infoblox

2025 DNS THREAT LANDSCAPE REPORT

The Cyberthreat Fog:

How Malicious Actors Use DNS to Deceive and Evade

Over the past year, threat actors have rapidly advanced their use of deception—scaling operations and leveraging AI to target individuals, organizations, and evade threat research. Infoblox Threat Intel has observed a new level of professionalism and speed in the way actors launch Domain Name System (DNS)-sourced cyberattacks, which affect consumers, businesses, and government agencies alike.

To defend effectively, security teams must understand the threats they face. Gaining insight into adversarial DNS techniques, the actors behind them, and the risks they pose is essential to strengthening defense strategies.

This report draws on vast volumes of real-time DNS telemetry, cutting-edge analytics, and decades of threat expertise to provide a unique perspective on how attackers exploit DNS. It also outlines the business implications and highlights DNS-based intelligence as a critical layer of modern cyberdefense.

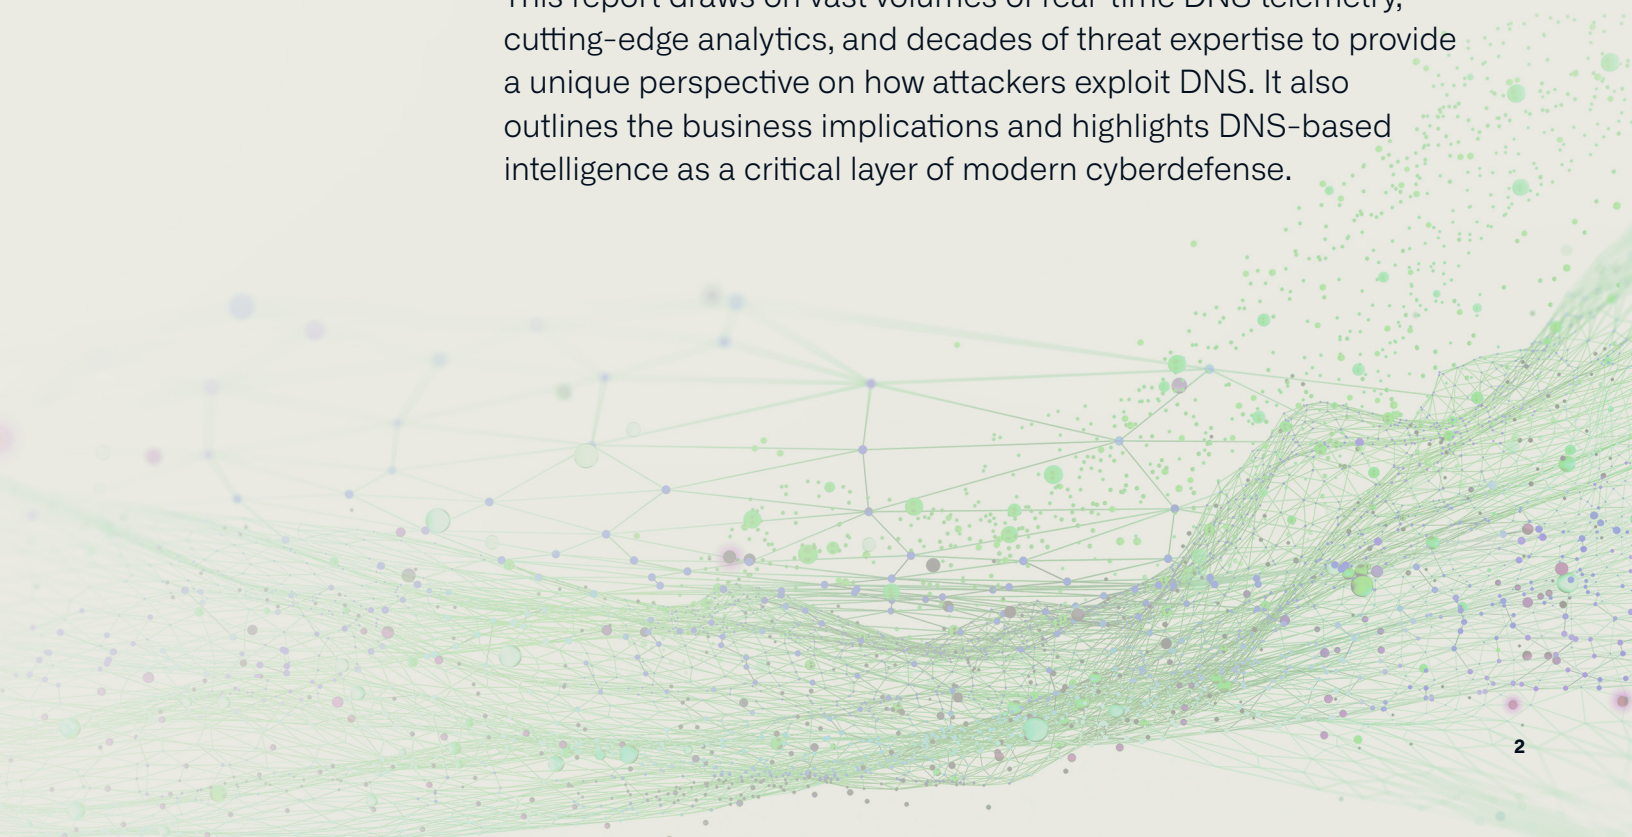

TABLE OF CONTENTS

THE UNTAPPED POTENTIAL OF DNS INTELLIGENCE

People often refer to DNS as the phonebook of the internet because it translates domain names into IP addresses. Every digital interaction begins with a DNS request, making it a high-fidelity source of telemetry for network operations by providing in-depth visibility into which digital assets are initiating connections over the internet.

DNS is also utilized by malicious actors when phishing, scamming, for detection evasion, and during data extraction. Consequently, analyzing DNS traffic and domain usage is foundational for security analysts. DNS data can be reshaped into predictive threat intelligence by holistically collecting pre-attack telemetry, enriching the data, analyzing it against baselines, and executing deep threat hunts. These insights offer defenders a comprehensive view of adversarial infrastructures, targeted victims, and tactics—before the attacker strikes.

"DNS unlocks a unique vantage point into past threat activity, which in turn serves as a crystal ball—revealing the precursors to future cyberthreats."

— **Dr. Renée Burton**
Head of Infoblox Threat Intel

As a result, DNS offers much more than just name resolution and has become both an enforcement point for enterprise security policy and an indicator of potential malicious activity on a network. Organizations like the National Institute of Standards and Technology (NIST) and the Cybersecurity & Infrastructure Security Agency (CISA) have recognized this critical—and early—role that DNS plays in cybersecurity and have highlighted its preemptive security potential in the recently proposed NIST Special Publication (SP) 800-81 Rev. 3.[1]

This report addresses four key questions:

What are the key DNS observations from the past 12 months?

Who are the DNS threat actors and what recent activities have been discovered?

What are the main malicious tactics behind DNS techniques and why are they dangerous?

What are the key challenges for defenders, and what opportunities does DNS-based threat intelligence offer?

1 Secure Domain Name System (DNS) Deployment Guide, National Institute of Standards and Technology (NIST), April 10, 2025.

SECTION 1: TOP DNS THREAT OBSERVATIONS

100.8
million newly observed domains in one year

25.1%
of newly observed domains are malicious or suspicious

Ephemeral Nature of Domains

As of the end of May 2025, Infoblox was processing and analyzing 70 billion DNS queries daily from over 13,000 Infoblox environments, covering millions of IP addresses across all types of devices.

The fully anonymized data from 1,300+ Infoblox Threat Defense™ customers provides global as well as in-depth visibility into millions of internet interactions, spanning multiple client types, geographies, and industry verticals. Year over year, this DNS telemetry volume increased by 21 percent.

Within all collected data, Infoblox Threat Intel identified **100.8 million newly observed domains (second-level domains) in the past 12 months**. This high volume of new domains is often a result of fast-changing infrastructures, short-term advertising campaigns, and branding initiatives.

Control Evasion Via One-Time-Use Domains

More than one-quarter of the newly observed domains (over 25 million) were classified by Infoblox as malicious or suspicious. Threat actors continuously register, activate, and deploy massive numbers of new domains to evade detection controls. Because it is difficult to identify and classify such large volumes of domains, attackers are able to fly under the radar, bypass blocking mechanisms, and leave minimal forensic evidence.

The isolated usage of identified threat-related domains—both malicious and suspicious—is also significant. Infoblox Threat Intel found that 95 percent of all threat-related domains were observed within a single network environment.

The objective behind this tactic is simple: bypass forensic-based defenses that rely on “patient zero” data by leveraging throwaway domains—of which attackers have unlimited supply.

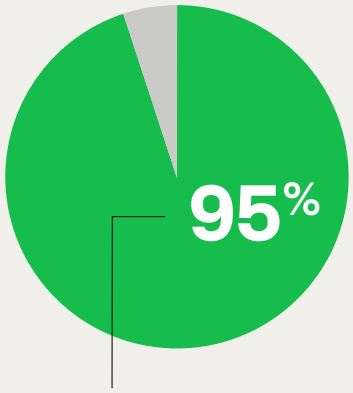

of threat-related domains were observed in only one customer environment.

Malicious Versus Suspicious Domains

- **Malicious domains** are confirmed threats supported by strong evidence. They do not age out and account for 1.6 percent of over 100 million newly observed domains.
- **Suspicious domains** are potential threats that lack conclusive evidence and account for 23.5 percent of all newly observed domains. If not confirmed, these indicators expire after a few months. Infoblox Threat Intel analysts continuously monitor these domains for new evidence. When additional indicators are discovered, the scores are updated, and suspicious domains may be reclassified as malicious.

Cloaking Via Domains Part of Traffic Distribution Systems

Adtech (short for advertising technology) refers to the tools, software, and platforms used to automate, manage, target, deliver, and analyze digital advertising. Traffic distribution systems (TDSs) are the platforms or mechanisms used—legitimately or maliciously—to redirect incoming internet traffic to different destinations based on predefined rules. Threat actors also adopted this technology, often referred to as **malicious adtech**.

82%

of customers queried a domain part of a traffic distribution system.

Over the past 12 months, **82 percent of all customer environments** queried domains that were part of TDS, much of which are operated by malicious adtech operators known for concealing harmful content, such as tailored phishing sites, scareware, scams, and infostealers.

These TDSs often consist of tens of thousands of domains, which are rapidly rotated to evade detection, delivering targeted malicious content to the ideal victims while cloaking that content from threat researchers.

Over time, Infoblox Threat Intel discovered over **1 million domains used by 168 malicious adtech operators** within their TDS infrastructure. These indicators span multiple DNS techniques, like hijacked domains, lookalikes, redirections, and algorithmically preregistered domain sets (registered domain name algorithms, or RDGAs). More on TDSs, how they work, and why they are dangerous in Section 3.

Domains Linked to Diverse Threat Types

As new threat-related domains are discovered, Infoblox threat researchers investigate the actors behind them and their underlying intent. The table on the next page presents a prioritized list of how actors use their domains for various malicious purposes.

Top 7 List: How Threat Actors Utilize New Domains	
1	**Engage in fraudulent activities and scams**, such as fake cryptocurrency investment sites.
2	**Host illegal content**, including gambling (particularly in regions like China) and adult material.
3	**Create phishing pages** designed to steal personal information or credit card data.
4	**Deploy malware.** Common examples include infostealers (e.g., Lumma Stealer), loaders via drive-by downloads (e.g., SocGholish), botnets, and ransomware (e.g., BlackBasta).
5	**Cloak their activities** via TDS and deliver various payloads or trick users into allowing unwanted browser notifications.
6	**Distribute potentially unwanted programs** (**PUPs**), such as scareware or unnecessary browser extensions.
7	**Conduct spam campaigns and distribute malicious emails**.

Table 1. Actors' purpose for newly observed domains.

Domain Popularity

Infoblox DNS telemetry also provides insights into domain type usage, offering clues about application popularity and the speed at which threat actors are becoming proficient at successfully pushing large volumes of weaponized domains in front of victims.

19 DAYS

Time needed for a TDS domain to become popular

Key Observations:

- Eight domain categories—such as content delivery networks (CDNs), technology providers, security vendors, business productivity tools, search engines, storage, cloud services, and net conferencing—account for the majority (approximately 70 percent on a given day) of all domains within customer DNS queries.
- In May 2025, domain queries related to personal internet usage—such as online shopping, gaming, and social media (e.g., TikTok and Facebook)—reached parity with those associated with professional collaboration platforms (e.g., Microsoft Teams, Slack). This illustrates the **growing overlap between professional and personal internet** use—an overlap that threat actors are acutely aware of.
- Adversaries continuously seek out weak attack surfaces—such as bring-your-own-device (BYOD) and mobile devices—and deceive users into performing high-risk actions aimed at extracting business-related data, including credentials. This trend was also highlighted in Verizon's 2025 Data Breach Investigations Report,[2] which affirms that no device is off limits and notes that **46 percent of stolen corporate credentials** originated from unmanaged or personal devices.

2 2025 Data Breach Investigations Report, Verizon.

- Infoblox Threat Intel observed domains part of TDSs becoming popular[3] in as few as 19 days, **2.35 times faster than in 2024 and 39 times faster than in 2020.** The speed at which TDS domains gain popularity—comparable to legitimate sites like `panerabread[.]com` or `draftkings[.]com`—illustrates how effectively weaponized domains are propagated and accessed by victims. Threat actors rapidly deploy large volumes of these domains in front of their targets, maximizing the impact of their campaigns while outpacing slower intelligence sources, such as open-source intelligence (OSINT) and forensic-based analysis.

SECTION 2: THREAT ACTORS AND RESEARCH

204K
total identified suspicious domain clusters

662
total identified DNS threat actors

10
new actors publicly disclosed in the past 12 months

The 100 million new domains discovered over the past year are not forces of nature—they are always caused by human actions and initiated for specific purposes. Infoblox Threat Intel continuously analyzes and investigates the actors behind threat-related domains by enriching collected telemetry and correlating common patterns

Since the start of its research, Infoblox Threat Intel discovered a total of 204,000 suspicious domain clusters, each sharing common threat elements, and has identified 662 unique threat actors. In the past 12 months alone, Infoblox researchers have publicly disclosed 10 new actors through various research reports and blog posts.

3 A domain is considered popular when it belongs to the subset of domains that account for the majority of customer traffic during a specific timeframe. It can range anywhere between 6,000 to 10,000 domains on a given day. For more information, see https://blogs.infoblox.com/wp-content/uploads/infoblox-whitelists-that-work.pdf.

The following list highlights key threat actors identified and publicly disclosed by Infoblox Threat Intel between July 1, 2024, and July 1, 2025.

Actor	Description
	This actor operates a malicious TDS that hijacks legitimate web traffic—primarily from compromised WordPress sites—and redirects it to scams, malware, and phishing content. VexTrio is considered one of the most pervasive and evasive actors in the threat landscape. Over the past 12 months, the actor was named in several reports for their relationship with affiliate hackers and is known for hijacking domains to supply their attack infrastructure. **Recently published reports:** • The Vexing and Vicious: The Eerie Relationship between WordPress Hackers and an Adtech Cabal • Pushed Down the Rabbit Hole
	This sophisticated DNS threat actor group specializes in hijacking abandoned cloud resources—such as Amazon S3 buckets and Azure endpoints—by exploiting misconfigured or forgotten DNS records, particularly dangling Canonical Name (CNAME) entries. Once Hazy Hawk gains control over these subdomains, it leverages the inherent trust of legitimate domains to host malicious content. Its operations often involve redirecting users through TDSs to deliver scams, malware, and deceptive push notifications. **Recently published reports:** • Cloudy with a Chance of Hijacking Forgotten DNS Records Enable Scam Actor
	This financially motivated threat actor has used hijacked domains for investment scams since February 2023. They embed these domains in short-lived Facebook ads across multiple continents, targeting victims in over 30 languages, including English, Italian, Polish, Turkish, and Spanish. The actor employs the Sitting Ducks attack vector to hijack reputable domains, which they use to protect their fraudulent sites from security researchers. As of October 2024, Infoblox has identified nearly 5,000 hijacked domains tied to this actor. **Recently published reports:** • Uncovering Actor TTP Patterns and the Role of DNS in Investment Scams • DNS Predators Hijack Domains to Supply Their Attack Infrastructure

	Reckless Rabbit is an investment scam actor that lures victims through malicious Facebook ads. It employs dictionary-based RDGAs and targets individuals in multiple countries, including Austria, Belgium, Denmark, France, Poland, Sweden, the United Kingdom, and others. The actor uses RDGAs and fake endorsements. **Recently published reports:** • Uncovering Actor TTP Patterns and the Role of DNS in Investment Scams
	This phishing actor runs investment scam campaigns that leverage dictionary-based RDGAs and spoof popular services. The actor operates their own domain cloaking service to perform user validation checks and targets Eastern European countries such as Romania, Russia, Poland, and others. **Recently published reports:** • Uncovering Actor TTP Patterns and the Role of DNS in Investment Scams
	This actor identifies abandoned cloud resources and repurposes them for various malicious activities. Hasty Hawk is known for hijacking domains that are used in charity-themed and DHL-themed campaigns distributed via Google ads. Hasty Hawk primarily uses “bulletproof” hosting networks such as Proton66 along with a TDS to direct users to the content. **Recently published reports:** • DNS Predators Hijack Domains to Supply Their Attack Infrastructure
	Vacant Viper operates the 404TDS, using it to deliver malware and other malicious content. Vacant Viper hijacks domains that are left vulnerable due to misconfigured DNS name servers—a flaw named “Sitting Ducks” by Infoblox researchers—and incorporates them into its malicious TDS infrastructure. **Recently published reports:** • Who Knew? Domain Hijacking Is So Easy
	This malicious adtech actor leverages WordPress vulnerabilities and distributes malware, phishing pages, fake apps, and unwanted content. They run an extensive TDS that incorporates push notifications, pop-ups, and redirects within a browser, serving ads even after the user leaves the initial page. **Recently published reports:** • The Vexing and Vicious: The Eerie Relationship between WordPress Hackers and an Adtech Cabal

Morphing Meerkat is global spam actor behind an advanced phishing-as-a-service (PhaaS) platform. This actor uses DNS MX records to identify the victim's email service provider and dynamically serve fake login pages. Morphing Meerkat exploits compromised WordPress websites as well as open redirect vulnerabilities on adtech servers.

Recently published reports:

- A Phishing Tale of DOH and DNS MX Abuse

ACTOR CASE STUDY: COORDINATION BETWEEN WORDPRESS HACKERS AND VEXTRIO VIPER CABAL

Infoblox recently uncovered a complex alliance between **WordPress hackers and a network of malicious adtech companies**, notably VexTrio's TDS.

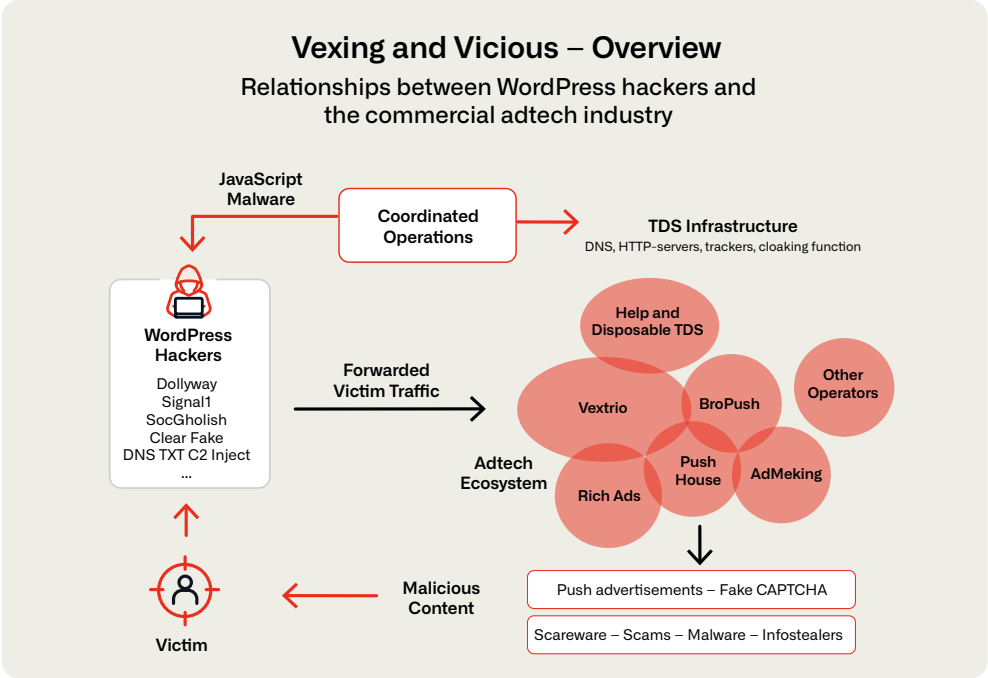

Figure 1. Relationship between WordPress hackers and commercial adtech industry

What Happened?

- **Quick Migration:** When VexTrio's TDS was disrupted in fall of 2024, multiple malware actors simultaneously shifted to a seemingly new TDS named "Help TDS." Further analysis revealed that Help TDS is not independent but closely linked to VexTrio, sharing infrastructure and software components.
- **Coordinated Operation:** Infoblox analyzed 4.5 million DNS TXT record responses from compromised websites over six months. This revealed two distinct command-and-control (C2) servers, both hosted on Russian-connected infrastructure, indicating a coordinated operation between WordPress hackers and the VexTrio cabal.
- **Involvement of Commercial Adtech Firms:** Several adtech companies, including

Los Pollos, Partners House, BroPush, and RichAds, were found to be intertwined with VexTrio's operations. These firms facilitated the distribution of malicious content via smartlinks and push notifications.

The investigation highlights the sophisticated and adaptive nature of cybercriminal networks leveraging compromised WordPress sites and commercial adtech infrastructures. It underscores the importance of DNS telemetry and collaborative efforts in uncovering and mitigating such threats.

SECTION 3: MALICIOUS DNS TECHNIQUES

Threat actors mentioned in Section 2 use DNS in various ways and with specific objectives in mind. Once Infoblox discovers a threat-related domain, analytics processes and expert reviews assign known malicious techniques to the domain. The table below provides an overview of the most common DNS techniques assigned by Infoblox Threat Intel to threat-related domains.

DNS Techniques and Threat-Related Domains Time frame: January 2025 through June 2025	
Domains generated by machines algorithms (RDGA, DDGA and DGA)	54.7 %
Domains used to redirect traffic	11%
CNAME or alias domains	5.8%
Lookalikes	5.1%
Hijacked domains	5.1%
Domains used in malicious SMS	4.2%
Domains created as part of a TDS	1.8%
Domains used for C2 and exfiltration	< 0.4%

Table 2. DNS techniques assigned to threat-related domains

Many of these techniques overlap during a threat campaign and become part of larger actor tactics to achieve their objectives. In this report, we dive deeper into four common DNS techniques, how they are used, and why they are dangerous:

- Usage of domains within TDSs
- Hijacking domains to steal trust
- Lookalike domains to deceive victims
- DNS tunneling for C2 and exfiltration

TRAFFIC DISTRIBUTION SYSTEMS PROVIDE A DANGEROUS LEVEL OF EVASION

DNS plays a central role in TDS by covertly redirecting users through multiple intermediary layers—often without their knowledge—based on various attributes like geolocation, device type, or security posture. DNS plays a foundational role in determining how and where network traffic is routed. Legit operators of TDS are mostly found in digital advertising or adtech. The name adtech (short for advertising technology) refers to the tools, platforms, and software used to manage, deliver, and analyze digital advertising campaigns.

Just like known legal advertising technology (e.g., Google AdSense), malicious adtech delivers the right content to the right audience at the right moment to increase the effectiveness of their campaigns. This type of cyberthreat is carried out by specialized organizations with many affiliates and deep pockets.

Top TDS Operators by Connection Share	
Actor Name	**Connection Share**
VexTrio Viper	72.8%
Vane Viper	68.4%
Venal Viper	72.5%
Undisclosed actor	64.8%
Vero Viper	60.5%
Tiano Gambling	50.9%

Table 3: TDS operators and the percentage of customer connection attempts they received

At the heart of these activities is a TDS that profiles victims and routes them to malicious advertisers while pointing threat researchers to a decoy site.

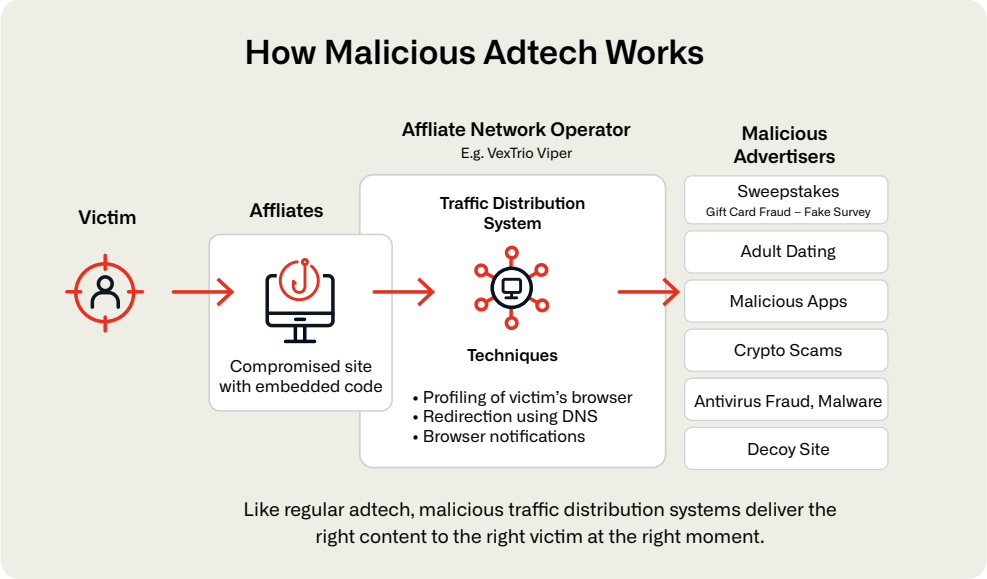

Figure 2: A high-level picture of the three players in malicious adtech; affiliates, operators, and malicious advertisers

There are multiple reasons why malicious adtech is harmful and should be an important area of focus for enterprise security teams:

Malicious Adtech Is a Fast-Growing, Underreported Threat Vector

Attackers are increasingly leveraging malicious ad networks as a low-cost service to deliver malware and other malicious content. These ads can lead to various types of attacks, including drive-by downloads, phishing sites, credential stealers, and exploit kits (see Table 4: TDS operators and delivered malicious content).

Because most of the security industry relies on a patient-zero approach—collecting telemetry during (e.g., sandboxing) or after (e.g., forensic-based intelligence) an attack—the resulting countermeasures are limited to artifacts discovered from that initial point of compromise. This limitation makes TDSs effective tools for evading detection, as actors continuously alter the malicious content they deliver and redirect threat researchers to decoy sites. Consequently, TDSs have become one of the most underreported threats in the cybersecurity industry.

Large Scale Infrastructures, Hard to Disrupt

Organizations operating malicious adtech often build infrastructure at considerable scale, including tens of thousands of fast-changing domains designed to redirect users and lure them into accepting browser push notifications. These operations are frequently compartmentalized into multiple entities to carry out cybercrime while avoiding legal scrutiny. Some operators, such as VexTrio Viper, have persisted for years, becoming highly profitable—and their activities show no signs of stopping.

Malicious Adtech Serves as a Gateway to Enterprise Risk

Malicious adtech deceives victims by mimicking popular brands or offering content they are eager to access, encouraging them to drop their guard and engage in high-risk interactions. Although these threats typically originate on consumer-facing sites, they can easily infiltrate corporate environments—exposing employees' personal devices to weaponized content. This allows threat actors to perform reconnaissance or impersonate enterprise notifications, elevating the risk to organizational networks.

DNS Operators	Malware	Scams	Phishing	Hijacked Domain
Vacant Viper	X	X		X
Vane Viper	X	X	X	
Vextrio Viper	X	X	X	X
Hasty Hawk			X	X
Sophisticated Chickens			X	X
Black TDS	X		X	
Parrot TDS	X			
R0bl0ch0n TDS		X		

Table 4. TDS operators and delivered malicious content

Example of TDS at Work:

When a victim visits a compromised site from a mobile device or endpoint, the operator may present a fake CAPTCHA to trick the victim into accepting browser push notifications from a malicious advertiser. These notifications can then deliver additional fraudulent content, such as prompts to download unverified software, share personal information, or enter organizational credentials.

As the TDS profiles incoming victims, SOC analysts or threat researchers using common security tools may not detect these notifications or malicious content—they may instead be redirected to a decoy site displaying legitimate material.

Because of the overlap between professional and personal internet usage, malicious advertising technology has become a significant contributor to cybercrime, especially to mobile devices, tablets, BYOD and unprotected assets.

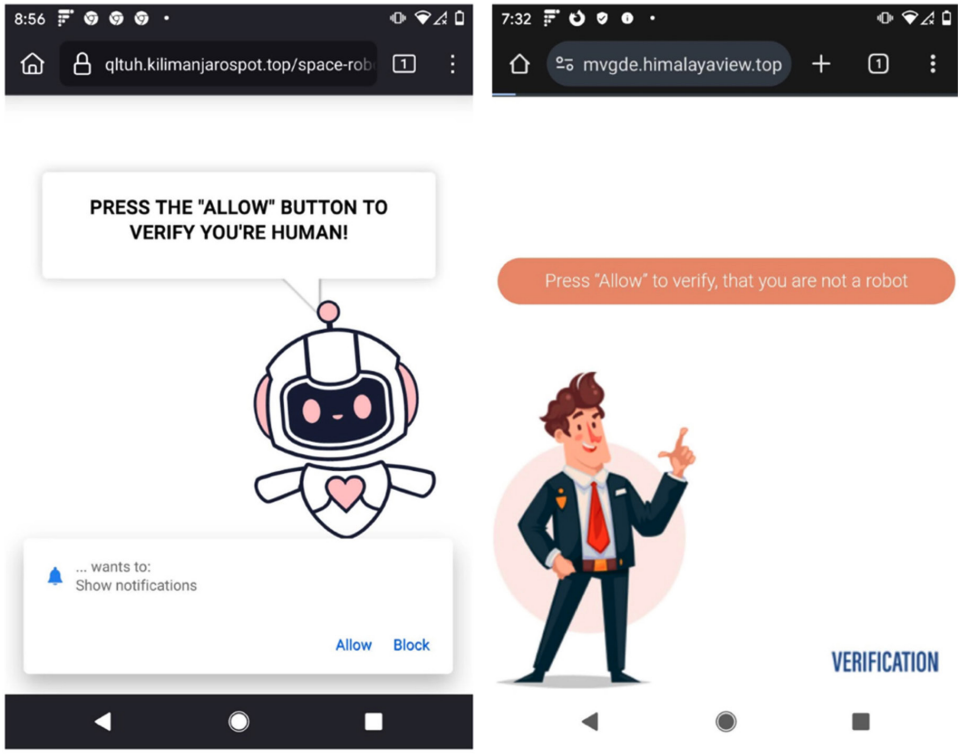

Figure 3: Examples of VexTrio Viper's landing page that leads the user to accept push notifications on their device; these were both seen when browsing to `germannautica[.]com`

Domains Used by Traffic Distribution Systems

Over the past 12 months, Infoblox Threat Intel discovered over 1 million indicators used by 168 malicious adtech operators within their TDS. These indicators span multiple techniques, including RDGAs, redirections, hijacked domains, lookalikes, and others.

TDSs used by malicious adtech operators can be quite large. Many include over 10,000 domains, with some exceeding 100,000. However, the size of a TDS does not necessarily correlate with its pervasiveness or threat level. Vigorish Viper operates a vast and growing network of 170,000 active domains but primarily targets victims in China, Hong Kong, and Macau. Venal Viper, though not among the top five in size, is one of the most frequently queried in customer networks—65 percent of all Infoblox customers have queried a Venal Viper domain in the past 12 months.

Disrupting TDS

Harmful adtech using TDSs thrives because it masquerades as legitimate advertising, deceives victims, and evades detection by security tools that rely on identifying known malicious behaviors through simulations or patient-zero data. In contrast, DNS records can reveal when and how new malicious infrastructure is configured.

Researchers who leverage real-time and historical DNS data—combined with innovative data science—can identify suspicious or malicious domains before any payload is delivered, including those used in malicious adtech.

DNS-derived intelligence sheds light on the infrastructure behind the threat, such as how the TDS operates and redirects traffic. Unlike other security methodologies, DNS-based security implementations can proactively uncover malicious adtech and prevent internet-connected endpoints from interacting with it.

Put simply, by focusing on attacker infrastructure, DNS-based protection breaks the supply chain between malicious advertisers and victims—offering long-term protection rather than merely reacting to the latest payloads.

DOMAIN HIJACKING TO STEAL TRUST

Threat actors hijack existing domains primarily to exploit the credibility and trust associated with legitimate domains. Once under control of the adversary, hijacked domains can be used to create convincing phishing sites, get prioritized by search engines, bypass spam filters, or execute fraud.

Infoblox Threat Intel discovered multiple ways actors hijack domains and the tools they use to deceive users.

Sitting Ducks Attacks

Sitting Ducks attacks gained prevalence over the past years. In 2024, Infoblox Threat Intel estimated that more than **1 million domains are vulnerable to this attack**. During a deep research exercise in the second half of 2024, **70,000 domains were discovered hijacked out of a pool of 800,000** vulnerable domains. This highlights the scale of the problem and the need for robust security measures.

Multiple threat actors use these techniques systematically. The ease with which these attacks can be executed—combined with the difficulty security teams face in detecting them—makes them particularly dangerous.

Actors known to be exploiting this attack include VexTrio Viper, Vigorish Viper, Horrid Hawk, and Hasty Hawk. These groups have demonstrated the effectiveness of Sitting Ducks attacks, underscoring the need for heightened vigilance and improved security practices to counter these threats.

Dangling CNAMES

In early 2025, threat actors exploited redirection configurations on high-reputation domains such as `cdc[.]gov` and several U.S.-based universities. This was possible because organizations had decommissioned cloud applications (e.g., CDNs) hosted by third-party providers (such as Microsoft Azure) while leaving their DNS aliases (CNAME records) active.

Malicious actors like Hazy Hawk exploited this lapse in DNS hygiene by creating new content on the same CDN. The motive was simple: by leveraging the reputation of the original domain alias, they were able to trick Google and other search engines into indexing the malicious content and including it in search results.

LOOKALIKE AND TYPOSQUATTED DOMAINS DECEIVE USERS

Lookalike domains are slightly altered domain names registered to deceive users. They often impersonate legitimate brands, employee communications, supply chains, or other trusted partners, causing significant issues.

Attackers used lookalike domains in SMS messages, phone calls, direct messages on social media, emails, and QR codes. Recently, they targeted multi-factor authentication (MFA) due to its growing adoption by everyone from gamers to digital currency marketplaces. Other examples include bypassing enterprise MFA or abusing domain names from popular identity access platforms.

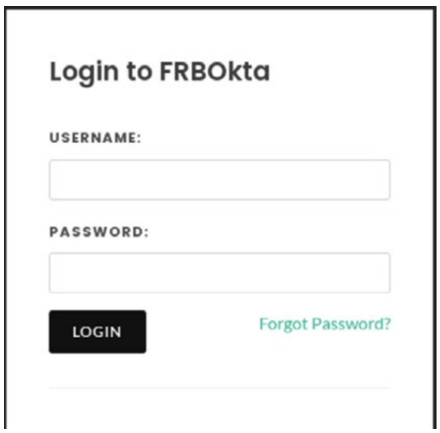

Figure 4. MFA message from lookalike domain

Lookalike domains have become a significantly larger problem as there are over **1,500 top-level domains, increasing costs for most organizations to monitor all variations.**

Additionally, organizations may have multiple groups registering domains and lack visibility into who is doing what. Security teams may think a similar-looking domain was created by the helpdesk or cloud application team, but the new domain may actually be configured by an actor to phish customers.

A lack of expertise within security teams often creates an appetite for quick fixes through managed services. Unfortunately, lookalike domains are not an easy problem to solve. Even mature security teams continue to encounter them, and effective monitoring requires significant diligence.

Infoblox: Identified Lookalike Techniques	
	**Homographs or homoglyphs** use visually similar characters from different character sets, such as Cyrillic or Greek (e.g., substituting "o" with "0"). The technique is effective because the inserted characters are not always clearly distinguishable.
	**Typosquats** include sneaky typing errors by registering domains that closely resemble popular websites (e.g., substituting "`amazonn[.]com`" for "`amazon[.]com`") to take users to a fraudulent website.
	**Combosquats** combine well-known brand or company names with other keywords, such as "mail", "security", or "support". Combosquatting is around 100 times more prevalent than typosquatting.
	**Soundsquats** are the most recent form of lookalike threats, using domain names that sound similar when spoken aloud (e.g., "`hsbsee[.]com`" instead of "`hsbc[.]com`"). It deceives users when using smart devices, such as Google Home, Siri, and Alexa.

DNS TUNNELING USED BY THREAT ACTORS, PENTESTERS AND LEGIT SECURITY TOOLS

DNS tunneling encodes data within DNS queries and responses, enabling covert communication that is often exploited for C2 operations and data exfiltration.

While Infoblox observed over 480 unique DNS tunneling domains in some months, an average of more than 100 unique domains related to DNS tunneling were discovered per month between June 2024 and June 2025. In addition to cybercriminal use, DNS tunneling is also employed in legitimate penetration testing and security tools. The following list provides an overview of prevalent DNS tunneling tools with C2 capabilities.

+100

unique DNS tunneling domains found monthly—benign and malicious

- **Cobalt Strike** is a widely used pentest tool featuring a DNS C2 module. Utilized by red teams and threat actors, it employs Hex-encoded queries with optional customizable prefixes like “post”, “api”, or “dx”.
- **Dnscat2** is a tool used for creating encrypted DNS tunnels. It is included within METASPLOIT, an open-source penetration testing tool.
- **DNS Exfiltrator** is a tool that encodes data into DNS queries for exfiltration, illustrating the potential misuse of DNS in practical scenarios. It uses TXT records, allows only one-directional communication, and is initiated via the command-line. Infoblox has not observed its use by a threat actor and considers it impractical due to the one-directional mechanism.
- **Sliver** is a cross-platform C2 framework with DNS tunneling capabilities, frequently utilized in adversary simulations and malicious campaigns.
- **Weasel** is a less-documented DNS tunneling tool developed by Facebook’s Red Team that supports stealthy data exfiltration and C2, typically used in niche red teaming engagements. It uses A and AAAA records for communications.
- **Pupy** is an open-source, multi-platform Remote Access Tool with DNS tunneling support, historically leveraged in espionage campaigns against government and corporate entities. It uses A records for communications.
- **Iodine** is a well-known tool for tunneling IPv4 traffic over DNS, used in penetration tests and sometimes abused in attacks, such as by nation-state actors for C2 purposes. Iodine uses A, TXT, CNAME, and MX records to communicate.
- **Several automated penetration testing tools** from vendors such as Cymulate and AttackIQ have emerged recently. Infoblox has discovered domains related to these vendors within customer networks.
- **Antivirus and antispam tools** also use DNS as a mechanism to look up if a domain or file hash may be malicious. A query may be of the form: “<domain>.<guid>.<avdomain”or “<file hash>.<guid>.<avdomain>” with response being NXDOMAIN if the domain or file hash is not in a known malware or spam list, or 127.0.0.X if it is in such a list.

Security Teams Need a Scalpel to Stop DNS Tunneling

Understanding and mitigating DNS tunneling is essential for protecting enterprises from cyberthreats and ensuring compliance with regulatory requirements, such as the Payment Card Industry Data Security Standard (PCI DSS), the Health Insurance Portability and Accountability Act (HIPAA), and the General Data Protection Regulation (GDPR). Due to the widespread use of DNS tunneling tools, many security teams struggle to effectively monitor and control DNS traffic.

Infoblox often detects DNS tunneling in networks, even those with next-generation firewalls or Secure Access Service Edge (SASE) type technologies. While these technologies have improved in detecting DNS tunneling, several complexities remain. CDNs, the use of new lookalike domains, and the expansion of legitimate DNS C2 tools complicate the detection and blocking of all C2 activities.

As a result, security teams require precise, targeted tools rather than broad, generalized measures. To address this challenge, Protective DNS solutions that leverage active threat actor tracking and continuously updated machine learning techniques are essential.

SECTION 4: CHALLENGES FOR DEFENDERS

In addition to the traditional adversarial DNS techniques, like TDSs, domain hijacking, lookalike domains, and DNS tunneling, defenders—whether they are SOC analysts, risk managers, or CISOs—face a growing array of challenges.

This section provides an overview of key trends like the use of adversarial AI, brand protection, and increasing pressure from new compliance mandates. Most importantly, it highlights opportunities offered by DNS-derived threat intelligence to combat these challenges.

88%

of AI-generated malware evades detections[4]

ADVERSARIAL AI BYPASSES EXISTING SECURITY CONTROLS

Generative AI (GenAI)—particularly large language models (LLMs)—is driving a transformation in cybersecurity. Adversaries are increasingly drawn to GenAI because it lowers the barrier to creating deceptive and convincing content. They use it to enhance the effectiveness of intrusion techniques such as social engineering and detection evasion.

To compensate for these new AI challenges, security teams need a new level of truth—such as DNS-based telemetry—that cannot be altered or obfuscated by AI and provides sufficient transparency in the chain of custody.

Recent Examples of Malicious AI: Deepfake Scams

At the end of 2024, the FBI warned that criminals were using generative AI to commit fraud at scale, making their schemes more believable.[5] GenAI tools like voice cloning significantly reduce the time and effort needed to deceive targets with seemingly trustworthy audio messages. Particularly concerning is the ease with which cybercriminals can access these tools, combined with the lack of security safeguards. Voice cloning has been used in various scenarios, including large-scale deepfake videos for cryptocurrency scams and the imitation of voices during targeted phone calls.

Case Study: Reckless Rabbit Usage of Deepfakes to Target Japanese-Speaking Victims

Infoblox Threat Intel reported in September 2024 on a YouTube account-hijacking campaign using deepfake videos of Elon Musk for crypto scams. A similar technique has now been adopted by a tracked actor known as **Reckless Rabbit**, who directly embeds deepfakes into fraudulent websites.

Reckless Rabbit recently shifted focus to **Japanese-speaking users**, promoting fake investment schemes through AI-generated news articles. These sites feature deepfake videos of public figures, like **Elon Musk** and **Masayoshi Son,** along with fabricated positive reviews to boost credibility.

4 Criminals Use Generative Artificial Intelligence to Facilitate Financial Fraud. FBI Alert Number: I-120324-PSA, December 3, 2024

5 AI Could Geneate 10,000 Malware Variants, Evading Detection in 88% of Case, Lakshmanan, Ravie, The Hacker News, December 23, 2024.

Previously, the actor targeted **Eastern European users** using RDGA-based domains and Facebook ads to lure victims to fake news content made up of simple text and images.

Figure 5. Recently discovered deepfake page

Reckless Rabbit uses fake articles featuring **deepfake videos with Japanese captions**, impersonating major outlets like Yomiuri Shimbun. These articles promote a fake investment platform called **"Finance Legend"** with a registration button that redirects to a contact form. The actor likely follows up with victims to solicit deposits by promising high returns.

AI-Powered Chatbots

Actors often select victims carefully by gathering intelligence about their interests, setting them up for highly personalized scams. After initial reconnaissance, they craft smishing messages that lead victims into chatbot-driven conversations. These conversations can continue for weeks and may include unusual steps, such as asking for a thumbs-up on YouTube or a repost on social media—tactics designed to assess the victim's susceptibility. With each positive interaction, the actor manipulates a fake "account balance" to rise. When the victim attempts to cash out, the actor requests access to their cryptocurrency account—abusing the trust built over time to steal the victim's funds. AI-powered chatbots enable actors to automate these conversations and scale their operations efficiently.

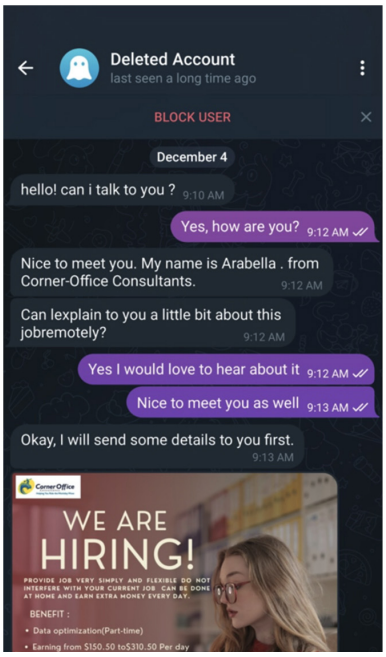

Figure 6. Example adversarial chat messaging, using a mix of AI/LLMs along with semi-automated chatbot interactions

Code Obfuscation and Evasion

Threat actors increasingly use GenAI to obfuscate, repurpose, and assemble malware in new ways to evade detection. This approach accelerates the creation of threat campaigns and reduces the technical skills required to build effective infection chains. According to HP Wolf Security research, email-based threat evasion has increased by approximately 11 percent.[6] Meanwhile, a prominent security vendor recently reported that a greedy LLM algorithm flipped its own malware classifier model's verdict from malicious to benign in **88 percent** of cases[7]—a significant indicator of how effectively adversarial AI can exploit current detection models..

PROTECTING BRAND AND ORGANIZATIONAL REPUTATION

Brands and organizational reputation are strategic assets. A strong reputation builds customer trust, enhances market credibility, attracts partners and investors, and supports long-term brand equity. According to Forbes, "Reputation is consistently ranked by corporate leaders as their most valuable asset."[8] However, protecting a brand within DNS presents several challenges:

- **Limited Visibility Beyond the Perimeter:** Monitoring domains requires tracking not only one's own domains but also thousands of potential lookalikes or impersonations. For example, Infoblox detected 28,331 lookalike domains in May 2025.
- **Human-Crafted Lookalikes Remain Hard to Detect:** Lookalike domains are carefully selected and imitated by humans, often surpassing the detection capabilities of automated systems.
- **Manual Domain Monitoring Strains Resources:** Security teams often lack the resources to manually monitor alerts and respond effectively. Without automation, domain monitoring becomes a high-effort, low-efficiency task.
- **Jurisdictional Barriers Hinder Enforcement:** 87 percent of discovered high-risk domains are registered with entities sanctioned by the Office of Foreign Assets Control (OFAC), where U.S. or European Union (EU) laws do not apply. As a result, domain and website takedowns are often ineffective.

28,331
lookalike domains detected by Infoblox in May 2025

To overcome these obstacles, security and marketing teams must partner with DNS experts who have deep visibility into global DNS usage and can leverage DNS-based intelligence. This collaboration enables them to monitor, detect, and remediate threats to digital assets that reflect the organization's reputation or brand.

COMPLIANCE PRESSURES AND DNS CHALLENGES FOR SECURITY TEAMS

Network and security teams face increasing pressure from evolving best practices and new mandates, such as **EU NIS2** and **NIST SP 800-81 Rev. 3**, which apply across sectors and require broader oversight—including DNS infrastructure.

6 Hackers Use Image-Based Malware and GenAI to Evade Email Security, Coker, James, Infosecurity Magazine, January 16, 2024.

7 AI Could Generate 10,000 Malware Variants, Evading Detection in 88% of Case, Lakshmanan, Ravie, The Hacker News, December 23, 2024.

8 The Importance Of Brand Reputation: 20 Years To Build, Five Minutes To Ruin, Blanchard, Paul, Forbes, December 27, 2019.

These frameworks introduce several challenges:

- **Operational Complexity:** NIS2 mandates risk assessments, 24-hour incident reporting, and continuous monitoring—requirements that are difficult for teams lacking centralized visibility or automation. NIST SP 800-81 Rev. 3 further requires deploying dedicated DNS servers and encrypting internal and external DNS traffic.
- **Fragmented Tooling:** Existing tools are often fragmented across on-premises, cloud, and remote environments, creating policy mismatches and visibility gaps. DNS policies (e.g., response policy zones, or RPZs) must be consistently applied to avoid disruptions.
- **Limited Resources:** SOC teams are overwhelmed by alert volume and lack contextual insight. NIS2's emphasis on early detection and rapid response puts additional strain on already overstretched teams-especially those lacking DNS-layer visibility.
- **Budget Constraints:** Compliance requires investment in tools, training, and DNS logging. Yet, organizations must justify these costs amid tighter budgets, even as DNS logging is critical for forensics and incident response.

Security teams require a straightforward approach to meeting new compliance requirements. Activating predictive threat intelligence and implementing controls at DNS level not only simplify compliance with NIST SP 800-81 Rev.3 and NIS2 but also aligns with broader security frameworks such as the NIST Cybersecurity Framework (CSF) and Zero Trust. Most importantly, it enhances global threat prevention, visibility, and reduction in security operations efforts.

NEXT STEPS

Infoblox offers security practitioners multiple options to explore our expert-produced threat intelligence and protect their environment with predictive intelligence.

For Threat Researchers:

- Learn more about Infoblox Threat Intel research at https://www.infoblox.com/threat-intel/.
- Talk to us on Mastodon at infobloxthreatintel@infosec.exchange.
- Access our research and indicators on GitHub at https://github.com/infobloxopen/threat-intelligence/.

For Security Teams:

- Request a DNS Security Workshop at https://info.infoblox.com/sec-ensecurityworkshop-20240901-registration.html.
- Learn more about Infoblox Threat Defense at https://www.infoblox.com/products/threat-defense/.

TERMINOLOGY USED

Adtech: short for **advertising technology,** refers to the **software, tools, and platforms** used by brands, agencies, publishers, and platforms to plan, execute, manage, and analyze **digital advertising campaigns.** It is the backbone of the online advertising ecosystem.

BYOD: bring your own device

C2: command and control

CDN: A content delivery network is a **network of geographically distributed servers** that work together to deliver digital content (like websites, videos, images, and scripts) **quickly, reliably, and securely** to users based on their location.

CNAME: Canonical Name record is a type of **DNS** (**Domain Name System**) record that **maps one domain name** (**an alias**) **to another domain name** (**the canonical name).** It's used to point one domain or subdomain to another domain, instead of pointing directly to an IP address.

DDGA: dictionary domain generation algorithm

DDI: **DNS, DHCP, and IP address management** (**IPAM**)—three critical network services that work together to provide **automated and centralized management of IP address spaces and name resolution** across enterprise networks.

DGA: domain generation algorithm

DNS: Domain Name System

DNS queries: A **DNS query** (Domain Name System query) is a request made by a device (usually a computer or mobile phone) to translate a **human-readable domain name** (like www.google.com) into a **machine-readable IP address** (like 142.250.190.68) so it can connect to the correct server on the internet.

GDPR: General Data Protection Regulation

HIPAA: Health Insurance Portability and Accountability Act

LLM: large language model

MFA: multi-factor authentication

MX Abuse: This involves malicious activities that exploit or misuse MX (mail exchange) records.

NIST: National Institute of Standards and Technology

NOD: newly observed domains

OFAC: **Office of Foreign Assets Control**, which is a division of the **U.S. Department of the Treasury.** It administers and enforces **economic and trade sanctions** based on U.S. foreign policy and national security goals.

OSINT: open-source intelligence

PCI DSS: Payment Card Industry Data Security Standard

PhaaS: phishing-as-a-service

RDGA: registered domain generation algorithm

SASE: Secure Access Service Edge

TDS: traffic distribution system

infoblox.

Infoblox unites networking and security to deliver unmatched performance and protection. Trusted by Fortune 100 companies and emerging innovators, we provide real-time visibility and control over who and what connects to your network, so your organization runs faster and stops threats earlier.

Corporate Headquarters
2390 Mission College Blvd, Ste. 501
Santa Clara, CA 95054

+1.408.986.4000
www.infoblox.com

Version: 20250730