infoblox

WHITEPAPER

DNS BEST PRACTICES

Author:
Craig Sanderson,
Principal Cyber Security Strategist
at Infoblox

TABLE OF CONTENT

This White Paper seeks to provide Domain Name System (DNS) security best practices that accomplish two objectives: to secure the DNS protocol and infrastructure to mitigate misuse or misconfiguration and to utilize DNS to provide an additional layer of network security as part of a zero-trust and/or defense-in-depth security risk management approach.

1. INTRODUCTION

To access Internet resources by domain names (e.g., www.example.com) rather than these IP addresses (93.184.216.34), users need a system that translates these domain names to IP addresses and back. This translation is the primary task of the Domain Name System (DNS). In a modern organization's network deployments, DNS infrastructure is mission-critical to all network communication internally or to the Internet. If it fails, entire networks, along with their applications and users, can be brought down. Thus, DNS is a critical element in an organization's digital resiliency and should be regularly assessed or re-evaluated.

Such assessments should include an evaluation of whether the DNS servers have sufficient capacity to meet the needs of the organization, especially when adopting encrypted DNS protocols. Assessments should also include a review of the health and configuration of the servers to ensure the DNS infrastructure is optimized as part of a highly resilient and redundant architecture. In addition, organizations should implement controls within the DNS platform to leverage the enormous potential of DNS as a foundational layer of a cyber-security architecture.

The DNS platform is already in use by all types of clients on the network, including on-premises, in the cloud and on all IoT devices. Thus, any protection provided by DNS infrastructure benefits all clients that use that infrastructure for name resolution, regardless of the type of device. It is for this reason that organizations and governments should adopt DNS as a foundational component in their cyber security strategy.

2. IMPACT OF DNS ON CYBER RESILIENCY, DEFENSE-IN-DEPTH, AND ZERO TRUST

Recent developments in network security best practices have driven an increased focus on the concept of "defense-in-depth," the idea that no defensive measure is infallible and that, therefore the best defense comes from multiple layers of protection. This style of cyber defense yields a more flexible, scalable, and resilient system that is more resistant to compromise and is more closely aligned to zero-trust principles.

Zero Trust principles assume that no element, node, or service can be implicitly trusted. However, DNS often remains an overlooked element within the Zero Trust strategies, creating a significant gap when it is implicitly trusted.

To align DNS with Zero Trust principles, organizations should not assume but should verify their information sources by leveraging DNS to enforce security policies, prevent end-users and systems from accessing malicious or unauthorized resources, and provide asset visibility for the purpose of digital forensics and incident response.

This marks a shift in the role of DNS from a purely operational one to a wider-scale tool to provide Internet security and bolster network resilience.

Zero trust presents a shift from a location-centric model to an identity, context, and data-centric approach with fine-grained security controls between users, systems, applications, data, and assets that change over time. As DNS resolution may be open to endpoints querying any public domain names, organizations should secure DNS to prevent it from being used as a route to bypass zero-trust network controls. This could involve the implementation of protective DNS services and/or specific DNS resolver configurations for certain classes of endpoints. DNS security is particularly useful in protecting the security of IoT devices and operational technology present in manufacturing and critical infrastructure, which have been challenging to deploy end-point security solutions.

3. DNS THREAT VECTORS

DNS infrastructure is a common threat vector for attack campaigns.[1] Threats to the DNS service can cascade into significant operational failures or loss of data, integrity, and confidentiality. In addition, many threat actors, such as ALPHV Blackcat, utilize malicious domains to launch ransom attacks at scale.[2] Below is a non-exhaustive list of examples of compromises that could jeopardize not only the DNS services but also all systems that organizations rely on:

a) **Distributed Denial of Service:** A malicious attacker could send a large volume of queries to perform a denial of service (DoS) attack against a DNS server/service. The attacker could use numerous third-party DNS clients to aid in the attack (distributed DoS or DDoS).

b) **Unauthorized Configuration Change:** The platform-level configuration file that enables DNS communication could be corrupted or subject to unauthorized modifications due to inadequate protections, resulting in disruptions varying from the breakdown of communication among DNS hosts to the complete failure of the DNS service itself.

c) **Dangling CNAME Exploitations:** When a DNS CNAME record links two domain names together, there is the risk that the parent domain of the canonical name the record points to does not remain registered by the target organization. As a result, threat actors can register the parent domain and thus, DNS resolutions will now ultimately resolve to the threat actor's-controlled domain. Another possible way CNAME records can be exploited is if the canonical name resolves to an IP address that is no longer in use by the domain owner and the attacker can gain control of that IP address, then that could be leveraged to conduct attacks.

d) **Lame Delegation Exploitation:** Domains are delegated from one level of the DNS hierarchy to another, most often from a top-level domain to an organization registering the domain. When delegation is set up, it is necessary to specify the authoritative DNS servers for the domain. Delegating a zone to any nameserver that is not authoritative for that zone is known as Lame Delegation. Lame delegation can result in domain hijacking in some circumstances. When a subdomain is delegated to a DNS hosting provider and the contract for providing DNS services for that domain lapses, threat actors are able to hijack resolution for that subdomain by contracting with the provider that controls the servers targeted by the delegation to host that subdomain under their control. This then enables the threat actor to redirect resolution requests to their own infrastructure, benefiting from the positive reputation of the domain name.

e) **Lookalike Domain Exploitation:** Threat actors extensively leverage lookalike or typosquat domains to impersonate target organizations. By leveraging the positive reputation of legitimate organizations, threat actors vastly increase the success rate of their phishing and malware campaigns. These lookalike domains can include subtle variations of legitimate domains or can use text or character substitution to register a domain that a user would likely believe to be owned by the legitimate organization.

f) **DGAs and Ransomware-as-a-Service:** Domain Generation Algorithms (DGAs) enable malicious actors to generate numerous domain names for command-and-control (C2) communication, facilitating malware distribution and evasion of detection. Ransomware-as-a-service (RaaS) such as Blackcat use DNS for C2 communication.

4. DNS SECURITY BEST PRACTICES

The original intent of DNS was to distribute information such as host and IP address mappings, mail routing information, and more, so it has not traditionally been viewed as a tool for securing network communications. However, the role DNS plays in enabling nearly all network

[1] https://www.cisa.gov/news-events/news/cisa-launches-its-protective-dns-resolver-general-availability-federal-agencies

[2] https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-353a
https://blogs.infoblox.com/threat-intelligence/dns-early-detection-breaking-the-blackcat-ransomware-kill-chain/

communications today makes it an effective tool for not only monitoring but also managing those communications. That is why DNS security has evolved from protecting the DNS infrastructure and protocols alone to becoming a critical security control and a key component of the security strategy of the whole organization. It should be noted that "DNSSEC" (which stands for a specific protocol called DNS Security Extension) is only one part of the wider idea of "DNS Security" and must be implemented alongside other best practices.

In the following sections, we discuss the three key components of DNS security best practices that should be considered, including:

- Employing Protective DNS
- Protecting the DNS Protocol
- Protecting the DNS Service and Infrastructure

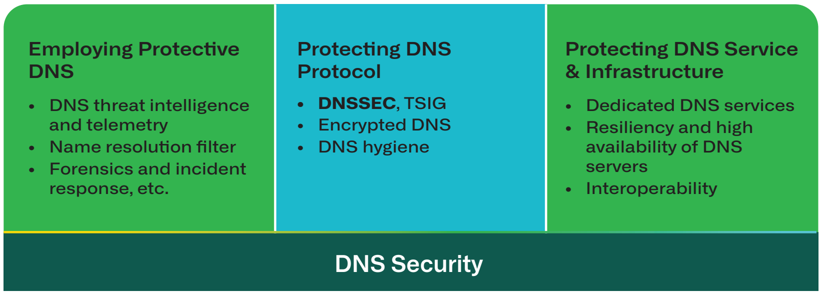

Employing Protective DNS

- DNS threat intelligence and telemetry
- Name resolution filter
- Forensics and incident response, etc.

Protecting DNS Protocol

- **DNSSEC**, TSIG
- Encrypted DNS
- DNS hygiene

Protecting DNS Service & Infrastructure

- Dedicated DNS services
- Resiliency and high availability of DNS servers
- Interoperability

DNS Security

4.1 Protective DNS

Protective DNS is a DNS service that is enhanced with security capabilities to analyze DNS queries and responses and take action to mitigate threats. Protective DNS prevents the delivery of malware, ransomware, phishing and other web-link-centric attacks that attempt to deliver spyware and viruses as well as blocks access to malicious websites. Protective DNS can be provided as a service from a vendor, deployed on internal DNS infrastructure, or a combination of the two. There are potential benefits to using a combination of externally provided Protective DNS with internally deployed Protective DNS. While this approach may not be applicable in all cases, we recommend that this combined hybrid scheme is utilized where feasible. As evidenced by the deployment of Protective DNS platforms by governments around the world and the DNS4EU initiative in the European Union,[3] the use of Protective DNS has become a DNS best practice.

The outcomes of employing Protective DNS should include:

Blocking or redirecting harmful traffic in real time before malicious activity starts.

Blocking categories of traffic. This is typically traffic that doesn't conform to an organizations policies.

Visibility into real-time DNS query and response history for digital forensics.

Integration with an organization's wider security ecosystem.

Facilitating an organizations responsibility to comply with regulatory or contractual requirements for blocking traffic to disallowed sites.

[3] https://www.joindns4.eu/
https://www.ncsc.gov.uk/information/pdns
https://www.forgov.qld.gov.au/information-and-communication-technology/cyber-security/cyber-security-services/system-and-email-hardening/protective-dns-service/implement-a-protective-dns-service
https://www.cisa.gov/resources-tools/services/protective-domain-name-system-resolver

a. Threat Intelligence and Telemetry

The attack surface of organizations is continually evolving, so being able to monitor and disrupt malicious communications as early as possible is of paramount importance. DNS is one constant even as other networking technologies emerge and is well positioned to protect users in all environments, from organizations to mobile endpoints. In addition, DNS is often the earliest connection to threats and can stop complex multi-stage attacks before they progress. DNS, as a security control point, is not limited to any single type of threat, unlike other mechanisms in the security stack. It can protect users and organizations from scams, credential theft, ransomware, and data exfiltration.

This approach requires threat intelligence and the ability to integrate that into the DNS resolver. Threat intelligence is leveraged in a DNS infrastructure via mechanisms such as Response Policy Zones, or RPZs, and can be seamlessly integrated into the DNS resolution chain via a number of architectures.

b. Name Resolution Filtering

Name resolution filtering is a term used to refer to DNS infrastructure that applies policy to DNS resolution. These policies are typically related to security: for example, a Protective DNS implementation might refuse to resolve a set of domain names known to be used in phishing campaigns or that identify malware command-and-control infrastructure. Instead of resolving these domain names to IP addresses or other types of data, Protective DNS generally returns some other form of DNS response, such as NXDOMAIN (which means "Non-existent Domain"), indicating that the domain name being looked up doesn't exist. Protective DNS implementations can also log queries for domain names that trigger policy because those queries may indicate infection by malware or other malicious activity.

Protective DNS is generally implemented by using either RPZs configured on an on-premises DNS service, a cloud-based, secure recursive DNS service, or some combination of the two.

Secure recursive DNS services are also available on the Internet. The services generally offer a web-based control panel that allows administrators to customize the resolution policy applied to queries from the organization's clients. The cloud approach offers greater scalability, storage and computing power, but has disadvantages, such as a loss of confidentiality, higher latency, and challenges in quickly and accurately attributing DNS queries to their source. Organizations may consider using a combination of on-premises and cloud-based secure DNS services in tandem to secure the benefits of each. The best choice for a given DNS deployment varies depending on the specific needs of the network and its users.

c. DNS for Digital Forensics and Incident Response

Government agencies and regulated enterprises should implement robust DNS traffic logging mechanisms to meet compliance requirements. Logging should capture both current and historical DNS traffic to enable digital forensics and incident response.
These DNS logs should be integrated with other system logs to facilitate correlation with other logs of network activity, enhancing visibility and auditability.

In cases where full DNS traffic logging is determined to be too resource-intensive, organizations may consider using cloud-based solutions, efficient logging methods
(e.g., DNSTAP), or selective logging. However, it is imperative that DNS queries and responses associated with domains classified as malicious or unauthorized by Protective DNS services are always logged to support security and compliance objectives. To ensure rapid notification of queries that might indicate infection or malicious activity, organizations should integrate Protective DNS logs from their name servers or their secure recursive DNS service with their SIEM or log analysis platform.

4.2 Protecting the DNS Protocol

The DNS protocol is essential for locating services to enable web browsing, email and other mission-critical applications. As a result, traditional security platforms, such as next-generation firewalls, often pass DNS traffic unhindered and uninspected. Threat actors have increasingly turned to DNS as an exfiltration vector. Per the U.S. Cybersecurity and Infrastructure Security Agency (CISA), "DNS infrastructure is a common threat vector for attack campaigns."[4] Threat actors often embed stolen data in DNS packets, relying on the DNS infrastructure to relay the stolen data to the threat actor-controlled DNS servers. DNS platforms are ideally positioned to evaluate recursive DNS requests they receive for attempts to exfiltrate data.

a. Protecting the Integrity of DNS Services

Protecting the integrity of the DNS protocol is critical to ensuring the security and reliability of internet communications. DNS Security Extensions (DNSSEC) play a crucial role by using cryptographic signatures to authenticate DNS responses, preventing attacks like cache poisoning and man-in-the-middle interceptions. Additionally, Transaction Signature (TSIG) enhances DNS security by enabling mutual authentication and integrity verification between DNS servers, using shared secret keys to safeguard zone transfers and dynamic updates from unauthorized changes. Implementing these mechanisms, alongside best practices like rate limiting, monitoring for anomalies, and enforcing strict access controls, helps maintain the trust and resilience of DNS infrastructure against evolving cyber threats. In addition, organizations should ensure that users do not deliberately or inadvertently use unauthorized public, Internet-based DNS services.

b. Usage of Encrypted DNS and Authentications to Protect the Protocol

Communication between stub resolvers and the recursive DNS servers they query has traditionally been unencrypted. The DNS messages exchanged by stub resolvers and DNS servers have a binary encoding, but that encoding is widely understood and easily decoded. This communication has, therefore, been subject to both interception and spoofing, which can reveal sensitive information or allow an attacker to redirect unsuspecting users to malicious sites.

To address these threats, the Internet Engineering Task Force (IETF) has developed several enhancements to DNS, commonly collectively known as Encrypted DNS. All these protocols encrypt communications between stub resolvers and recursive DNS servers and optionally allow recursive DNS servers to authenticate themselves to stub resolvers, addressing the threats of interception and spoofing.

c. DNS Hygiene and Best Practices

Exploitation of misconfiguration and lapsed domain/DNS resolver registration is relatively simple for threat actors to execute and can result in serious compromise of DNS integrity.

Threat actors have proven that attacks such as phishing are far more likely to succeed if they are linked to domains owned by trusted organizations. As a result, they often register lookalike domains that look similar to but are not owned by the target organization. More concerning, poor authoritative domain hygiene can allow threat actors to take control of domains owned by a trusted organization.

4.3 Protecting the DNS Service and Infrastructure

DNS appliances, like other network appliances, are purpose-built and optimized for ease of management, security, and performance. General-purpose servers cannot match the tuning that these appliances offer and, as such, expose organizations to significant risks, given the criticality of DNS as a network service.

[4] https://www.cisa.gov/news-events/news/cisa-launches-its-protective-dns-resolver-general-availability-federal-agencies

Organizations commonly use Windows servers that host DNS and DHCP alongside Active Directory (AD), another mission-critical identity infrastructure. AD manages user access and permissions and stores critical information, creating a large attack surface that is vulnerable to cyberattacks.

Without this separation of duties, if an attacker follows a common escalation path and targets Active Directory, the DNS service and the network it relies on are at risk. As DNS takes on a more significant role in an organization's cyber security strategy, a dedicated DNS Server will be essential to mitigate these risks.

Similarly, threat actors often seek to take advantage of vulnerabilities discovered in DNS that can impact the integrity of the DNS system or result in denial of service.

For further guidance on secure DNS, refer to the most recent version of NIST SP 800-81: Secure Domain Name System (DNS) Deployment Guide.

5. SUMMARY & RECOMMENDATIONS

This White Paper discusses the critical importance of DNS infrastructure in maintaining an organization's digital resiliency and cyber security. It highlights the need for regular assessments of DNS servers to ensure they have sufficient capacity and are configured optimally. It also emphasizes the role of DNS as a foundational layer in cybersecurity architecture, benefiting all clients using the infrastructure.

To address the evolving DNS security threats, we make the following high-level recommendations to network and security owners:

- Employ Protective DNS, wherever technically feasible, to provide additional network-wide security capabilities that include:
  - Blocking of harmful or malicious traffic in real time,
  - Filtering out categories of traffic that do not conform to an organization's policies,
  - Real-time and historical DNS query and response data to facilitate digital forensics and incident response.
  - Integration with the wider security ecosystem as part of a defense in depth or zero trust approach, and
  - Facilitation of an organization's responsibility to comply with regulatory or contractual requirements for blocking traffic to disallowed sites (copyright violations, legal restrictions, etc.).
- Encrypt DNS traffic, both internal and external, wherever feasible.
- Deploy dedicated DNS servers to reduce the attack surface.
- Follow all technical guidance on ensuring your DNS deployments and the DNS protocol are as secure and resilient as possible.

For Additional Information, please contact the Infoblox team at ga@infoblox.com.

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: 20250131v3