Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 25 July 2026

H. Duan
Tsinghua University
M. Liu
Zhongguancun Laboratory
B. Liu
Tsinghua University
C. Lu
Zhongguancun Laboratory

21 January 2026

Considerations for Protective DNS Server Operators
draft-liu-dnsop-protective-dns-02

Abstract

Protective DNS is a defense mechanism deployed on recursive resolvers to prevent users from accessing malicious domains. For domain names in the blocklist, it rewrites DNS resolution responses to point to secure destinations (e.g., safe servers) to prevent users from accessing malicious entities.

Owing to its effective defenses against common cyber attack behaviors— such as command-and-control (C2) communications of malware— Protective DNS deployment has surged via various initiatives. Not only have renowned DNS service providers adopted this defense, but some countries have also launched national-scale deployments. Meanwhile, studies analyzing Protective DNS have identified implementation diversity.

Thus, this document aims to provide specific operational and security considerations for Protective DNS. It is intended primarily for entities seeking to deploy Protective DNS for defensive purposes, offering deployment and security considerations.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://MingxuanLiu.github.io/Protective-DNS-Draft/draft-ietf-dnsop-protective-dns.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-liu-dnsop-protective-dns/.

Source for this draft and an issue tracker can be found at https://github.com/MingxuanLiu/Protective-DNS-Draft.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Protective DNS (also termed as PDNS) is a lightweight defensive measure deployed at recursive resolvers to proactively rewrite DNS resolution responses for malicious domains, thereby preventing users from accessing malicious resources. Specifically, when a client initiates a domain name resolution request, PDNS first performs a security check on the target domain—determining whether the domain poses a security risk by matching it against blocklists. If the domain is identified as malicious, PDNS uses DNS rewriting technology to intercept the resolution request and return a secure response (e.g., a safe server address controlled by the PDNS service provider), blocking users from establishing connections with malicious resources. For domains outside blocklists, PDNS will perform normal recursive queries and return authentic DNS responses.

The defensive benefits offered by PDNS have spurred extensive deployment efforts. Large DNS resolution service providers have increasingly deployed PDNS on their recursive servers. Moreover, some countries or regions have initiated national-level infrastructure deployments of PDNS. The UK's National Cyber Security Centre (NCSC), for instance, launched a PDNS service in 2017 to enhance cyber defenses, which has since been used by central government departments, local authorities, schools, and emergency services [UK-NCSC-PDNS]. Notably, statistics from the UK NCSC indicate that approximately 7.2 million individual users utilized the system in 2023 [UK-NCSC-PDNS-Usage]. In 2022, PDNS in UK processed 810 billion DNS queries and blocked 11 billion queries involving 420,000 domains, accounting for approximately 2% of all queries [UK-Defence].

As the deployment of PDNS continues to grow, efforts have been made to systematically analyze PDNS services, including deployment status, operational mechanism and security implications. A recent work has revealed that 9% of open recursive resolvers exhibit PDNS behaviors [NDSS24], and finds significant discrepancies among providers in terms of rewriting policies, blocklist selection and performance. Particularly, security risks such as implementation flaws, over-blocking, and dangling resources have been identified.

This document is primarily intended for readers familiar with Protective DNS technology and somewhat aware of the potential impacts that deploying such technologies may entail. Moreover, existing documents [USENIX24], [SAC127] are recommended for legal considerations. On this basis, this document focuses on discussing technical considerations at the deployment level.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Background

3.1. Overview of Protective DNS

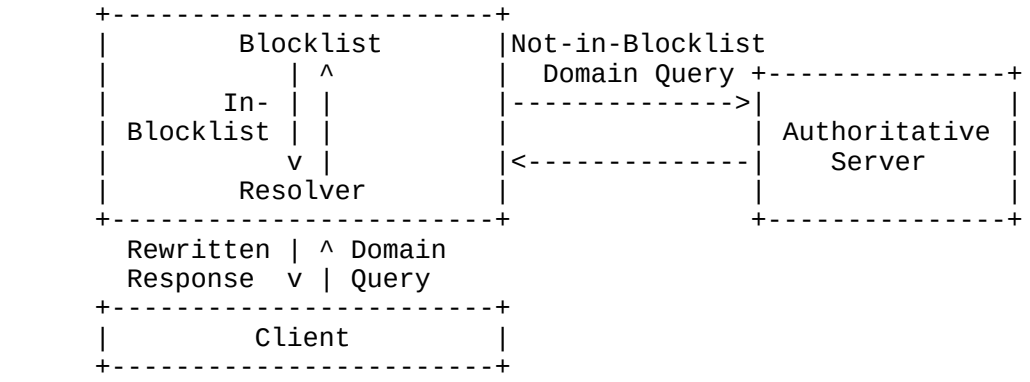

Figure 1: The workflow of Protective DNS.

Figure 1 shows the workflow of Protective DNS (PDNS). Protective DNS is deployed on a recursive resolver. When the PDNS resolver receives a DNS query for a domain name, it first matches the domain against its maintained blocklist. The resolver then makes a decision based

on the blocklist lookup result. If the domain is found in the blocklist, PDNS rewrites the DNS response to resolve the query to a "secure" result (e.g., special-purpose addresses), effectively preventing the client from accessing the corresponding malicious resource. Conversely, if the domain is not in the blocklist, the resolver returns a normal response by querying authoritative servers or using local cache results to respond to the client [RFC1034], [RFC1035]. Thus, the two functional components that underpin the critical role of PDNS are the Blocklist and the Rewriting Policy.

Blocklist. PDNS determines whether to rewrite the resolution result of a domain name based on its presence in the blocklist. Blocklist sources include multiple aspects: commercial threat intelligence (TI), open-source TI, vendor-maintained domain blocklists, and user complaints. The types of malicious domains included in blocklists vary by vendor definition, encompassing but not limited to: malware, botnet command-and-control (C2), phishing, fraud, and adult content. Additionally, PDNS deployments implement blocklist lookup in two primary forms:

1. Local Lookup: storing directly on the PDNS server, allowing direct queries against the local blocklist file.

2. Remote Lookup: performing lookups via network interfaces (e.g., DNSBL [RFC5782]).

Rewriting Policy. Upon retrieving blocklist matching results, the PDNS server rewrites resolution responses for domains in the blocklist. rewriting strategies exist in multiple forms. Based on empirical analysis of leading Protective DNS vendors, this document summarizes five specific rewriting policies.

1) Using the secure IP addresses in A record controlled by the provider:

```
malicious_domain.com        A 10 controled_IP;
```

2) Using IP addresses with special purposes, such as the reserved address like 0.0.0.0, link local address like 192.168.0.1, loopback addresses like 127.0.0.1, and so on:

```
malicious_domain.com        A 10 127.0.0.1;
```

3) Utilizing the CNAME record to rewrite the request to the domain name controlled by the provider:

```
 malicious_domain.com    CNAME    10    controled_domain.com;
```

4) Using an empty Answer field in the response to prevent users from accessing malicious resources:

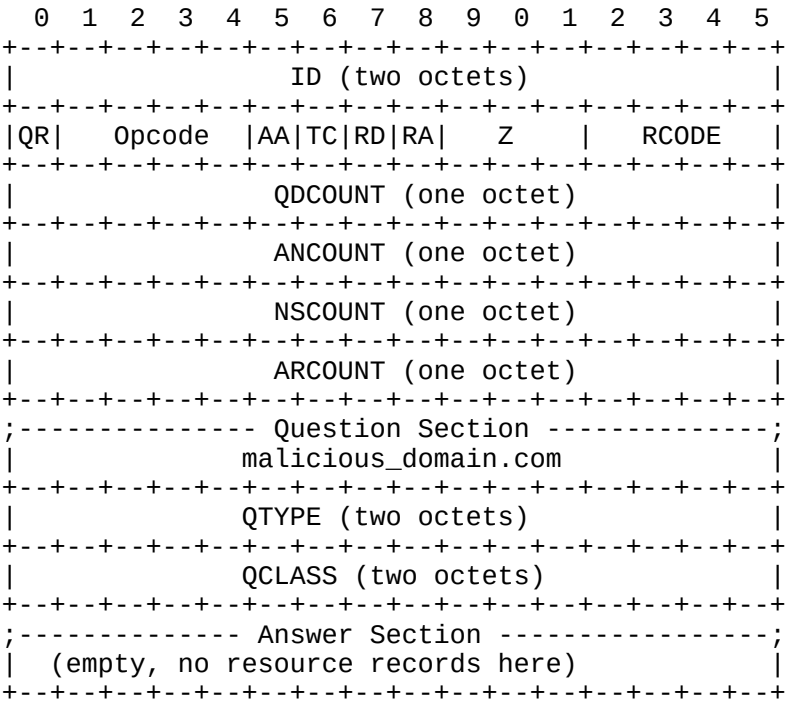

Figure 2: Format of empty Answer fileld in Protective DNS.

5) Using special response codes for the reply, such as NXDomain, ServerFail, etc:

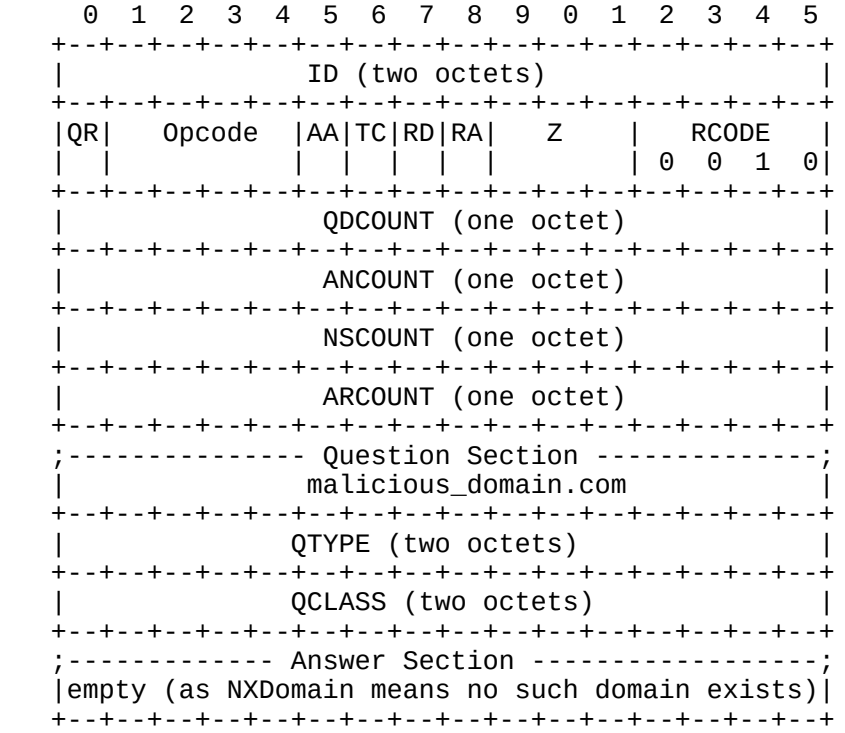

Figure 3: Format of NXDomain Response in Protective DNS.

Additionally, primary implementations of the rewriting module include, but are not limited to:

1. DNS Response Policy Zones (RPZ) [I-D.ietf-dnsop-dns-rpz]: Implemented as zone files [RFC1034], [RFC1035], specifying both whether rewriting is required and providing domain-specific rewrite results.

2. Domain Lists: This format consists of one domain per line, specifying only the rewrite requirement for each domain.

3.2. Relationship between Other Technologies

Relation ship bewteen PDNS and DNS Blockling. Protective DNS constitutes a subset of DNS Blocking (generally considered as synonymous with DNS Filtering [SAC127]), sharing the goal of security defense by leveraging known blocklists (e.g., domain blocklists) to prevent user exposure to malicious resources. DNS Blocking is a broader security concept encompassing any rewriting operations of DNS resolution traffic to restrict access to specific domains. Its

applicability spans diverse scenarios and architectural layers, specifically including defense at any DNS resolution role (stub in client, recursive resolver, authoritative server) and other security use scenarios (e.g., spam filtering, gateway firewall protection). In contrast, PDNS represents a specific implementation of DNS Blocking, typically deployed on recursive resolvers.

Several prior documents define specific techniques within DNS Blocking:

1. DNS Blocklist (DNSBL)[RFC5782]: DNSBL represents a DNS-based blocklist lookup technology, serving as a concrete implementation of blocklist lookup within the broader DNS Blocking paradigm (including PDNS). First, it maintains blocklists of IP addresses or domain names associated with malicious activities. Then, it utilizes DNS queries to determine whether to block related traffic.

2. DNS Response Policy Zones (RPZ) [I-D.ietf-dnsop-dns-rpz]: Beyond remote blocklist querying (e.g., DNSBL), local deployment offers another blocklist deployment strategy. RPZ achieves localized blocklist query via a zone file containing rewrite instructions for each malicious domain, often utilized in scenarios such as DNS Firewall. RPZ supports multiple triggering mechanisms. Among these, QNAME triggering is the primary mode adopted by PDNS.

Relationship between PDNS and Censorship. While censorship shares the objective with DNS blocking, i.e., preventing end users from accessing specific resources, its blocking strategies are far more diverse. Domain blocklist-based blocking represents just one type of DNS-layer mitigation, alongside techniques such as DNS poisoning. Cross-layer mechanisms at network layers include TCP reset (e.g., sending forged RST packets), IP blocking, HTTPS man-in-the-middle attacks, and Deep Packet Inspection (DPI) for payload analysis.

Relationship between PDNS and Other Domain Protection. Defense mechanisms at the domain level are diverse, with protective actions occurring across various roles in the domain ecosystem. During the initial domain registration process, registries and registrars can use domain seizure to remove malicious domains from registration data, thereby preventing their continued harm to users. Registries and registrars can also employ sinkhole technology to redirect some malicious domains to domain black holes. The Protective DNS defense mechanism addressed in this document is primarily deployed on recursive resolvers.

4. Deployment Status

Due to the significant protective efficacy of Protective DNS, existing academic efforts have evaluated its deployment status through measurement. The results indicate a growing adoption trend of PDNS across large-scale DNS service providers, operational recursive resolvers in the wild, and national-level deployments.

National Deployment. First, several countries and regions have even deployed national Protective DNS projects, including:

1. DNS4EU [DNS4EU]: Launched in January 2023 under the auspices of the European Union Agency for Cybersecurity (ENISA), this initiative serves as an alternative to prevailing public DNS resolution services. It is designed to deliver protective, privacy-compliant, and resilient DNS capabilities, thereby enhancing the EU's digital sovereignty and security. Any device connected to DNS4EU resolvers that attempts to access a malicious domain (e.g., hosting malware or related to phishing content) is immediately blocked, preventing potential harm. Leveraging real-time threat monitoring and defense mechanisms, the project ensures that malicious domains identified in one EU jurisdiction are countered across multiple member states to contain their spread. DNS4EU provides DNS security assurances to critical sectors including EU citizens, public institutions, government entities, and operators, while also supporting voluntary adoption or opt-out by EU residents. Recently, DNS4EU has also opened public query interfaces accessible to the citizens of the European Union [DNS4EU-Public].

2. PDNS in UK [UK-Defence]: Supported by the UK's National Cyber Security Centre (NCSC), this initiative advises private enterprises and government agencies to adopt Protective DNS to safeguard their IT assets and network security. By blocking access to known malicious domains, it significantly reduces the effectiveness of ransomware, phishing, botnet, and malware attacks. According to NCSC statistics, in 2022, the project's PDNS service processed 810 billion DNS queries and blocked 11 billion queries involving 420 thousand distinct domains [UK-Defence].

3. PDNS in US [US-Protect]: This initiative, launched by the Cybersecurity and Infrastructure Security Agency (CISA) in 2022, is designed to provide security safeguards for the United States' national critical infrastructure. The project protects the federal government by blocking network traffic at the DNS resolution layer from reaching potentially malicious destinations, enhancing resilience against intrusions and

attacks. When Protective DNS detects a DNS request matching threat intelligence indicators, the service blocks, redirects, or sinkhole the query response to a secure endpoint, while sending alerts to the source agency and CISA. The initiative mandates the use of this Protective DNS service by all Federal Civil Executive Branch (FCEB) agencies and offers limited availability to infrastructure participants in pilot programs.

Public DNS Provider. Existing work, by combining surveys of publicly available documentation from widely adopted DNS service providers with active testing, confirms that two-thirds of public DNS providers already offer Protective DNS services. Notably, practices vary significantly across vendors, including differences in the types of malicious domains defended against and rewrite policies. In terms of deployment strategies, the majority of vendors employ hybrid deployments, providing both Protective DNS and regular DNS resolution on a single resolver IP address, such as Comodo DNS and OpenDNS. A smaller subset deploys Protective DNS and regular DNS resolution on separate resolver IPs, and in some cases, even different resolver servers from the same vendor may defend against distinct types of malicious domains. For example, Cloudflare operates standard DNS services on 1.1.1.1, while PDNS servers on 1.1.1.2 and 1.1.1.3 implement differentiated protection: 1) 1.1.1.2 focuses on malware defense; 2) 1.1.1.3 defends against both malware and adult content.

Open Resolver. Furthermore, existing studies using active scanning of the IPv4 address space in 2023 confirm that approximately 9% of recursive resolvers have deployed Protective DNS capabilities, covering almost two-thirds of the world’s countries. By analyzing the types of malicious domains defended against, research has found that malware, botnet, phishing, and spam are the most common categories, with observable blocklist overlap across different recursive resolution services. In terms of rewrite policies, using safe IP addresses and specialized IP addresses represents the most prevalent rewriting approach.

5. Operational Considerations

Considering that deployment is the first step in using and even maintaining the security of Protective DNS, in this section, we propose a series of operational considerations that cover multiple aspects of deployment practice, including blocklist selection, rewriting strategy construction, performance impact assessment, and explanatory offering

5.1. Operational Consideration 1: Blocklist Selection

One of the necessary conditions for Protective DNS to achieve its defensive capability is to maintain a blocklist that includes a series of domain names to be blocked. PDNS providers should maintain one or more policy-driven domain lists derived from threat intelligence, regulatory sources, and organizational policy. These may include domains associated with security threats, inappropriate content, privacy risks, or policy violations. Blocking behavior may vary by context, including NXDOMAIN synthesis, redirection, or other response modes aligned with regulatory and operational requirements.

First, when collecting blocklists, PDNS providers should explicitly define the blocked domain types based on the intended use case, taking into account multiple factors such as national defense requirements and policy regulations with telecom and media regulators. Common blocked domain types fall into 6 primary categories.

1) The most prevalent block type are malicious domains, including but not limited to malware, botnet, phishing, spam, and tracking domains.

2) Certain domains are blocked for content control purposes, such as adult sites, gambling platforms, and piracy-related domains. This category is often tied to national policies governing PDNS deployments— for instance, gambling is illegal in some jurisdictions.

3) Some domains are blocked to comply with privacy protection regulations, e.g., those associated with trackers and advertising networks.

4) Domains that are legitimate but inappropriate for children (e.g., pornographic, gambling, and gaming sites) are blocked for educational or parental control purposes.

5) Domains posing potential data leakage risks or offering unauthorized services are blocked in accordance with national policies.

6) In security defense scenarios, domains linked to known suspicious indicators are blocked based on domain resolution correlations. Examples include domains resolved by suspicious or unauthorized upstream DNS servers, or those mapping to IP address ranges hosting malware, proxies, bulletproof hosting, or other undesirable infrastructure.

Second, after identifying the types of domains to block, PDNS providers should construct blocklists from appropriate sources, such as self-collected data derived from traffic analysis and user feedback, open-source threat intelligence feeds, and commercial threat intelligence services. Typically, PDNS providers rely on a combination of multiple threat intelligence sources, including but not limited to government-mandated lists, internal telemetry data, and customer-defined filters. These data collection sources vary widely by region, vertical, and operational goal.

Third, PDNS operators should select an appropriate blocklist deployment approach based on operational context, including device resource constraints and network access patterns. Remote querying approaches (e.g., using DNSBL) necessitate proactive consideration of potential privacy implications and the impacts of network instability. Local deployments (e.g., RPZ-formated local zone file) necessitate a thorough assessment of local resource limitations. Specifically, the Blocklist scale deployed on Protective DNS should be carefully defined based on the system’s processing capability. Blocklist size directly affects both the response efficiency of Protective DNS and hardware resource consumption (e.g., CPU, memory) on the hosting device. Blocking behavior may vary by context, including NXDOMAIN synthesis, redirection, or other response modes aligned with regulatory and operational requirements.

5.2. Operational Consideration 2 - Rewriting Policy Construction

Based on empirical analyses of popular Protective DNS providers, five primary rewriting approaches have been identified in practice, see the Section of Overview of Protective DNS. Each rewriting strategy caters to specific security scenarios, requiring providers to select appropriate approaches based on their application requirements, specifically as follows:

1) Secure IP addresses: Under this policy, the rewritten target address is a server controlled by the PDNS provider, enabling PDNS providers to monitor traffic. Specifically, the controlled server acts as a "honeypot" managed by the provider, capturing DNS traffic of malicious domain names for further analysis of threat behaviors (e.g., malware communications). However, this approach incurs operational overhead for PDNS providers, who need to actively monitor the status of these servers. Additionally, this scenario necessitates consideration of privacy risks arising from traffic monitoring, such as the security concerns of user traffic collected for surveillance purposes.

2) Specialized IP addresses: Due to the non-routable nature of these IP addresses on public networks, they are better suited for scenarios with strict privacy protection requirements, such as when users do not want any third parties to track their network behavior. However, PDNS operators should consider potential risks when using these specialized IP addresses. For example: a) 192.168.0.1 is typically used for local area network devices. Such configurations may lead to unintended access to internal network devices if clients mistakenly connect to them; b) 127.0.0.1 is commonly used for local inter-process communication. Redirecting to this address may cause clients to attempt to connect to local services, which could be exploited if vulnerabilities exist— such as through port scanning or service spoofing. This is particularly risky when users mistake local services for external ones, potentially exposing sensitive information or enabling attacks. Additionally, this approach lacks transparency for PDNS users, as they cannot obtain explicit explanations on resolution results (i.e., resolved IPs) similar to secure IP indicators. However, rewriting information can be provided in DNS resolutions— for example, through explanations in extended fields (i.e., EDE [RFC8914]), with detailed content elaborated in Section of Operational Considerations 4.

3) Secure CNAME: This strategy, similar to using controlled IP addresses, enables providers to dynamically monitor traffic. However, providers should remain vigilant against dangling resource record risks arising from improper management, details of which are discussed in Security Considerations.

4) Empty Answer Section: This strategy represents a minimalist rewriting approach, simply returning an empty answer section. However, it undermines the transparency of PDNS services for users and may escalate to more aggressive implementations— such as refusing to return resolution responses— which entail denial-of-response risks detailed in the Security Considerations.

5) Special Response Codes (Rcodes): This strategy is compatible with regular DNS error scenarios, which help prevent malware from detecting the defensive mechanisms of PDNS. However, such practices may also undermine the transparency of Protective DNS services, making it difficult for users to understand the principles behind rewriting operations and to distinguish between defensive rewrites from PDNS and genuine failures.

Second, PDNS operators should consider the impact of TTL configurations and appropriately configure the TTL values for rewritten records. On one hand, an overly long TTL may lead to delayed updates of defense strategies for malicious domains. On the other hand, a too-short TTL triggers frequent DNS queries, increasing PDNS server load and potentially degrading performance.

5.3. Operational Consideration 3 - Performance Impact

As Protective DNS services introduce an additional query step on whether a domain is malicious during standard DNS resolution, operators should anticipate potential impacts on DNS resolution performance. Specifically, factors such as blocklist deployment method (remote vs. local), scale, and domain matching techniques (e.g., hash matching) can affect performance. Experimental results show that loading a blocklist into memory with five million malicious domains can still be maintained within 10-second response times, but exceeding this scale may result in loading times exceeding 10 seconds. Additionally, experimental studies on Public DNS Providers have found that 70% of blocked domains receive PDNS responses within 0.2 seconds. Meanwhile, for providers offering both PDNS and non-PDNS services, the response time difference typically does not exceed 2 ms. Therefore, when implementing PDNS, operational references from public PDNS provider experiments should be addressed— specifically, the impact of blocklist loading and domain matching on query latency.

5.4. Operational Consideration 4 - Offering Explanation

Protective DNS operates as a complete black-box service for users. Regardless of the rewriting strategy employed, users only perceive the blocking effect— i.e., the inability to access a domain. While providers can refine blocklist quality to minimize false positives, the inevitable presence of false positives significantly degrades user experience. Users may encounter unexplained domain inaccessibility that is indistinguishable from prevalent DNS tampering (e.g., censorship, man-in-the-middle attacks). Therefore, providing explanations for Protective DNS services can enhance user experience and alleviate privacy concerns. Operators should anticipate that omitted explanations may lead users to misperceive service instability or even switch to competing DNS providers.

This explanation of PDNS can be realized through multiple approaches. First, to indicate that blocking originates from Protective DNS, service providers may offer a dedicated landing page to explain their protective services, helping users confirm that the observed DNS rewriting originates from the provider’ s Protective DNS. Empirical analyses show some PDNS providers redirect DNS queries to a secure IP address hosting a page that notifies users of potential malicious

website access. However, using a provider-controlled IP introduces risks of dangling resource takeover, detailed in Security Consideration. Second, the EDE (Extended DNS Error) [RFC8914], [I-D.ietf-dnsop-structured-dns-error] protocol can be employed to specify in the extension fields that rewriting results from Protective DNS defense against malicious domains. Existing work has shown that some PDNS implementations already utilize the EDE field for indication— for example, using "PROHIBITED", as illustrated in Figure 4.

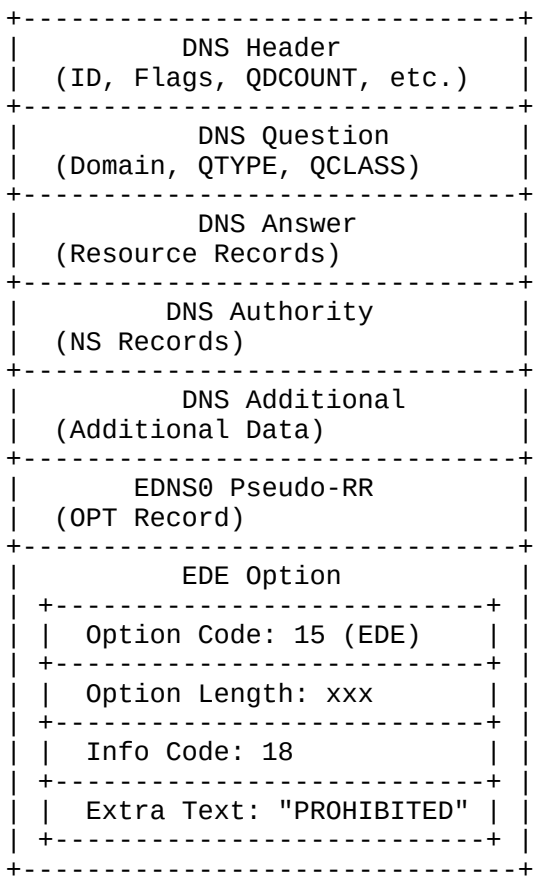

Figure 4: Example of EDE usage in Protective DNS.

Moreover, providing user appeal channels on explanation pages, such as an email address, could mitigate negative impacts of potential false positives.

6. Security Considerations

Furthermore, by integrating the operational considerations, we propose security considerations to enhance the security of Protective DNS on the basis of improving its practicality. We outline specific considerations covering multiple dimensions. For each factor, we provide specific recommended mitigations, including rewriting policy flaws, over-blocking, dangling resource risks, interaction with data integrity protection and fallbacks.

6.1. Security Consideration 1 - Rewriting Policy Flaws

To prevent flaws in the protection function or even bypassing, service providers should consider the following three aspects.

Redundant Rdata. According to measurements of Protective DNS services, the configurations of Rdata in rewritten records by some providers have defects. Specifically, along with the rewritten records, several PDNS providers may also include the original malicious records in the DNS response. For local stub resolvers of users, the selection of the resolution result is uncontrollable, and users still have a high probability of accessing malicious resources. Therefore, Protective DNS providers should avoid such redundant configurations to ensure the completeness of the defense effectiveness.

$ORIGIN malicious_domain.com malicious_domain.com A controled_IP;
malicious_domain.com A original_malicious_IP;

Missing Record type. While A records are the most common type of DNS resolution and are often the primary focus of defensive configuration by service providers— since they directly point users to malicious resources— empirical measurements have revealed that some Protective DNS providers fail to protect less common query types, such as TXT records. In these cases, the provider may return original responses, potentially exposing users to hidden threats. This oversight could be exploited to bypass PDNS protections, particularly when malicious domains embed harmful instructions within less scrutinized record types. Therefore, PDNS providers should proactively consider the potential impacts of missing record type configurations.

$ORIGIN malicious_domain.com malicious_domain.com A controled_IP;
malicious_domain.com CNAME controled_domain;

Policy Coverage. In addition to the defensive configuration of the response results, Protective DNS providers should ensure that the defensive functions are effective in all functional scenario.

Specifically, encrypted DNS should also have the same defensive effect as non-encrypted DNS, to prevent malicious domain names from bypassing the defense by merely using encrypted DNS. Additionally, IPv6 scenarios should also be considered.

6.2. Security Consideration 2 - Dangling Resources

Some Protective DNS providers use self-controlled domains or IPs as rewriting strategies. However, mismanagement of these rewritten resources may lead to takeover risks from Dangling Resources, specifically:

1. If the rewritten IP is a cloud service IP, obsolete cloud IP addresses pose a takeover risk.

2. If the rewritten CNAME domain expires, there is an expired domain takeover risk.

3. If the rewritten CNAME domain belongs to a third-party service, subdomain takeover risks may arise.

Upon gaining control of vulnerable rewritten resources (e.g., those at risk of expiration), attackers can trigger DNS queries for malicious domains to the compromised Protective DNS server via phishing tactics. They can then return modified malicious content to victims, enabling connections between victims and attacker servers.

Therefore, Protective DNS service providers should exercise due diligence when using third-party network resources. On one hand, they should consider the financial and management costs of regular maintenance of these resources. On the other hand, they need to periodically verify the status of these third-party services to avoid dangling resource risks.

6.3. Security Consideration 3 - Over-Blocking

Protective DNS rewriting should minimize the impact of over-blocking, as this introduces significant collateral damage in two primary aspects.

Blocklist Construction. First, Protective DNS service providers should avoid errors in blocklists, as blocklist errors directly cause collateral damage to benign domain names. Second, over-generalizing target domains for blocking in Protective DNS may also lead to collateral damage. Using keywords as blocklist entries exacerbates the likelihood of false positives, causing unintended blocking of benign domains and degrading PDNS availability. Employing wildcard domains in blocklists similarly introduces false positives.

Meanwhile, blocking at the second-level domain (SLD) or top-level domain (TLD) levels can also trigger false positives—for example, cloud services often host user-specific services on subdomains, so blocking the apex domain of such a cloud service would impact numerous unrelated services. Thus, blocking at the fully qualified domain names (FQDNs) could minimize collateral damage. Finally, providers should promptly update blocklists to avoid false positives from delayed updates.

$ORIGIN malicious_domain.com malicious_domain.com A controled_IP; (FQDN) "phishing" in domain A controled_IP; (Keyword) *.malicious_domain.com A controled_IP; (Wildcard Domain) *.com A controled_IP; (SLD/TLD Level Domain)

Blocking Policy. The primary defense objective of Protective DNS is to prevent users from accessing any malicious resources, i.e., intercepting as many malicious domains as possible. However, empirical analysis has shown that some Protective DNS implementations exhibit over-blocking collateral damage from aggressive blocking. Measurements reveal that certain Protective DNS services apply extreme defensive strategies to queries for one or more malicious domains, temporarily blocking all domain resolution for the client—including legitimate domains. This introduces denial-of-response (DoR) risks, as attackers can exploit this behavior to impose DoR attacks on arbitrary victims. Specifically, sending a set of malicious domain queries with spoofed source IP addresses can force the victim’s client to lose all DNS resolution capabilities, effectively executing a denial-of-service attack.

Therefore, PDNS service providers should exercise caution when implementing aggressive defensive strategies and consider the potential impact of such approaches in advance. Meanwhile, Protective DNS providers should preconfigure defense mechanisms against potential denial-of-resolution (DoR) risks. Specifically, when a client initiates a large volume of DNS queries exceeding a defined threshold for malicious domains to a Protective DNS server, providers should evaluate the impact of directly blocking all DNS query responses from the client for a period of time. To effectively mitigate denial-of-response attacks, providers can send oversized DNS responses to enforce TCP fallback, thereby thwarting DoR attacks constructed via IP spoofing.

Most critically, operators should strive to avoid controversial blocklist formats to minimize the impact of potential false positives. First, PDNS should refrain from using keywords as Blocklist entries, as this exacerbates the likelihood of introducing false positives and undermines PDNS availability. Second, PDNS providers should avoid using wildcard domains in Blocklists, as such

practices may also lead to false positives. To maximize the mitigation of false positives, mitigation at the minimum subdomain granularity (i.e., FQDN) may minimize collateral damage.

6.4. Security Consideration 4 - Interaction with data integrity protection

For blocked domains, Protective DNS rewrites data in the DNS responses and breaks their data integrity by design. Particularly, when domains are DNSSEC-signed, PDNS will not be able to return validated responses for blocked query names and MUST NOT set the AD bit when this occurs [RFC4035]. For clients that have DO bit or CD bit set in DNS queries, meaning they wish to receive DNSSEC RRs in response messages [RFC4033], PDNS will not be able to provide DNSSEC RRs when the query names are blocked, because no actual queries are sent to authoritative servers. In such cases, error messages (e.g., via EDE [RFC8914]) is recommended in responses for diagnosing purposes.

6.5. Security Consideration 5 - Fallback

As Protective DNS introduces new components, such as blocklists, service providers should consider fault diagnosis for denial-of-service (DoS) failures in individual components and corresponding fallback mechanisms to ensure performance stability. For example, in scenarios involving remote blocklist queries, providers should proactively diagnose the availability of remote blocklist interfaces on a regular basis. If remote blocklist query services become unavailable due to network issues or other causes, and no fallback mechanism is in place, this may render the provider’ s DNS query services inoperable. Thus, providers should predefine fallback mechanisms— such as reverting to normal DNS resolution procedures.

7. IANA Considerations

This document has no IANA actions.

8. References

8.1. Normative References

[I-D.ietf-dnsop-dns-rpz]
Vixie, P. A. and V. Schryver, "DNS Response Policy Zones (RPZ)", Work in Progress, Internet-Draft, draft-ietf-dnsop-dns-rpz-00, 9 March 2017, <https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-dns-rpz-00>.

[I-D.ietf-dnsop-structured-dns-error]
Wing, D., Reddy.K, T., Cook, N., and M. Boucadair, "Structured Error Data for Filtered DNS", Work in Progress, Internet-Draft, draft-ietf-dnsop-structured-dns-error-15, 5 May 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-structured-dns-error-15>.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <https://www.rfc-editor.org/info/rfc1034>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <https://www.rfc-editor.org/info/rfc4033>.

[RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005, <https://www.rfc-editor.org/info/rfc4035>.

[RFC5782] Levine, J., "DNS Blacklists and Whitelists", RFC 5782, DOI 10.17487/RFC5782, February 2010, <https://www.rfc-editor.org/info/rfc5782>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8914] Kumari, W., Hunt, E., Arends, R., Hardaker, W., and D. Lawrence, "Extended DNS Errors", RFC 8914, DOI 10.17487/RFC8914, October 2020, <https://www.rfc-editor.org/info/rfc8914>.

8.2. Informative References

[Canada-Protect]
"Canadian shield offers dns-based protection against malware and phishing attacks", July 2021, <https://www.cira.ca/en/canadian-shield/faq-public/>.

[Cisco] "DNS Security Your New Secret Weapon in The Fight Against Cybercrime", February 2024, <https://umbrella.cisco.com/blog/dns-security-your-new-secret-weapon-in-your-fight-against-cybercrime>.

[DNS4EU] "DNS4EU", May 2025, <https://www.joindns4.eu/>.

[DNS4EU-Public]
"DNS4EU for Public", 2025, <https://www.joindns4.eu/for-public>.

[ICANN-DNSBL]
"How Choice of Reputation Blocklists Affects DNS Abuse Metrics", July 2025, <https://www.icann.org/en/blogs/details/how-choice-of-reputation-blocklists-affects-dns-abuse-metrics-07-07-2025-en>.

[NDSS24] "Understanding the Implementation and Security Implications of Protective DNS Services", March 2024, <https://www.ndss-symposium.org/ndss-paper/understanding-the-implementation-and-security-implications-of-protective-dns-services/>.

[SAC127] "SAC127 DNS Blocking Revisited", May 2025, <https://itp.cdn.icann.org/en/files/security-and-stability-advisory-committee-ssac-reports/sac127-dns-blocking-revisited-16-05-2025-en.pdf>.

[UK-Defence]
"Active Cyber Defence", 2022, <https://www.ncsc.gov.uk/files/ACD6-full-report.pdf>.

[UK-NCSC-PDNS]
"NCSC announces new partnership for PDNS delivery", April 2024, <https://4thplatform.co.uk/2024/04/19/ncsc-announces-new-partnership-for-pdns-delivery-2-2/>.

[UK-NCSC-PDNS-Usage]
"Experts in Domain Name System (DNS) Services", 2025, <https://nominet.uk/dns-services/>.

[US-Protect]
"Protective domain name system services", May 2022,
<https://www.nsa>.

[USENIX24] "Two Sides of the Shield: Understanding Protective DNS adoption factors", August 2024,
<https://www.usenix.org/conference/usenixsecurity23/presentation/rodriguez>.

Acknowledgments

TODO acknowledge.

Authors' Addresses

Haixin Duan
Tsinghua University
Beijing
China

Mingxuan Liu
Zhongguancun Laboratory
Beijing
China

Baojun Liu
Tsinghua University
Beijing
China

Chaoyi Lu
Zhongguancun Laboratory
Beijing
China