doi: 10.1109/ISCC.2013.6755060.

[4] V. Paxson, M. Christodorescu, M. Javed, J. Rao, R. Sailer, D. L. Schales, M. Stoecklin, K. Thomas, W. Venema, and N. Weaver, "Practical comprehensive bounds on surreptitious communication over DNS," in *Proc. 22nd USENIX Security Symposium (USENIX Security 13)*, Aug. 2013, pp. 17–32.

[5] S. Shafieian, D. Smith, and M. Zulkernine, "Detecting DNS tunneling using ensemble learning," in *Proc. Network and System Security*, 2017, pp. 112–127.

[6] J. Liu, S. Li, Y. Zhang, J. Xiao, P. Chang, and C. Peng, "Detecting DNS tunnel through binary-classification based on behavior features," in *Proc. 2017 IEEE Trustcom/BigDataSE/ICESS*, 2017, pp. 339–346, doi: 10.1109/Trustcom/BigDataSE/ICESS.2017.256.

[7] A. Das, M.-Y. Shen, M. Shashanka, and J. Wang, "Detection of exfiltration and tunneling over DNS," in *Proc. 2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2017, pp. 737–742, doi: 10.1109/ICMLA.2017.00-71.

[8] A. Almusawi and H. Amintoosi, "DNS tunneling detection method based on multilabel support vector machine," *Security and Communication Networks*, vol. 2018, no. 1, p. 6137098, 2018, doi: 10.1155/2018/6137098.

[9] I. Homem, P. Papapetrou, and S. Dosis, "Information-entropy-based DNS tunnel prediction," in *Proc. Advances in Digital Forensics XIV*, 2018, pp. 127–140.

[10] J. Steadman and S. Scott-Hayward, "DNSxD: Detecting data exfiltration over DNS," in *Proc. 2018 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, 2018, pp. 1–6, doi: 10.1109/NFV-SDN.2018.8725640.

[11] C.-M. Lai, B.-C. Huang, S.-Y. Huang, C.-H. Mao, and H.-M. Lee, "Detection of DNS tunneling by feature-free mechanism," in *Proc. 2018 IEEE Conference on Dependable and Secure Computing (DSC)*, 2018, pp. 1–2, doi: 10.1109/DESEC.2018.8625166.

[12] C. Liu, L. Dai, W. Cui, and T. Lin, "A byte-level CNN method to detect DNS tunnels," in *Proc. 2019 IEEE 38th International Performance Computing and Communications Conference (IPCCC)*, 2019, pp. 1–8, doi: 10.1109/IPCCC47392.2019.8958714.

[13] J. Ahmed, H. Habibi Gharakheili, Q. Raza, C. Russell, and V. Sivaraman, "Monitoring enterprise DNS queries for detecting data exfiltration from internal hosts," *IEEE Transactions on Network and Service Management*, vol. 17, no. 1, pp. 265–279, 2020, doi: 10.1109/TNSM.2019.2940735.

[14] S. Chen, B. Lang, H. Liu, D. Li, and C. Gao, "DNS covert channel detection method using the LSTM model," *Computers & Security*, vol. 104, p. 102095, 2021, doi: 10.1016/j.cose.2020.102095.

[15] N. Ishikura, D. Kondo, V. Vassiliades, I. Iordanov, and H. Tode, "DNS tunneling detection by cache-property-aware features," *IEEE Transactions on Network and Service Management*, vol. 18, no. 2, pp. 1203–1217, 2021, doi: 10.1109/TNSM.2021.3078428.

[16] H. Bai, W. Liu, G. Liu, Y. Dai, and S. Huang, "Application behavior identification in DNS tunnels based on spatial-temporal information," *IEEE Access*, vol. 9, pp. 80 639–80 653, 2021, doi: 10.1109/ACCESS.2021.3085500.

[17] J. Steadman and S. Scott-Hayward, "DNSxP: Enhancing data exfiltration protection through data plane programmability," *Computer Networks*, vol. 195, p. 108174, 2021, doi: 10.1016/j.comnet.2021.108174.

[18] M. Zhan, Y. Li, G. Yu, B. Li, and W. Wang, "Detecting DNS over HTTPS based data exfiltration," *Computer Networks*, vol. 209, p. 108919, 2022, doi: 10.1016/j.comnet.2022.108919.

[19] S. Wang, L. Sun, S. Qin, W. Li, and W. Liu, "KRTunnel: DNS channel detector for mobile devices," *Computers & Security*, vol. 120, p. 102818, 2022, doi: 10.1016/j.cose.2022.102818.

[20] K. Žiža, P. Tadić, and P. Vuletić, "DNS exfiltration detection in the presence of adversarial attacks and modified exfiltrator behaviour," *International Journal of Information Security*, vol. 22, no. 6, pp. 1865–1880, Dec. 2023.

[21] Y. Ozery, A. Nadler, and A. Shabtai, "Information-based heavy hitters for real-time DNS data exfiltration detection," in *Proc. Network and Distributed System Security (NDSS) Symposium*, 2024, doi: 10.14722/ndss.2024.24388.

[22] G. Gao, W. Niu, J. Gong, D. Gu, S. Li, M. Zhang, and X. Zhang, "GraphTunnel: Robust DNS tunnel detection based on DNS recursive resolution graph," *IEEE Transactions on Information Forensics and Security*, vol. 19, pp. 7705–7719, 2024, doi: 10.1109/TIFS.2024.3443596.

[23] "New Threat Actor Group DarkHydrus Targets Middle East Government," Accessed: May 10, 2026. [Online]. Available: https://unit42.paloaltonetworks.com/unit42-new-threat-actor-group-darkhydrus-targets-middle-east-government/

[24] "DNS Tunneling in the Wild: Overview of OilRig's DNS Tunneling," Accessed: May 10, 2026. [Online]. Available: https://unit42.paloaltonetworks.com/dns-tunneling-in-the-wild-overview-of-oilrigs-dns-tunneling/

[25] "CIC-Bell-DNS-EXF 2021 — Datasets — Research — Canadian Institute for Cybersecurity — UNB," Accessed: Sep. 5, 2025. [Online]. Available: https://www.unb.ca/cic/datasets/dns-exf-2021.html

[26] S. Mahdavifar, A. Hanafy Salem, P. Victor, A. H. Razavi, M. Garzon, N. Hellberg, and A. H. Lashkari, "Lightweight hybrid detection of data exfiltration using DNS based on machine learning," in *Proc. the 2021 11th International Conference on Communication and Network Security*, 2022, pp. 80–86, doi: 10.1145/3507509.3507520.

[27] "Arno0x/DNSExfiltrator: Data exfiltration over DNS request covert channel," Accessed: Sep. 7, 2025. [Online]. Available: https://github.com/Arno0x/DNSExfiltrator

[28] "Public Suffix List," Accessed: Sep. 6, 2025. [Online]. Available: https://publicsuffix.org/

[29] "yarrick/iodine: Official git repo for iodine dns tunnel," Accessed: Sep. 7, 2025. [Online]. Available: https://github.com/yarrick/iodine

[30] Y. Chen, M. Antonakakis, R. Perdisci, Y. Nadji, D. Dagon, and W. Lee, "DNS noise: Measuring the pervasiveness of disposable domains in modern DNS traffic," in *Proc. 2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 2014, pp. 598–609, doi: 10.1109/DSN.2014.61.

[31] S. Hao and H. Wang, "Exploring domain name based features on the effectiveness of DNS caching," *SIGCOMM Comput. Commun. Rev.*, vol. 47, no. 1, pp. 36–42, Jan. 2017, doi: 10.1145/3041027.3041032.

[32] P. Mockapetris, "Domain names - implementation and specification," RFC 1035, Nov. 1987, doi: 10.17487/RFC1035.

[33] "iagox86/dnscat2," Accessed: Sep. 7, 2025. [Online]. Available: https://github.com/iagox86/dnscat2