DNSAFE
by Admiresty Corporation

WHITEPAPER

The Case for DNS-Layer Security

How filtering at the DNS layer stops threats before they reach your network

Published by Admiresty Corporation · dnsafe.net

Table of Contents

01 Executive Summary
02 The Evolving Threat Landscape
03 Why Traditional Security Misses DNS
04 How DNS-Layer Security Works
05 The Case for DNS as the Primary Control Point
06 ThreatGrid: Intelligence-Driven Filtering
07 Real-World Use Cases
08 What to Look for in a DNS Security Solution
09 Getting Started with DNSAFE
10 Conclusion

01 Executive Summary

Cyberattacks have become the defining operational risk of the modern era. Yet despite billions spent on endpoint protection, firewalls, and email gateways, one of the most exploited attack vectors remains largely unguarded in most organisations: the Domain Name System.

DNS is the phonebook of the internet — the protocol every device uses to translate a human-readable address like *malware-download.xyz* into an IP address before any connection is made. Because DNS queries precede all network traffic, filtering at this layer allows threats to be blocked at the earliest possible point — before a malicious payload is ever downloaded, before a phishing page is ever rendered, before ransomware ever calls home.

This whitepaper makes the case for DNS-layer security as a foundational control for businesses of all sizes, outlines how DNSAFE's ThreatGrid-powered platform implements it, and provides practical guidance for MSPs, IT managers, and security teams evaluating DNS filtering solutions.

91%	82%	<2ms	99.99%
of malware uses DNS to operate	of phishing domains are ≤30 days old	DNSAFE average query latency	platform uptime SLA

02 The Evolving Threat Landscape

The threat landscape has shifted dramatically over the past decade. Attackers have moved from opportunistic, broad-spectrum campaigns to targeted, multi-stage operations that bypass traditional defences with increasing ease.

The Numbers Tell the Story

DNS-based attacks have grown in both volume and sophistication. Research from multiple threat intelligence organisations consistently shows:

- **91% of malware uses DNS** during at least one stage of an attack — for command-and-control (C2) check-ins, data exfiltration, or payload retrieval.
- **Phishing domains are registered fresh** — 82% of phishing campaigns use domains less than 30 days old, specifically to evade reputation-based blocklists.
- **DNS tunnelling is rising** — attackers encode data inside DNS queries to exfiltrate information even through tightly locked-down firewalls.
- **Domain Generation Algorithms (DGAs)** allow malware to automatically generate thousands of possible C2 domains, making traditional blocklisting ineffective.
- **Ransomware phone-home before encrypting** — virtually all ransomware strains make DNS lookups for C2 infrastructure before beginning the encryption routine.

The Cost of a Breach

The financial consequences of a successful cyberattack extend far beyond immediate remediation. Downtime, regulatory penalties, reputational damage, and legal liability compound the impact for businesses of every size. For MSPs, a single compromised client can create cross-tenant risk and existential reputational damage.

Impact Category	Average Cost	Primary Driver
Data breach	$4.45M	Regulatory + remediation
Ransomware incident	$1.85M	Downtime + recovery
Business email compromise	$137K avg loss	Wire fraud / credential theft
DNS-based attack	$942K	Service disruption + data loss

Sources: IBM Cost of a Data Breach Report 2024; Sophos State of Ransomware 2024; IDC DNS Threat Report.

03 Why Traditional Security Misses DNS

Most organisations invest heavily in endpoint protection platforms (EPP), next-generation firewalls (NGFW), and email security gateways. Each of these controls is valuable — but none of them are designed to intercept threats at the DNS layer.

The Security Stack Gap

Control	What It Protects	DNS Gap
Endpoint AV / EDR	Files, processes, memory	Doesn't inspect DNS queries; misses fileless attacks over DNS
Next-Gen Firewall	Ports, protocols, IP ranges	DNS (port 53) typically whitelisted; no domain intelligence
Web Proxy / SWG	HTTP/HTTPS traffic	Requires agent; doesn't cover non-browser traffic or IoT
Email Gateway	Email-borne threats	No visibility into what happens after a user clicks a link
DNS Security	**Every DNS query**	**Covers all devices, all protocols, no agent required**

The table above illustrates a consistent pattern: traditional security tools operate *downstream* of the DNS resolution process. By the time an endpoint AV product detects a malicious file, or a NGFW blocks a connection to a known-bad IP, the initial DNS handshake has already occurred — and often the damage is already done.

04 How DNS-Layer Security Works

DNS-layer security operates on a simple but powerful principle: every internet connection begins with a DNS lookup. By intercepting and evaluating those lookups before a connection is established, threats can be stopped before they have any opportunity to execute.

The Resolution Flow

When a device attempts to connect to any internet resource — whether visiting a website, launching an application, or a background process phoning home — the sequence is:

1	**Device issues DNS query** The operating system or application asks "what is the IP address of this domain?"
2	**Query reaches DNSAFE resolver** Instead of going to a public resolver (8.8.8.8, 1.1.1.1), the query is directed to DNSAFE's resolver infrastructure.
3	**ThreatGrid evaluation** The domain is evaluated against ThreatGrid intelligence: threat score, category classification, custom policy rules, and allow/block lists.
4	**Decision in <2ms** If the domain is permitted, the real IP is returned and the connection proceeds normally. If blocked, NXDOMAIN or a redirect to the block page is returned instead.
5	**Zero connection made** Because the IP is never returned for a blocked domain, the device never establishes a connection to the malicious infrastructure.

No Agent Required

One of the most significant operational advantages of DNS-layer security is that it requires no software installation on end-user devices. Changing a single DNS setting — on a router, in a group policy, or on an individual device — is sufficient to protect every device on that network. This makes deployment radically faster and simpler than any endpoint-based alternative, and it means IoT devices, smart TVs, gaming consoles, and BYOD devices are all protected without any agent management overhead.

05 The Case for DNS as the Primary Control Point

DNS-layer security is not a replacement for a comprehensive security stack — it is the most cost-effective first layer of defence available. The argument for DNS as the primary control point rests on four pillars:

- **Universality**
  Every internet-connected device uses DNS. Filtering at this layer provides coverage across 100% of devices without any agent deployment — including IoT, OT, mobile, and unmanaged devices that cannot run endpoint security software.

- **Speed of Intervention**
  Because DNS is the first step in any connection, filtering here stops threats before any malicious content is downloaded, any credentials are entered, or any C2 communication is established. Compare this with endpoint AV, which must wait for a file to be written to disk before it can act.

- **Cost Efficiency**
  DNS filtering delivers a disproportionately high security return per dollar spent. Deployment requires no hardware, no agents, and minimal ongoing administration. For MSPs, DNS filtering is the highest-margin security service in the stack.

- **Complementarity**
  DNS security works alongside — not instead of — endpoint, email, and firewall controls. It eliminates the threat categories that other tools miss while adding zero friction to the user experience or IT overhead.

Key insight: A business that deploys DNS filtering as its only security control will prevent the majority of commodity malware infections, phishing attacks, and ransomware delivery attempts. It is the single highest-leverage security investment available.

06 ThreatGrid: Intelligence-Driven Filtering

The effectiveness of any DNS filtering solution is only as good as the intelligence behind it. DNSAFE is powered by ThreatGrid — a multi-source threat intelligence engine purpose-built for DNS-layer protection.

How ThreatGrid Works

ThreatGrid aggregates signals from multiple data sources, processes them in a unified scoring engine, and updates domain classifications continuously:

- **40+ threat intel feeds** — Commercial, open-source, and government threat intelligence feeds are ingested, normalised, and deduplicated in real time.
- **Passive DNS telemetry** — Query patterns across the DNSAFE resolver network surface emerging threats — domains that suddenly spike in query volume are flagged for investigation even before feed updates.
- **Domain Generation Algorithm (DGA) detection** — Machine-learning models identify algorithmically generated domains used by malware families to evade static blocklists.
- **Newly Registered Domain (NRD) tracking** — 82% of phishing campaigns use fresh domains. ThreatGrid monitors all new domain registrations and applies elevated risk scoring automatically.
- **Category classification** — 80+ content categories are maintained and updated daily, covering everything from adult content and gambling to cryptomining and anonymisation services.
- **15-minute update cycle** — Domain scores and classifications are pushed to all DNSAFE resolvers every 15 minutes, ensuring near real-time response to emerging campaigns.

Domain Intelligence

Every domain in the ThreatGrid database carries a composite threat score (0–100) derived from infrastructure analysis, registrant behaviour, hosting patterns, and observed malicious activity. DNSAFE customers can look up any domain through the Domain Intel tool in the portal to see the full scoring breakdown, category classification, and historical activity — before deciding whether to allow or block it.

07 Real-World Use Cases

For Businesses

- **Ransomware prevention** — Ransomware strains universally use DNS to locate C2 servers before encrypting files. DNS filtering blocks C2 communication at the query level, preventing the ransomware from activating even if it reaches the endpoint.
- **Phishing protection** — Newly registered phishing domains are caught by NRD filtering and ThreatGrid classification before users can visit them — even on the first query.
- **Policy enforcement** — Content categories (social media, streaming, gambling) can be toggled per policy, with schedule-based rules for flexible work environments.
- **Compliance support** — DNS filtering contributes to SOC 2, Cyber Essentials, and HIPAA compliance postures by providing demonstrable controls against malware and data exfiltration.

For MSPs

- **Multi-tenant management** — Manage DNS policies for all clients from a single portal. Per-tenant isolation ensures one client's rules never affect another.
- **High-margin upsell** — DNS security is the easiest security service to sell, deliver, and renew. Typical margins run 60–70% with minimal ongoing support overhead.
- **Automated threat response** — SIEM integration and PSA connector (ConnectWise, Autotask) allow DNS threat events to automatically create tickets and trigger workflows.
- **Roaming protection** — Laptop users off the corporate network remain protected via the DNSAFE roaming client — no VPN required.

For Home Users

- **Whole-family protection** — Configure DNS on the home router to protect every device — including smart TVs, game consoles, and IoT devices that can't run security software.
- **Parental controls** — Block adult content, gambling, and other age-inappropriate categories with no ongoing maintenance. SafeSearch enforcement and YouTube Restricted Mode work at the DNS level.
- **Ad and tracker blocking** — DNS-level blocking of advertising and tracking domains improves page load speeds and reduces exposure to malvertising campaigns.

08 What to Look for in a DNS Security Solution

Not all DNS filtering solutions are equal. When evaluating options, organisations should assess the following criteria:

01 **Intelligence quality and freshness**
How many threat feeds does the solution aggregate? How frequently are classifications updated? A solution that updates daily is orders of magnitude less effective than one that updates every 15 minutes against fast-moving campaigns.

02 **Coverage breadth**
Does it cover standard DNS (port 53), DoH, and DoT? A solution that only filters port 53 can be trivially bypassed by any application that uses encrypted DNS.

03 **Deployment flexibility**
Can it protect roaming devices without a VPN? Does it support router-level deployment for networks with unmanaged devices? Is there a zero-agent option?

04 **Logging and visibility**
Does it provide per-query logs with enough context to investigate incidents? Can logs be forwarded to a SIEM? Log retention period is also a key compliance consideration.

05 **Policy granularity**
Can policies be applied per-location, per-device, or per-user group? Are schedules supported? Can allow/block rules be applied at different levels of the hierarchy?

06 **MSP-readiness**
For service providers: is there true multi-tenant isolation? White-label capability? PSA and RMM integrations? Consolidated billing? These features determine whether the product is genuinely built for MSPs or just a single-tenant solution with a "portal."

07 **Performance and reliability**
DNS filtering adds a step to every internet connection. Any latency at the resolver level is felt by every user. Look for sub-5ms resolution times and a credible uptime SLA.

09 Getting Started with DNSAFE

DNSAFE is designed to be operational in minutes, not weeks. There is no hardware to provision, no agents to deploy before you can start, and no long procurement cycle. The following steps will have your network protected within the hour.

Quick Start (Under 10 Minutes)

1. **Sign up** Create an account at my.dnsafe.net (home users) or portal.dnsafe.net (MSPs and businesses). No credit card required for the 14-day trial.
2. **Create a policy** Define which threat categories and content categories to block. DNSAFE's recommended policy template enables all security categories with a single click.
3. **Register your network** Add your public IP range as a network location and assign your policy to it.
4. **Update your DNS** Set your router or DNS server's primary resolver to 3.12.124.91. All devices on the network are now protected.
5. **Verify coverage** Visit dnsafe.net/check to confirm your DNS is routing through DNSAFE. The page will display your registered location name.

Resolver Addresses

Protocol	Address	Port
Standard DNS	`3.12.124.91`	53 (UDP/TCP)
DNS-over-TLS (DoT)	`3.12.124.91`	853 (TCP)
DNS-over-HTTPS (DoH)	`https://api.dnsafe.net/dns-query`	443 (HTTPS)

Resources

- Product website: dnsafe.net
- Documentation: docs.dnsafe.net
- API reference: api.dnsafe.net/docs
- MSP partner portal: portal.dnsafe.net
- DNS check tool: dnsafe.net/check
- Sales enquiries: sales@dnsafe.net
- Partner programme: partners@dnsafe.net

10 Conclusion

The case for DNS-layer security rests on a straightforward premise: every threat that uses the internet must first use DNS. Filtering at this layer — before any connection is established, before any payload is delivered, before any credentials are stolen — is the most efficient, most comprehensive, and most cost-effective security control available.

Traditional security tools remain necessary and valuable. But they operate downstream of DNS, leaving a critical gap that attackers have learned to exploit systematically. The 91% of malware that uses DNS to operate does so precisely because most organisations leave this channel unmonitored.

DNSAFE closes this gap with enterprise-grade threat intelligence, real-time filtering, and a deployment model that protects every device on a network — managed or unmanaged — with a single configuration change. For MSPs, it delivers the highest-margin security service in the stack. For businesses, it eliminates the most common attack vector. For home users, it puts world-class protection on every device without complexity.

Start your 14-day free trial

No credit card required. Setup takes under 10 minutes.

dnsafe.net

About Admiresty Corporation

Admiresty Corporation is the company behind DNSAFE, an enterprise DNS security platform serving MSPs, businesses, and home users globally. Powered by ThreatGrid, DNSAFE provides DNS-layer threat filtering, content controls, and network intelligence through a fully managed cloud platform. Learn more at dnsafe.net.