DNSFilter

THE ROLE DNS PLAYS IN GOOD IT INFRASTRUCTURE

2023 EDITION

DNS TOUCHES EVERYTHING

DNS is a protocol that is usually taken for granted. It's sometimes overlooked because it's so ubiquitous. And yet when something goes wrong, it's the first thing we blame—often, rightfully so.

While DNS is not given the same level of attention as other elements of your IT infrastructure, it's one of the few components that likely interacts with every other piece of your IT setup.

DNS can either act as a **shield** and be the first line of defense in your network security infrastructure, or it can be your first point of **vulnerability.** Your inclusion (or exclusion) of protective DNS in your IT infrastructure determines which one it will be.

Understanding DNS vulnerabilities and mitigating them is essential as your business faces increasingly sophisticated and malicious attacks from malware, ransomware, phishing, and next-generation threats (e.g. Botnets and Cryptomining)—attacks that may compromise the health of your IT infrastructure and the data within. These DNS-based malicious attacks can also cost your business millions. The average cost of an enterprise data breach in 2022 was $4.35M[1].

This analysis examines what role DNS plays in your IT infrastructure and the actions you can take to protect your DNS network.

WHAT IS DNS?

Popularly referred to as the “phone book of the internet,” DNS stands for Domain Name System.

The Domain Name System maps domain names (websites) to their numerical IP addresses so that users don’t need to memorize a series of numbers.

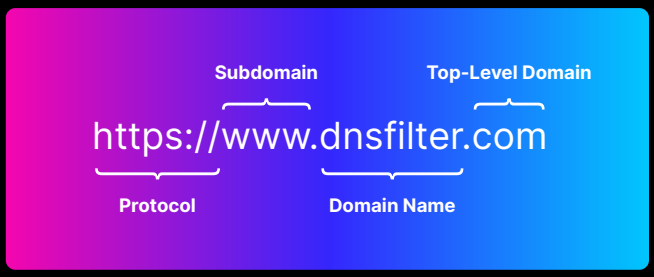

When you type a URL in a web browser, you’re really asking a DNS server, “What is the IP address of this website?” The DNS server responds with the IP address and takes you to the website you’d like to visit. This is an incredibly simplified version of what DNS does—it can get a little more complicated. For instance, DNS also determines which mail servers should be used for a domain, or if content servers need to be contacted to retrieve images or video.

DNS is a hierarchical and decentralized naming system that ensures all names are completely unique. And because DNS servers are distributed (i.e., not located in a single place), it has allowed the internet to grow the way it has.

In contrast, ARPAnet, the precursor to the internet, was not distributed. A single HOSTS.TXT file housed all of the information about the hosts. Everything was routed through this *single* file, which meant if anything happened to that file or the servers it ran on, everything would become inaccessible. This was not a scalable solution, contributing to the creation of DNS as an alternative—and eventual foundation—of the modern internet.

There are two types of DNS servers: **Authoritative** and **Recursive.**

The IP address of websites are stored on authoritative name servers.

When someone enters a domain, the *recursive* DNS server sends this query to the following places:

1. The root name server
2. The TLD (Top Level Domain) server
3. The authoritative name server

The authoritative name server is the part of this equation that has the answer to “What is the IP address of this website?” The recursive DNS is then “resolved” and the site is accessed in the browser.

Once a user has visited a site, the IP of that site will be cached locally until at some point when it’s refreshed. This means that recursive DNS, when information is cached, can sometimes operate without contacting the authoritative DNS server.

When you implement protective DNS and content filtering, this is done through recursive DNS. Recursive DNS is sometimes referred to as a “DNS resolver” since it “resolves” DNS queries. If you attempt to access a malicious domain with protective DNS in place, that domain will not resolve at all and instead you’ll see a page letting you know that the content is blocked.

HOW DNS CAN BE USED FOR BAD?

Despite its ubiquity, and that every single Google search and navigation to Facebook.com relies on it, DNS is also used by bad actors. Every day roughly 100,000 new domains are registered and many to be used in a malicious way (e.g. Spear-phishing). Since it is used in everything we do online it can also be used for nefarious reasons. In addition, the DNS protocol dates back to 1981 and throughout the years is now just getting more attention to further improving the service by using encryption.

EXPOSURE TO MALICIOUS SITES

Nearly every online attack that has ever succeeded could be considered a DNS attack, as it used the DNS protocol to: Operate, identify its targets, and spread. Cybersecurity Ventures reports that there is a cyber attack on businesses every 11 seconds[2]. They're predicting that frequency to increase to **every 2 seconds by 2031.** From Q4 2022 to Q1 2023, DNSFilter saw a 281.88% increase in identified malicious phishing, malware, crypto and botnet traffic. And in 2022 alone, data of more than 422 million individuals[3] was breached—an increase of 40% YoY. A single breach can put a company at incredible financial and legal risk, as well as create a PR nightmare.

Cyber attacks come in many different forms rooted in different motives, from attention seekers to profiteering nation-states. With the increase in remote work and the surge in smartphone usage, BYOD (bring your own device) and edge devices are a persistent and growing aspect (and weakness) of your organization's infrastructure. These devices are no longer protected by company infrastructure, leaving employees more open to attacks. Unfortunately, 51% of organizations do not utilize any cybersecurity measures to prevent attacks[4]. And IT professionals report that the biggest cybersecurity risks to their company are endpoints such as smart phones and laptops.

The major DNS attacks that users are vulnerable to include:

PHISHING

Every minute, nearly $18,000 is lost because of a phishing attack[5]. These attacks make up 80% of reported security incidents. That isn't all that surprising when you consider that 3.4 billion phishing emails are sent each day[6].

Phishing often starts with an email designed to create trust by mimicking a reliable source or brand. The strategy is to trick the target into handing over personal information, often by clicking a link and sharing sensitive information such as passwords, banking details, or credit card numbers.

MALWARE AND RANSOMWARE

In 2022, records impacted by **ransomware more than doubled** compared to the year prior[7]. Malware is a broad term that is short for “malicious software.” It can be spread through forced downloads, phishing schemes, or malicious ad content. Ransomware is a type of malware that enables hackers to encrypt user files and then demand a ransom. By 2031, ransomware will cost organizations $265 billion annually[8].

DDOS (DISTRIBUTED DENIAL OF SERVICE) ATTACKS

In 2022, DDoS attacks increased by 67% YoY. And network-layer Mirai botnet attacks increased 405% from one quarter to the next[9]. DDoS of a DNS server is a deliberate attempt to overload the server by making a large number of requests to prevent queries initiated by genuine users from being resolved. This results in business down-time (averaging 24 days in 2022), leading to revenue loss of $5.1 million per organization[10].

CACHE POISONING AND DNS SPOOFING

Cache poisoning is a class of malicious attack that can happen outside of your own security boundary. It is the direct interference with the DNS records or a DNS cache to misdirect a user's request.

By editing a DNS resolver's locally-held copy of the website’s IP address listings, it is possible to redirect the user to an alternative site. This is known as DNS spoofing or cache poisoning. By mapping a false IP address to a genuine site or device, the IP redirect may be hijacked and the user taken to an irregular location.

TUNNELING

Because DNS traffic is essential for internet use, it is typically allowed through any firewall (both inbound and outbound).

DNS tunneling uses DNS requests as a command and control channel for malware. Inbound DNS traffic can carry commands to resident malware. Outbound traffic can "exfiltrate" (take out) sensitive data or provide

responses to the requests from infected hosts.

There are few restrictions on what data a DNS request may contain, thus allowing data packets to be exfiltrated in this way. Being that the protocol was designed to look for domain names of websites, which could be almost anything, these fields can be hijacked to carry data.

These requests are designed to go to attacker-controlled DNS servers, ensuring that they can receive the requests and respond in the corresponding DNS replies.

The reality of the security issues network administrators face goes far deeper than preventing such active and aggressive attempts to misdirect users.

TYPOSQUATTING

Typosquatting takes advantage of human error when typing website URLs. It's commonly used in phishing schemes. Cybercriminals use typos (such as "paaypal.com") and host fully-branded sites ready to capture user data by fooling the user into believing they reached their intended destination.

During the 2020 election, at least 500 domains related to each candidates' campaign were discovered. At DNSFilter, our AI discovers new domains every week commonly mimicking sites like Chase, PayPal, and Microsoft.

IMPROPER CONFIGURATIONS

If a device, website, or service does not have its DNS record set up correctly, a router will not be able to find an IP address, and services will fail. That means every printer, router, computer, server, and switch is impacted by DNS.

By taking advantage of an improperly-configured DNS setup, attackers can avoid proxies or web filtering. Also, data exfiltration via DNS may go unobserved as the data is hidden within normal network traffic and such data packets pass unchallenged by the firewall.

DNS SUFFIXES

It pays to play nice with your own DHCP (Dynamic Host Configuration Protocol) setup. The TCP/IP (Transmission Control Protocol/Internet Protocol) settings for each network interface can be assigned a unique DNS suffix, either statically, or dynamically via DHCP. The browser first calls the primary suffix, if that is not set it will apply any connection-specific Suffix.

However, this procedure may be overridden, because the TCP/IP settings for network interfaces share an optional set of DNS suffixes that are known as a "search list". Search list entries are resolved in favor of, i.e. take priority over, any primary suffix, its parents, and connection-specific suffix. For this reason, be careful when implementing DNS suffixes for internal use.

POINT TO THE RIGHT DNS RESOLVERS

One misconfiguration that can leave you vulnerable is pointing to the wrong (or multiple) DNS resolvers. This usually occurs after you've chosen a DNS service provider to filter out malicious content.

The issue here is when the DNS resolvers are either not set, or you are pointing to more than one DNS resolver and mixing resolvers.

For instance, with mixed revolvers, something as simple as leaving Google's public DNS can create unintended issues. Administrators may think this configuration is a good solution and might act as "failover" in case one resolver fails—but it doesn't. Most systems are going to use every resource available to them. You might have some requests going to your protective DNS (such as DNSFilter) and others going to Google. This will be clear in your reporting when you go to look at your network activity and realize a portion of your network is being neither filtered nor scanned for threats.

This is also an important consideration when using local domains. Be sure to specify local domains that need to be resolved internally.

APPROPRIATE TTL INTERVAL

TTL (time to live) is used by DNS resolvers to know when to return to the authoritative record to check for updates.

If you are doing regular domain updates, temporarily lower your TTL on a record prior to making changes/server moves. This way, your old record doesn't get stuck in recursive resolver caches. Remember, you may have to wait your full TTL until your new TTL registers.

ENCRYPTION

DNS is mostly unencrypted. DNS sends requests in plain text, and it's not encrypted by default in most scenarios. This means in instances where encryption is not applied you are vulnerable to:

- Spoofing – Forged DNS requests usually come in the form of a man-in-the-middle attack where a malicious actor will temporarily redirect users to a fake login page to collect personal information or login credentials
- Tracking – When untrustworthy entities can view your DNS requests and collect information on you, this data can be sold to advertisers

The two warring options for DNS encryption are DNS-over-HTTPS (DoH) and DNS-over-TLS (DoT). Other protocols put forth in the past that don't offer end-to-end encryption, but are still valid security options to use with DNS, are DNSSEC and DNSCrypt.

As far as *security* is concerned, DoH and DoT are essentially identical. The differences are in how they are implemented.

DoH operates on port 443 while DoT operates on port 853. Implementing DoH in browsers for example, makes sense because other web traffic is already using that port. This works well because DoH is very concerned with privacy, but from a sysadmin perspective that means you lose valuable network insight because that encrypted DoH traffic is simply blending into the already noisy HTTP traffic on port 443. Because of this, even if you have content filtering on your company network, DoH could actually bypass that filtering and act as a workaround.

At DNSFilter, **we prefer DoT** to ensure that DNS requests that are handled outside of the browser environment (e.g., Slack messages, links embedded in Excel or other desktop applications) are encrypted. DoT also has slightly lower latency than DoH, meaning DNS requests are resolved faster.

Every platform implements DoH differently, and the pendulum of opinion will swing according to each infrastructure's needs. This is why DoH **remains a moving target.**

DNS + THE REST OF YOUR NETWORK

23%

Employees using password protection on all personal devices

Your IT infrastructure is a complicated and ever-changing environment. If you implement a new BYOD (bring your own device) policy to help your team transition more smoothly to remote work, you've just changed your infrastructure. If you let a freelance designer add their laptop to your network, you've just changed your infrastructure. If you allow an employee to add unapproved security software to their laptop, you've just changed your infrastructure.

In one survey, only 23% of employees use basic password protection on all their personal devices[11]. And 39% of those same workers use personal devices to access corporate data, often via services and applications hosted in the cloud. This exposes the need for companies to apply protection as broadly as possible across their infrastructure.

DNS is one of the broadest ways to apply protection because it means blocking threats (and unwanted content) essentially at the source—or as close to it as possible. With every new connection or change to your IT infrastructure, DNS is involved.

DNS + VPNS

The exponential increase in remote working has put many enterprise-level VPNs under massive pressure—about 60% of internet users utilize a VPN at least once per week[12].

Many corporations had to quickly rethink their strategy when their remote workforce exceeded the intended maximum capacity of the VPN infrastructure. 71% of organizations were forced to increase their VPN capacity due to a surge in remote work at the start of the 2020 pandemic[13].

There are essentially two VPN approaches to choose from:

- Split-tunnel
- Full-tunnel

With a full-tunnel VPN, *all* traffic is going through the VPN. That includes every DNS query. And in a full-tunnel approach, applying any sort of threat protection or content filtering would need to be done at the network level.

This is the easiest way to set up a VPN, but it can also impact performance.

Take for instance a new employee who is working remotely in San Francisco while their company is located somewhere on the East Coast. That employee is set up with a VPN on their work computer so they can access internal file servers, be protected with the company firewall, and have content filtering at the network level.

The issue is that a full-tunnel VPN is their DNS requests are traveling across the country and accessing servers that are nowhere near them. The same issues might occur for employees traveling to another country for a conference. Everything will move much slower because of how far their DNS requests have to travel because of their VPN solution.

At DNSFilter, we tend to recommend that users implement a split-tunnel VPN. With a split-tunnel VPN, internal DNS requests will go to company servers. This way, if a user just needs to access a website, that DNS query does not need to travel very far. The DNS resolver will route users to local servers and return the request much faster.

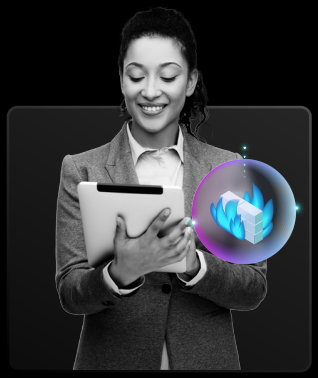

DNS + FIREWALLS

Firewalls and protective DNS offer different things. Protective DNS is concerned with the web content that is allowed on your network while firewalls prevent your corporate networks from exposure to external threats. They also operate at different layers of the OSI (Open System Interconnection) model.

Some firewalls, however, do have some content filtering capabilities. This can lead to conflicts when combined with your DNS filtering solution. Both tools are necessary layers of protection within your IT infrastructure. If they are *both* attempting to apply content filtering to domain requests, it may result in failure of a domain to resolve or inconsistent resolution as we mentioned earlier.

Another consideration for how your firewall interacts with DNS is circumvention of DNS filtering. Some employees may attempt to change the DNS settings on their device to bypass company filtering. In this case, firewalls can be configured to forward *all* DNS queries to a service like DNSFilter to ensure that company filtering is applied correctly.

DNS + RMM AND MDM TOOLS

Remote networking monitoring tools, sometimes referred to as RMM (Remote Monitoring & Management), are another key infrastructure component impacted by DNS. RMM and MDM tools enable Managed Service Providers in particular to control various programs, deployed on various devices, all from a single location. Many RMMs from providers will also bundle certain tools to make it easier.

These tools are reliant on DNS to run and communicate with applications running on system administrators' computers. As an example, tools like N-Able that provide a firewall aspect of their product need to alert users when **their own DNS names change**. Some users may block domains that fall into certain categories, and if these domains happen to fall into those categories then their product will fail to work.

To assist network administrators to act in response to security threats from a distance, DNSFilter can be deployed using **various RMM and MDM tools**. They can also connect to the DNSFilter API to generate reports in their RMM or MDM solution.

The fine control that DNSFilter provides over edge devices held by Roaming Clients means that you can **remotely**:

- Initiate a new filtering policy or schedule
- Disable individual agents
- Enable auto-registration of edge devices

WHAT DOES "GOOD" IT INFRASTRUCTURE LOOK LIKE?

Good infrastructure is not a product of simply implementing a recommended tech stack and being done with it. The *approach* you take towards IT and cybersecurity is just as important as the tools you put in place—if not, moreso.

A LAYERED APPROACH SECURITY

Layered security means that you use several methods to protect your organization at different layers of vulnerability. This tightens security by protecting individual assets through a multitude of methods.

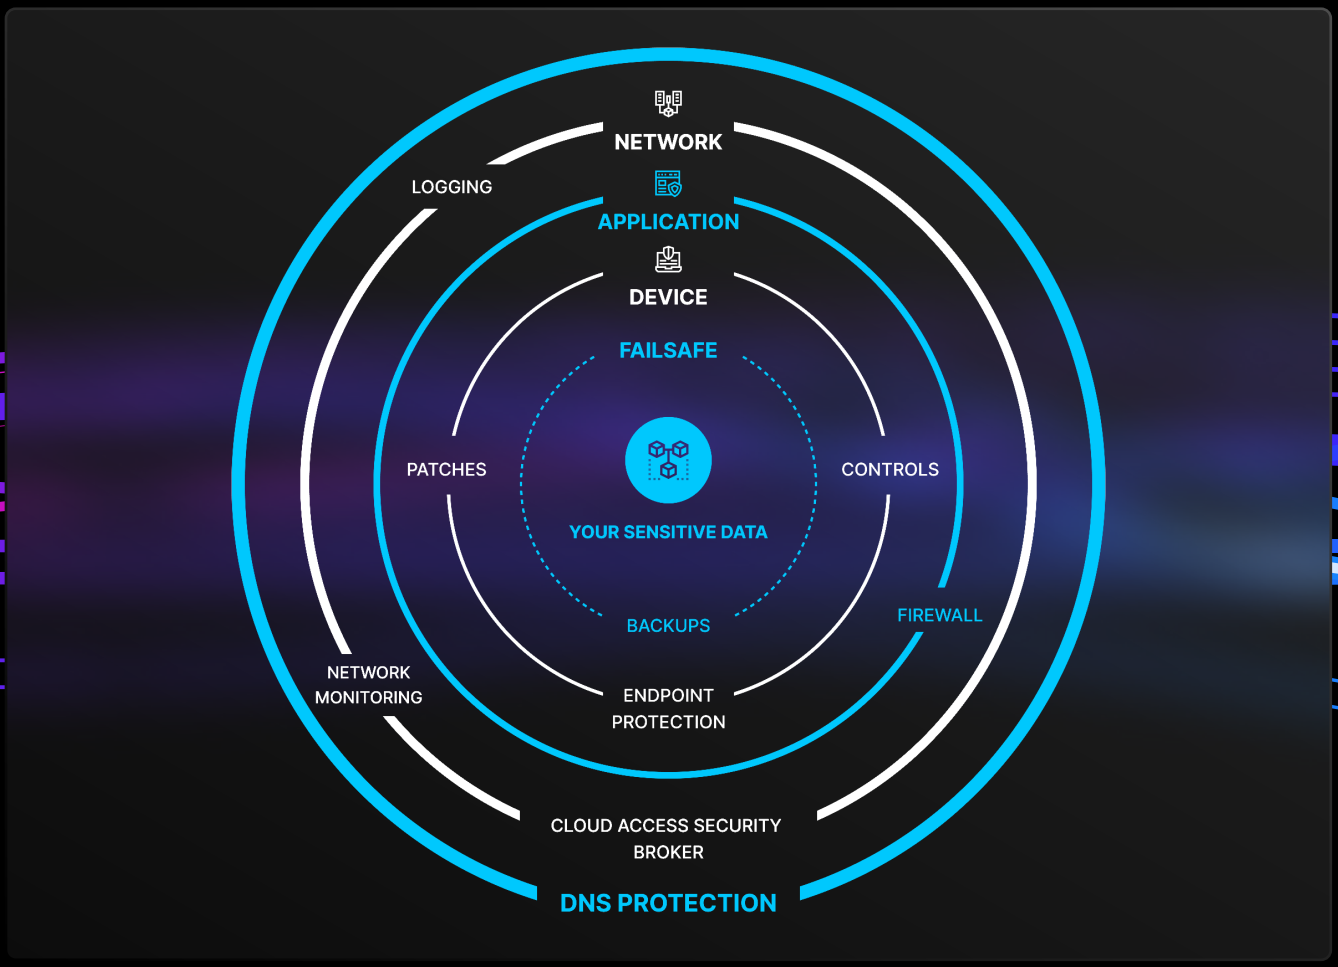

A layered approach means that if something does go wrong, you have “backups” in place to increase the likelihood that your organization will be protected.

Let’s look at an organization that has the following in place:

- Password manager
- Protective DNS

The password manager protects employees from entering credentials on websites that are deceptive. If a user is linked to a website impersonating Chase.com and asks for Chase login credentials, the password manager won’t recognize that website and will not populate the credentials.

But what if the password manager fails in this instance? The employee, frustrated by not seeing their password manager work, can go into the password manager tool and pull their username and password to enter into the fraudulent website.

If they also have protective DNS implemented, they’ll never get to that fake Chase.com page in the first place—it will be blocked as “deceptive.”

These two security measures clearly provide very different features. But in the end, they’re both protecting end users from becoming the victim of a cyber attack. They’re just implementing their security at two different layers. Protective DNS is implemented at the internet layer, securing every domain request. The password manager works at the “people” layer, preventing users from making a mistake and entering credentials on the wrong page.

ZERO TRUST

A layered security approach is an aspect of the popular “Zero Trust” philosophy. Zero Trust is a concept created by John Kindervag while he was Principal Analyst at Forrester Research in 2010.

The concept is relatively simple: Organizations should not trust *anything* inside or outside of their network by default. Everything on the network needs to be verified before it is granted access. This is in stark contrast to the old mentality that security should only be concerned with external threats. Hackers who gain access to internal infrastructure are able to navigate systems easily and make changes (or extract data) with incredible ease when the infrastructure itself is never asking for further credentials from these hackers.

Because of this change in attack strategy, a model needed to be developed that could protect organizations even when the threat is *inside* the organization.

For this concept to actually work, the components of your IT infrastructure need to be Zero Trust in nature. And your infrastructure needs to be *finite*. All additions to your infrastructure need to come from IT. You need to be able to keep an eye on shadow IT and prevent the usage of certain applications if they are not Zero Trust. This means:

- Disallowing administrative privileges on workstations outside of IT
- Setting up multi-factor authentication (MFA) wherever possible
- Creating and sharing a policy for implementing new IT software

Layering your security measures allows you to validate that what you're accessing is a trusted resource at *multiple* levels. It also allows you to find discrepancies within your infrastructure. For instance, is there something your firewall is blocking that another aspect of your network is allowing?

Zero Trust demands that your infrastructure always be asking, "Should I trust this resource?" But it also gives you additional data about the nature of your infrastructure and how all of the tools that are part of your multi-layered approach are working together to keep threats out and trusted resources in.

THE TOOLS YOU NEED

With all that being said, there are certain tools that are absolutely necessary for your IT infrastructure to function in a way that protects your employees.

Let's consider what the layers are that need to be protected:

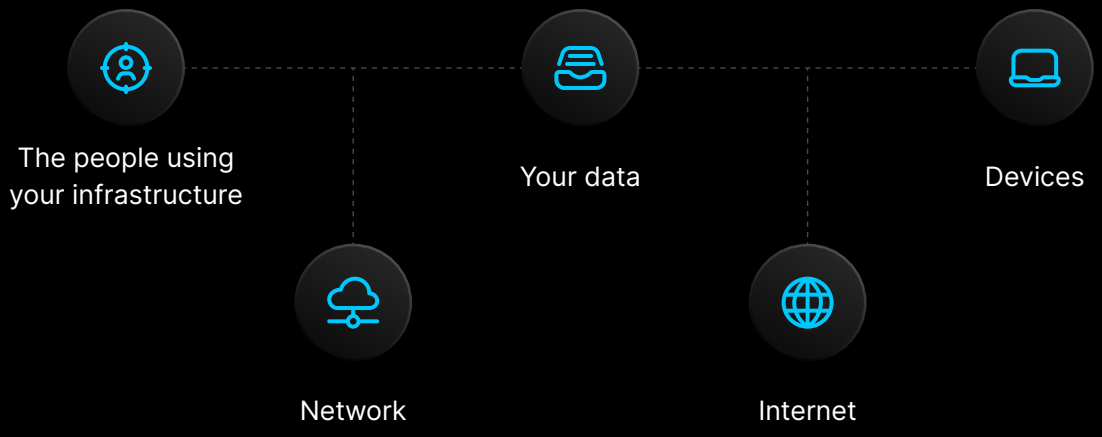

How do you protect your users with a Zero Trust mindset? We've already mentioned Password Managers, multi-factor authentication, and disabling administrative privileges on workstations. When it comes to tools, these are the *bare minimum* that you need to have in place to be a Zero Trust cybersecurity-conscious company.

When it comes to protecting your data, you need to actively protect the location of the data. This can be applications or devices. Not only should you have MFA and secure passwords in place to protect this data, wherever it is, but it should be backed up somewhere else equally secure.

For devices, you need comprehensive endpoint protection. This includes antivirus software, and also encompasses threat detection, device management, and response. But protective DNS, when deployed at the device level, is also a part of your endpoint protection.

Your network *needs* to have a firewall. It's not something you can neglect in your IT infrastructure. And when deployed at the network level, protective DNS is also part of your network security.

And finally, no matter how it's deployed, protective DNS secures users at the internet level.

It can be argued that multiple tools can have an impact at varying layers of your security approach, but protective DNS is the solution that has the greatest impact across the most possible layers. It also acts as a failsafe for various components of your infrastructure, including your password manager, antivirus software, and firewall.

SETTING UP PROTECTIVE DNS

Now that you understand the role that DNS plays in your IT infrastructure, and the necessity of protective DNS for both a multi-layered and Zero Trust approach to cybersecurity, here is what you need to consider when choosing a protective DNS solution.

FILTER TRAFFIC THROUGH DNS

DNS is your first line of defense. As mentioned at the beginning of this whitepaper, your infrastructure can either experience its first vulnerability at the DNS layer, or the DNS layer can act as a shield protecting the rest of your architecture. Because of that, protective DNS is absolutely necessary.

Filtering traffic through DNS on your network or at the device level allows you to block malicious site content, such as malware and phishing. It also enables companies to meet compliance standards such as CIPA, or whatever internal standards that company may have.

The most effective protective DNS solutions are able to detect new threats on their own, even if they've never been seen before.

AVOID FREE RESOLVERS

It may be tempting to implement a free DNS resolver that allows you to enable content filtering, but that will not act as the DNS shield you need it to.

The NSA recommends that all enterprise networks leverage enterprise-grade DNS resolvers in order to benefit from their cybersecurity defenses, facilitate safe access to local network resources, and protect data held on the internal network.

One of the main reasons the NSA urges enterprises to stay away from free resolvers is the ability to apply DNS encryption. Encryption is not a default with DNS resolution, and the free resolvers don't necessarily make this feature available.

To take advantage of all of the possible protections that DNS filtering can give companies, it's best that they rely on vendors that are held accountable by their customers. This promotes innovation of the product and encourages the programmers behind these tools to keep the software updated—meaning no new vulnerabilities are introduced because of out-of-date code.

CHOOSE RELIABILITY

As a dedicated DNS resolving service, DNSFilter is committed to providing layer after layer of **redundancy**.

DNSFilter supports multiple routing paths via Anycast. This is managed using the BGP (Border Gateway Protocol), which supports our IP space from multiple locations. With over 50 data centers worldwide, you can be guaranteed that you are provided the path with the lowest possible latency.

This is a far superior approach to DNS resolution provided by unicast services, i.e., ones that resolve all requests using a single server.

DNSFilter is composed of two separate Anycast networks: DNS1 and DNS2. If you submit a request and the DNS1 server does not answer (because of an outage, latency, etc.), then DNS2 will return your request instead—with no impact on speed.

Our DNS1 and DNS2 networks are completely different in their composition by design. We use different hosting providers, data centers, and server architectures. This is the same strategy that is applied by the public authoritative DNS servers of the internet.

CONCLUSION

Your IT infrastructure is both dependent on and intricately merged with DNS. Everything that happens online is part of DNS. Employees working from home are reliant on a growing number of web applications, exposing companies to more threats than ever before—and DNS is a critical point of vulnerability that you need to protect.

Without protective DNS in place, DNS can open up vulnerable parts of your network. But when you implement security via DNS, you turn DNS into a shield and add a much-needed layer to your Zero Trust IT security stack. In this way, DNS can become the strongest security tool in your infrastructure.

Don't take DNS for granted and leave your network (and employees) vulnerable to an attack. Adopt a Zero Trust mentality and start filtering every single domain that users attempt to access on your network. It could mean the difference of experiencing a catastrophic data breach, or simply going about your day.

Ensure it's the latter with a 14-day trial of DNSFilter.

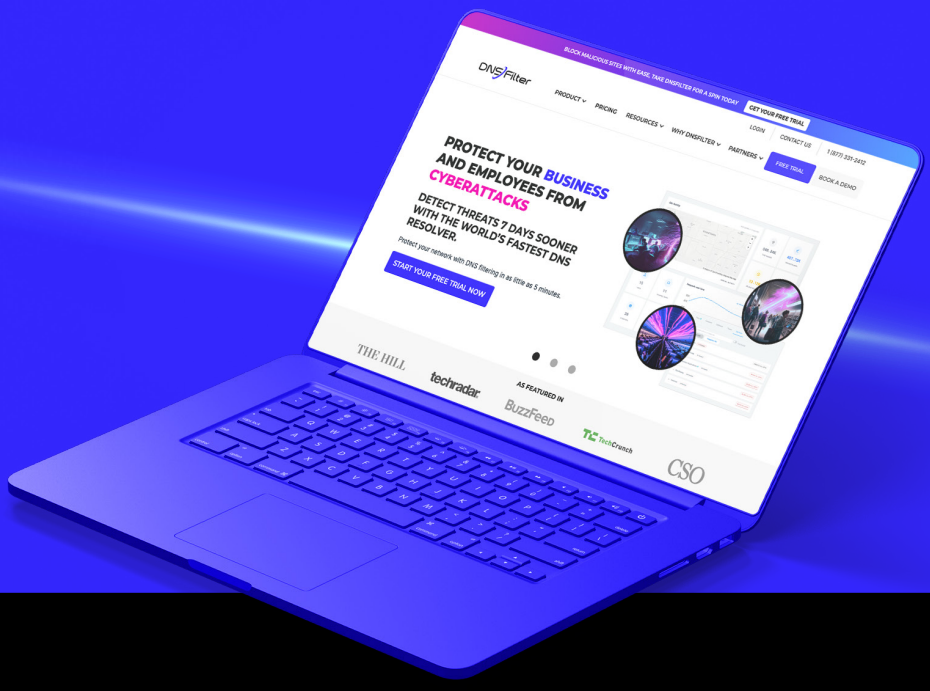

SOURCES

1. Cost of a Data Breach Report 2022
   https://www.ibm.com/downloads/cas/3R8N1DZJ

2. Top 10 Cybersecurity Predictions And Statistics For 2023
   https://cybersecurityventures.com/top-5-cybersecurity-facts-figures-predictions-and-statistics-for-2021-to-2025/

3. Data Breaches Hit Lots More People in 2022
   https://www.cnet.com/tech/services-and-software/data-breaches-hit-lots-more-people-in-2022/

4. 51 Small Business Cyber Attack Statistics 2023 (And What You Can Do About Them)
   https://www.getastra.com/blog/security-audit/small-business-cyber-attack-statistics/

5. New Security Report Breaks Down Increase in Cyber Attacks Due to Remote Work; Lack of Training, Overwhelmed IT Departments are the Main Issues
   https://www.cpomagazine.com/cyber-security/new-security-report-breaks-down-increase-in-cyber-attacks-due-to-remote-work-lack-of-training-overwhelmed-it-departments-are-the-main-issues/

6. The Latest 2023 Phishing Statistics (updated May 2023)
   https://aag-it.com/the-latest-phishing-statistics/#:~:text=Yes%2C%20phishing%20is%20the%20most,emails%20are%20sent%20every%20day

7. Ransomware attacks declined in ’22 but more records being compromised
   https://www.securityinfowatch.com/cybersecurity/article/21292765/ransomware-attacks-declined-in-22-but-more-records-being-compromised

8. Global Ransomware Damage Costs Predicted To Exceed $265 Billion By 2031
   https://cybersecurityventures.com/global-ransomware-damage-costs-predicted-to-reach-250-billion-usd-by-2031/

9. 20+ DDoS attack statistics and facts for 2018-2023
   https://www.comparitech.com/blog/information-security/ddos-statistics-facts/

10. Ransomware, carding, and initial access brokers: Group-IB presents report on trending crimes
    https://www.group-ib.com/media-center/press-releases/gib-2021-2022-report/

11. 43 Interesting Password Statistics in the Cybersecurity World
    https://www.g2.com/articles/password-statistics

12. 2023 VPN Usage Statistics
    https://www.security.org/vpn/statistics/

13. VPN Risk Report
    https://recursos.bps.com.es/files/991/04.pdf

DNSFilter