A Guide to

Protective DNS Security

What Is **DNS?**

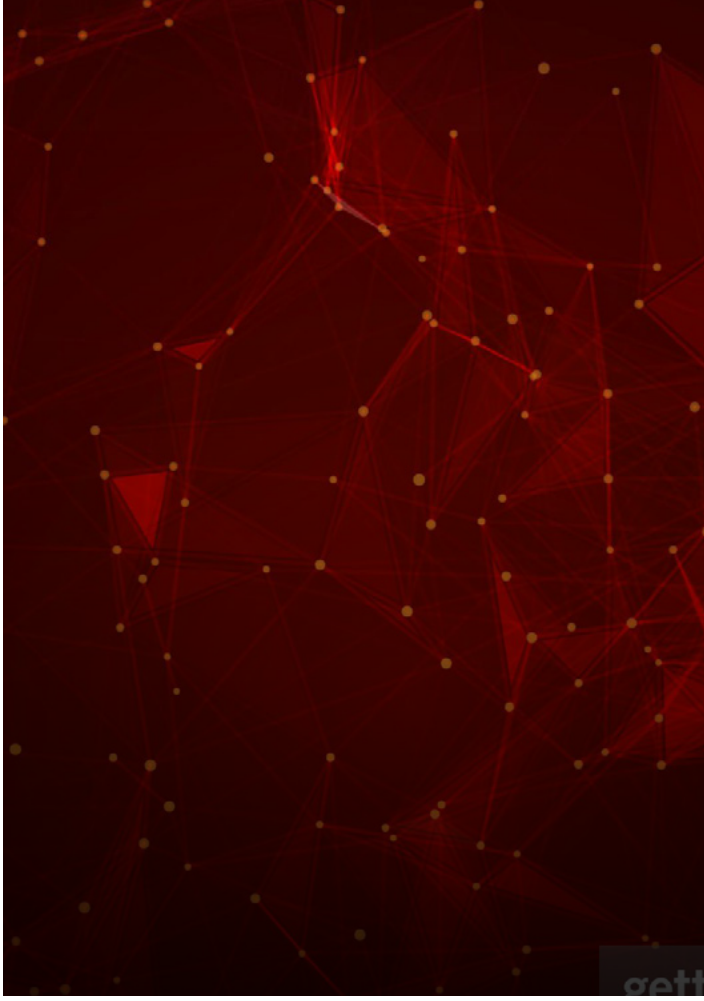

The Domain Name System (DNS) is the hierarchical and decentralized naming system used to identify computers, services, and other resources reachable through the Internet or other Internet Protocol (IP) networks. The resource records contained in the DNS associate domain names with other forms of information. These are most commonly used to map human-friendly domain names to the numerical IP addresses computers need to locate services and devices using the underlying network protocols but have been extended over time to perform many other functions as well.

The Domain Name System has been an essential component of the functionality of the Internet since 1985. An often-used analogy to explain the Domain Name System is that it serves as the phone book for the Internet by translating human-friendly computer hostnames into IP addresses. For example, the domain name www.example.com translates to the addresses 93.184.216.34 (IPv4) and 2606:2800:220:1:248:1893:25c8:1946 (IPv6).

The DNS can be quickly and transparently updated, allowing a service's location to change without affecting the end users, who continue to use the same hostname. Users take advantage of this when they use meaningful Uniform Resource Locators (URLs) and e-mail addresses without having to know how the computer locates the services.

In fact, almost nothing on the Internet would function without the use of the DNS. It is seamlessly used every minute of every day by people, devices, and services around the world to identify and locate their required connections somewhere on the global Internet.

As such, DNS is critical for everyday use; however, it is also utilized by nefarious and malicious bad actors. It is estimated that over ninety percent of all malware, ransomware, and the like utilizes domain names for command and control [See report here].

Bad actors need to communicate with their malware to provide instructions, guide lateral movement, perform data exfiltration, exchange binaries, perform encryption, and other nefarious activities – they do this by communicating with a known (or programmatically changing) Internet resource known as "command and control" or C2.

Utilizing domain names allows these bad actors to quickly change IP addresses and/or locations for their C2 to try and stay ahead of attempts to block them; thus, the malware needs to use the DNS every time it beacons out to communicate. It is both the simplicity and the ubiquity of the DNS which means it is vital for correct operations of almost everything on the Internet, and easily exploitable by bad actors.

What Is **Protective DNS?**

Given that the DNS can be used for both legitimate and nefarious lookups, Protective DNS solutions exist to help the administrator and/or client of the DNS determine whether a given request is good, suspicious, or nefarious.

Unlike so-called "DNS security" solutions which are aimed at protecting the DNS itself from attacks like DDOS, Protective DNS does not protect the DNS at all but rather provides an intelligence and security layer on top. Organizations such as CISA and the US Government encourage the use of Protective DNS to strengthen one's cyber defense posture and add a "protective shield" around an organization as part of a security-in-layers approach.

As of mid-2023, it is a recommended part of a SASE architecture, and cyber insurance carriers are starting to ask if organizations have a Protective DNS solution or not. Various national governments, including but not limited to Australia, Canada, the United States, and the UK, have even launched national Protective DNS solutions, either for government or private use. With the continued expansion of the attack surface, as well as a rise in network-connected devices that are unable to run traditional endpoint security solutions, Protective DNS solutions are increasingly critical.

What to Consider When **Selecting a Protective DNS Solution**

Selecting a Protective DNS solution is like finding a new car – there are many different options, but you need to find the one that addresses your specific needs and requirements. There are many different considerations that should go into your decision, including but not limited to:

- Efficacy
- The Data Used to Make Decisions
- Deployment Architecture
- Available Integrations
- Configurability
- Scalability
- Performance
- Support for Key Standards: DoH, DoT, DNSSec
- Testimonials

Read on to learn more about each of these topics.

Efficacy

When thinking about the efficacy of the solution, you need to consider both its ability to properly identify malicious domains and infrastructure as well as the solution's false positive rate. Most importantly, you need to look to an independent, third-party assessor for a true impartial assessment of the solution, such as AV-TEST (avtest.org). Without independent validation of the efficacy, you may find widely differing results in real production use.

It may seem obvious, but ideally you want to select the solution with the highest possible level of efficacy combined with the lowest possible false-positive rate. The lower the level of efficacy, the more threats that will go unnoticed in your environment – a good solution should be rated at or above 80% effective in third-party testing. Similarly, you want to select a solution with the lowest possible false-positive rate. The higher the false-positive rate, the more that you will accidentally interrupt normal operations and need to involve IT administrators – a good solution should be rated at a 2% or less false-positive level.

The Data Used to Make Decisions

Blocking or non-blocking decisions are made by a Protective DNS solution using the data it has available at that point in time; when selecting a Protective DNS solution, it is important to look at not just the efficacy level but understand the origin, ownership, and longevity of the data used by the solution to make its decision.

A solution could have a high level of efficacy, but if the provider doesn't control a large percentage of its data, then you need to consider what happens to the level of efficacy should that data source disappear or no longer be accessible.

Second, lots of different feeds of "bad domains" exist on the Internet, but solutions that depend on these feeds (i) will all clearly make the same decision and (ii) generally only are able to make decisions after the feed is distributed and integrated into the solution, making these solutions highly dependent on third parties outside of their control for new information.

Ideally a provider should have a significant amount of data that is exclusive or unique, meaning that it is able to use that data and make decisions that others cannot, as well as being able to "control its own destiny" and not have to worry about the impact of a data source going away or otherwise becoming unavailable.

Deployment Architecture

A traditional Protective DNS deployment model is as an external DNS recursor for an enterprise or organization. In this architecture, the ideal Protective DNS solution can either be the enforcement agent (allow or block the resolution of specific domains) or operate in a passive "pass-through" mode and simply deliver alerts to another component, either for visibility, testing, or to enable some other component in the security stack to be the enforcement agent.

When deploying a Protective DNS solution that blocks traffic, it is critical that a device's DNS request be routed to the appropriate Protective DNS recursor regardless of the device's location. Depending on your specific architecture and configuration, you may want a Protective DNS solution that can integrate with your existing endpoint detection and response (EDR) solution, ensuring that the Protective DNS solution receives visibility of the DNS traffic from the EDR, applies its logic and analysis, and feeds verdicts and information back into the EDR.

In this model, the EDR solution itself would be the enforcement agent. Alternatively, you may want a Protective DNS solution that includes its own lightweight agent that ensures all DNS requests are properly routed to the Protective DNS solution for resolution, regardless of the device's physical location. In this model, the Protective DNS solution would be responsible for security enforcement. When considering this latter architecture, you should make sure that the provider's agent will operate on all of your device operating systems, including laptop/desktop as well as mobile devices.

Alternatively, depending on your use case, and other architectural considerations, you may want the ability to deploy the Protective DNS solution in an entirely passive manner – for example, when protecting a production network. In this way, the protective shield can still be extended around the network but without any potential or measured impact on service availability, latency, or performance.

In general, it is recommended to select a Protective DNS provider that offers maximum flexibility of deployment architecture, because you may not know how your needs and requirements will change. This also helps ensure that you select a "future-proof" solution that can continue to be utilized regardless of other changes in the network or security stack.

Additionally, you need to think about your organization's and your country's data retention policies and laws. You may require a Protective DNS solution that provides enough deployment flexibility to ensure that personally identifiable information (PII) is only stored in-country and does not cross-national borders. Not all Protective DNS architectures support this key capability, so make sure to understand what your requirements are with regard to data policies before selecting a solution.

In general, it makes sense to always select a solution that provides the ability to keep confidential information in-country – as more and more countries implement data retention policies, you don't want to find yourself selecting a solution that can't implement this requirement should your country implement a new policy nationally.

Available Integrations

Generally, solutions only provide maximum benefit when they can be integrated together. Given that each organization may make different decisions around not just the security architecture but the selection of various components, it is important to select a Protective DNS solution that not only supports integrations into your key components today but provides a general-purpose API to enable necessary integrations in the future.

The API should be fully documented to allow you to build your own proprietary integrations as well – ask your proposed Protective DNS provider for a copy of their API documentation and ensure it meets your internal teams' needs.

Configurability

While some risks are common across all organizations (e.g., everyone most likely wants to avoid being infected with Emotet), not every organization thinks about risk in the same way or even tries to implement the same level of risk. It is important to select a Protective DNS solution that provides the flexibility and configurability that you require so you can adapt and customize it to your specific needs.

For instance, some organizations may want to (hypothetically) block all traffic from Russia (.ru). Others may want to block based on specific nameservers or other aspects of DNS infrastructure. Some may allow movie-streaming (and thus access) to pirated websites on their network; others may see this as a policy violation.

The ability to configure one's specific level of risk is vital to making a Protective DNS solution work properly for you – make sure that your Protective DNS solution of choice has a fully flexible and configurable policy manager, in addition to the typical ability to specify course-grained custom allow-and-deny lists.

Scalability

Your usage needs (e.g., the number of DNS queries processed in any given period) may change over time, sometimes rapidly based on changes in your business. A good Protective DNS solution can quickly and automatically scale up to handle both traffic spikes as well as long-term increases in the amount of DNS queries being processed.

Ask your proposed Protective DNS provider about their architecture and how it scales with traffic increases to ensure that you don't ever have an outage due to unintended or unforeseen traffic activity.

Performance

Similar to the sheer number of DNS queries being processed and ensuring that the solution can scale as that number grows, it is equally important to understand the performance of the solution in terms of latency. That is, how much time does it take for the solution to resolve a domain and make a block-or-not decision, and how does latency potentially change with differing traffic patterns?

While small latency differences are probably unnoticeable to the average individual browsing web pages (as generally the time to browse a web page is dominated by the content fetch time not the time to resolve the domain name itself), latency can sometimes be noticeable especially for automated systems and processes.

While not the most important consideration in selecting a Protective DNS solution – and you definitely do not necessarily need the solution with the absolute lowest latency measurement – it is important to understand what the latency measurement is and whether that will be viable in your environment and use case.

Support for Key Standards

Make sure that any Protective DNS solution you select implements DNS over TLS (DoT) and DNS over HTTPS (DoH). You never know when you may need to implement and/or control these protocols inside of your network, and you never want to be caught not being able to implement support for a known standard.

Similarly, DNS by itself is not secure. It was designed in the 1980s and security was not a forefront design principle. Realizing this, the Internet Engineering Task Force (IETF) organization responsible for the DNS protocol standards began working on a solution in the 1990s and the result was the DNSSEC Security Extensions (DNSSEC).

DNSSEC strengthens authentication in DNS using digital signatures based on public key cryptography. With DNSSEC, it is not DNS queries and responses themselves that are cryptographically signed, but rather DNS data itself is signed by the owner of the data. Make sure that any Protective DNS solution you implement supports DNSSEC.

Testimonials

Even though most organizations don't publicly announce what security solutions they use in their environment, every solution provider typically has some level of client testimonials. However, not all testimonials are created equal. Look at the sources of these testimonials and consider whether they apply to your use cases and environment or not – do they come from a known brand, for instance.

Testimonials are best when they come from an independent, unbiased source, such as, a government agency. A good example of a government testimonial on Protective DNS is CISA's memo on selecting a Protective DNS solution. In it, they do not recommend one solution versus another but if your proposed solution isn't even mentioned by CISA, you should at least ask the question why and what that might imply.

Summary

There are a lot of different Protective DNS providers and solutions, and we hope this guide will help provide valuable information to pick the right one for your organization. If you would like to speak with a technical expert for impartial advice, or just answer your questions, please reach out to our security consultation specialists.

CONTACT US
hyas.com/contact