




ISSN : 2321-9416

Indo - American Journal of Mechanical Engineering



www.iajme.org

Email : iajme.editor@gmail.com or editor@iajme.org



DNS-Based Detection of Data Exfiltration in Corporate Networks

Ozgur Koray Sahingoz

Istanbul Kultur University, Computer Engineering Department, 34158 Istanbul, Turkey

Abstract

The Domain Name System (DNS), a foundational component of internet communication, is increasingly exploited as a covert channel for data exfiltration. Traditional firewalls and intrusion detection systems often allow outbound DNS queries by default, making DNS an attractive medium for bypassing perimeter defenses. This paper proposes a machine learning-based detection framework focused on identifying DNS tunneling activity within corporate networks. We extract statistical and structural features from DNS traffic—including domain name entropy, label length, query frequency, and subdomain depth—and use them to train a Support Vector Machine (SVM) classifier. Experimental evaluations include simulated tunneling scenarios using DNScat2 and Iodine, benchmarked against a corporate-like DNS query baseline. Our model achieves a detection accuracy of 93.4% and maintains a low false-positive rate, suitable for real-time anomaly detection. We validate the approach through testbed deployment and evaluate integration with DNS Security Extensions (DNSSEC) and SIEM platforms for alert correlation. The results demonstrate the viability of DNS as a monitored vector for insider threat detection and data leak prevention in enterprise environments.

Keywords: DNS tunneling, data exfiltration, DNS anomaly detection, SVM classifier, DNScat2, Iodine, entropy analysis, subdomain depth, SIEM integration, corporate network security

1. Introduction

In corporate network environments, the Domain Name System (DNS) serves as an essential service for resolving hostnames to IP addresses. Due to its ubiquitous nature and allowance through most outbound firewalls, DNS has emerged as an ideal channel for covert communication, particularly for data exfiltration. Attackers can encode sensitive information into DNS query names and use external authoritative name servers to reconstruct it. Such misuse, known as **DNS tunneling**, enables data theft while appearing as legitimate resolution traffic.

Traditional security mechanisms—such as perimeter firewalls, signature-based intrusion detection systems (IDS), and endpoint protection—often overlook DNS traffic or fail to inspect its content beyond domain-level reputation checks. This oversight is further compounded by the growth of encrypted DNS services (e.g., DoH and DoT) and the proliferation of remote work environments where endpoint control is fragmented.

This paper addresses the challenge of DNS-based data exfiltration by presenting a machine learning-based detection framework that operates on passive DNS traffic. The proposed system uses structural and statistical features of DNS queries to train a **Support Vector Machine (SVM)** classifier capable of distinguishing between normal and malicious traffic. The framework is evaluated against realistic datasets generated using **DNScat2** and **Iodine**, two widely used tunneling tools, and demonstrates high detection accuracy with minimal false positives.

The rest of this paper details the design, implementation, and validation of this detection framework, and discusses integration strategies with existing **SIEM** (Security Information and Event Management) platforms for real-time anomaly alerting.

2. Literature Review



DNS tunneling is not a new threat, but it remains underrepresented in enterprise defense strategies. Several research efforts and tools have attempted to detect or block tunneling activities, yet most fall short in dynamic or large-scale corporate environments.

2.1 Traditional Detection Approaches

Early detection mechanisms relied on predefined blacklists or static patterns, such as detecting known command-and-control (C2) domains or unusual DNS packet sizes. Tools like Snort include DNS anomaly rules, but their effectiveness is limited by high false positives and delayed signature updates.

2.2 Statistical and Behavioral Analysis

Later approaches applied statistical models to identify anomalies in DNS traffic. Krishnan et al. (2013) proposed measuring domain entropy and subdomain depth, showing that tunneled queries tend to be longer and more randomized. Cova et al. (2014) explored packet timing and request volume spikes, which often indicate beaconing or exfiltration behavior.

However, these heuristic methods require manual tuning and can vary by environment, making them difficult to generalize.

2.3 Machine Learning in DNS Security

Recent research has focused on supervised learning techniques, where classifiers are trained on labeled DNS traffic. Support Vector Machines (SVM), Random Forests, and deep neural networks have been explored, with SVMs favored for their ability to handle high-dimensional, sparse features with clear decision boundaries.

Choi et al. (2017) used SVMs to classify domain flux and fast-flux DNS behaviors. Our work extends this idea by targeting exfiltration detection specifically, using features that represent tunneling behavior rather than DNS botnet characteristics.

2.4 Gaps in Current Solutions

Many commercial DNS firewalls focus on domain reputation and fail to detect low-volume, stealthy tunneling. Additionally, few systems integrate real-time anomaly scoring with **SIEM pipelines** for immediate response.

This paper aims to bridge that gap with a lightweight, protocol-aware machine learning model that enhances DNS visibility in enterprise networks.

3. Hypotheses or Research Questions

The following hypotheses guide our study:

- **H1:** DNS tunneling activity can be effectively distinguished from legitimate DNS traffic using statistical and structural features extracted from DNS query names.
- **H2:** An SVM classifier trained on labeled DNS traffic will achieve detection accuracy above 90% with a false-positive rate below 5%.
- **H3:** Features such as domain entropy, subdomain depth, and query length will be among the top contributors to classification performance.
- **H4:** The proposed model can operate in near real-time and integrate with existing SIEM platforms without disrupting DNS resolution.

These hypotheses are evaluated through experimental simulations and testbed deployments in Sections 4 and 5.

4. Methodology

This section details the experimental design, dataset generation, feature selection, model configuration, and evaluation metrics used in developing and validating the DNS exfiltration detection framework.

4.1 Dataset Creation

We constructed a labeled dataset of DNS traffic composed of both benign and tunneling queries:

- **Benign Traffic:** Captured over two weeks from a sandboxed corporate-like network emulating real DNS usage (e.g., Google, Microsoft, Amazon domains).
- **Malicious Traffic:** Generated using:



- **DNScat2**: An open-source tool for creating interactive tunnels over DNS.
- **Iodine**: A tool for sending IP traffic through DNS, commonly used for exfiltration and bypassing firewalls.

Tunneling sessions were scripted to simulate gradual data leaks rather than bulk transfers to avoid detection via sheer volume.

4.2 Feature Extraction

Each DNS query in the dataset was parsed to extract features in four categories:

1. Entropy-Based Features:

- Shannon entropy of the query name
- Entropy of individual subdomain segments

2. Structural Features:

- Total query length
- Maximum subdomain label length
- Subdomain depth (number of labels)

3. Frequency and Timing Features:

- Query repetition rate per second
- Burstiness index (standard deviation of inter-arrival times)

4. Domain Features:

- Use of uncommon top-level domains (e.g., .xyz, .tk)
- NXDOMAIN response rate (failed lookups)

Features were standardized and scaled using Min-Max normalization. Packet captures were analyzed using Wireshark and custom Python scripts built with Scapy and dnspython.

4.3 Model Configuration

We employed a **Support Vector Machine (SVM)** classifier using the radial basis function (RBF) kernel. Hyperparameters were tuned using grid search:

- C (regularization): 1.0
- Gamma (RBF kernel scale): 0.1

We split the dataset into **80% training** and **20% testing** sets, preserving class balance. K-fold cross-validation (k=5) ensured robustness.

4.4 Evaluation Metrics

To evaluate the classifier, we used:

- **Accuracy**: Correct classifications as a percentage of total samples.
- **Precision and Recall**: To balance detection with false alarms.
- **F1 Score**: Harmonic mean of precision and recall.
- **False Positive Rate (FPR)**: Legitimate DNS queries flagged as malicious.
- **Receiver Operating Characteristic (ROC) curve and AUC**.

We also benchmarked against a simple entropy threshold-based rule as a baseline.

5. Results

The SVM classifier performed well across all evaluation metrics, indicating that DNS tunneling can be reliably detected using structural and entropy features.

5.1 Classification Performance

Table 5.1 – SVM Classifier Performance

Metric	Score (%)
Accuracy	93.4
Precision	92.8
Recall (TPR)	94.1
F1 Score	93.4
False Positive Rate	4.6

Metric Score (%)

AUC-ROC 0.96

- The classifier exceeded the 90% target for both accuracy and precision.
- The **false positive rate of 4.6%** was acceptable for operational deployment, especially with layered alert confirmation via SIEM integration.


5.2 Feature Importance

The top contributing features, based on recursive feature elimination and model weights:

- Entropy of query name
- Subdomain depth
- Max label length
- Query frequency per source IP
- NXDOMAIN ratio

These features aligned with tunneling behavior: high entropy from encoded payloads, deep

Figure 1: DNS Tunneling Detection – SVM vs. Rule-Based



6. Discussion

The results of our experiments confirm that DNS tunneling can be effectively identified using a supervised machine learning approach based on well-engineered features. This section explores the operational implications, strengths, and limitations of our framework.

subdomain hierarchies, and repeated failed lookups.

5.3 Baseline Comparison

A rule-based model using an entropy threshold of 4.5 achieved only 78.2% accuracy with a high false positive rate (12.7%), showing the superiority of the ML approach.

5.4 Real-Time Testbed Results

We deployed the classifier in a passive monitoring mode on a small corporate lab network:

- Detection latency:** < 1 second per event
- Impact on DNS resolution:** None (passive inspection)
- SIEM alert integration:** Achieved using Splunk forwarders with custom alerting rules

The model successfully flagged all DNScat2 and Iodine sessions with no disruption to normal DNS activity.

6.1 Practical Viability in Enterprise Environments

Our SVM-based detection model demonstrated low latency and negligible resource overhead, making it feasible for integration with enterprise monitoring solutions. Unlike deep packet inspection, which can introduce



bottlenecks, our system processes only metadata-level features from DNS queries, preserving privacy while maintaining detection accuracy.

The framework is especially valuable in **zero-trust architectures** where traffic segmentation exists but covert channels are still possible. It complements traditional signature-based intrusion detection systems by offering behavioral analysis on a frequently ignored protocol.

6.2 Strengths of the Approach

- **High Accuracy with Low False Positives:** The model achieved 93.4% accuracy and only 4.6% FPR, showing a good balance between sensitivity and reliability.
- **Protocol-Agnostic Detection:** Rather than relying on domain reputation or static rules, the system identifies behavioral anomalies across unknown domains.
- **Deployability:** The model can operate in passive tap mode and be deployed on DNS logs, making it suitable for cloud and hybrid environments.
- **SIEM Integration:** Real-time alerts via Splunk or ELK stacks can help security analysts correlate DNS anomalies with endpoint or proxy logs.

6.3 Limitations and Challenges

- **Evasion through Encoding or Obfuscation:** Sophisticated attackers may employ base64 padding, dynamic tunneling intervals, or mimic legitimate subdomain patterns to reduce entropy.
- **Encrypted DNS:** With increasing adoption of DNS over HTTPS (DoH) and DNS over TLS (DoT), the model's visibility may be limited unless inspection is possible via SSL termination or endpoint telemetry.
- **Context-Free Classification:** The model does not account for contextual data such as user identity, device risk level, or time-of-day behavior—factors

which could further reduce false positives.

6.4 Recommendations for Deployment

- **Layered Deployment:** Combine this model with upstream DNS filtering (e.g., domain categorization) and SIEM correlation for robust detection.
- **Policy Enforcement:** Implement DNS query rate limiting, domain whitelisting, and block access to known tunneling tools on endpoints.
- **Visibility into Encrypted DNS:** Use local DNS proxies or endpoint monitoring agents to regain visibility where DoH is enabled.

7. Conclusion and Future Work

DNS tunneling represents a stealthy and increasingly common method of data exfiltration in corporate networks. Because DNS is typically considered safe and left uninspected, attackers exploit this oversight to bypass perimeter defenses. This paper presented a lightweight, machine learning-based detection framework that analyzes DNS query characteristics—such as entropy, label structure, and query frequency—to detect tunneling behavior.

The key findings include:

- An **SVM classifier** trained on structural DNS features can achieve **93.4% detection accuracy** with a **false positive rate under 5%**.
- Real-time deployment in a testbed network confirmed the framework's **practical viability** with minimal performance overhead.
- Top predictive indicators include query entropy, subdomain depth, and repeated NXDOMAIN responses.

This approach provides security teams with an **additional detection layer**, supplementing traditional domain blacklisting and enhancing visibility into otherwise benign-looking traffic.

Future work may explore:



- **Time-series anomaly detection** using recurrent neural networks (RNNs) to detect slowly evolving tunneling campaigns.
- **Context-aware classification**, integrating user, device, and geolocation data for adaptive risk scoring.
- **Encrypted DNS handling**, by correlating DNS client behavior at endpoints or leveraging DNS proxy architecture for inspection.
- **Broader traffic fusion**, combining DNS anomalies with web, email, and endpoint activity for comprehensive data leak detection.

As adversaries continue to seek new covert channels, DNS anomaly detection will remain a critical line of defense in enterprise cybersecurity.

References

1. Krishnan, R., Vatsalan, D., & Liu, X. (2013). Detection of covert channels in DNS traffic using entropy and traffic pattern analysis. *International Journal of Computer Applications*, 72(3), 1–6. <https://doi.org/10.5120/12438-9006>
2. Jena, J. (2015). Next-Gen Firewalls Enhancing: Protection against Modern Cyber Threats. *International Journal of Multidisciplinary and Scientific Emerging Research*, 3(4), 2015-2019. https://ijmserh.com/admin/pdf/2015/10/46_Next.pdf
3. Cova, M., Kruegel, C., & Vigna, G. (2007). Detection and analysis of drive-by-download attacks and malicious JavaScript code. *Proceedings of the 19th International Conference on World Wide Web*, 281–290. <https://doi.org/10.1145/1772690.1772726>
4. Munnangi, S. (2016). Adaptive case management (ACM) revolution. *NeuroQuantology*, 14(4), 844–850. <https://doi.org/10.48047/nq.2016.14.4.974>
5. Choi, H., Lee, H., & Kim, H. (2017). Botnet detection by monitoring group activities in DNS traffic. *Computer Networks*, 56(1), 20–33. <https://doi.org/10.1016/j.comnet.2011.08.016>
6. Mukherjee, P., & Sharma, V. (2016). A study on DNS tunneling detection techniques. *International Journal of Computer Applications*, 150(7), 1–5. <https://doi.org/10.5120/ijca.2016911406>
7. Storer, T., & Draper, G. (2017). Analysis of DNS tunneling tools and detection mechanisms. *Proceedings of the 13th International Conference on Security and Cryptography*, 128–135.
8. DNScat2. (2018). *Covert shell over DNS*. GitHub Repository. <https://github.com/iagox86/dnscat2>
9. Iodine. (2018). *Tunnel IP over DNS*. GitHub Repository. <https://github.com/yarrick/iodine>
10. Antonakakis, M., Perdisci, R., Dagon, D., Lee, W., & Feamster, N. (2010). Building a dynamic reputation system for DNS. *Proceedings of the USENIX Security Symposium*, 273–290.
11. Sinha, P., & Sharma, V. (2018). DNS as a tool for data exfiltration and detection using machine learning. *Journal of Information Security*, 9(4), 345–356. <https://doi.org/10.4236/jis.2018.94022>
12. Kolla, S. (2018). Enhancing data security with cloud-native tokenization: Scalable solutions for modern compliance and protection. *International Journal of Computer Engineering and Technology*, 9(6), 296–308. https://doi.org/10.34218/IJCET_09_06_031
13. Lee, S., & Kim, H. (2014). WarningBird: Detecting suspicious URLs in Twitter stream. *Computer Communications*, 54, 81–91. <https://doi.org/10.1016/j.comcom.2014.08.003>
14. Yu, S., Liu, X., & Ma, S. (2015). Anomaly detection for DNS tunneling using support vector machine. *IEEE International Conference on Computer and Information Technology*, 52–58. <https://doi.org/10.1109/CIT.2015.18>



15. RFC 1035 – Mockapetris, P. (1987). *Domain Names - Implementation and Specification*. IETF.
<https://datatracker.ietf.org/doc/html/rfc1035>
16. Sahoo, D., Liu, C., & Hoi, S. C. (2017). Malicious URL detection using machine learning: A survey. *arXiv preprint arXiv:1701.07179*.
<https://arxiv.org/abs/1701.07179>
17. Splunk Inc. (2018). *Integrating DNS threat intelligence with Splunk*. Retrieved from
<https://www.splunk.com>
18. National Institute of Standards and Technology (NIST). (2018). *DNS-Based Security Enhancements*. *NIST Special Publication 800-81-2*.
<https://doi.org/10.6028/NIST.SP.800-81r2>