EDNS NSID Option

Christian Kilb, Johannes Zirngibl*, Patrick Sattler*
*Chair of Network Architectures and Services, Department of Informatics
Technical University of Munich, Germany
Email: christian.kilb@tum.de, zirngibl@net.in.tum.de, sattler@net.in.tum.de

Abstract—**The DNS Name Server Identifier (NSID) Option is a DNS extension that helps disambiguate name servers which share IP addresses in anycast setups. NSID is useful for DNS analysis by name server administrators and researchers. This paper evaluates the usage of NSID by 89k name servers authoritative for 533k Alexa Top Sites. It analyzes in particular how widespread its usage is and how identifiers are chosen. The evaluation shows that about a third of name server IP addresses provide NSID support and two-thirds of Alexa domains have at least one name server with support. 81% of observed NSID values are valid UTF-8 strings. Every third NSID value is a domain name. In two out of three cases, these domains resolve back to the name servers' IP addresses.**

Index Terms—**Domain Name System, EDNS, NSID, network measurement**

1. Introduction

In DNS anycast setups, multiple name servers share a single IP address. DNS queries to that IP address are then answered by one of the name servers, for low latency ideally by one in close physical proximity to the client. The DNS response however does not make it apparent to the client which concrete anycast name server handled the query. When researching or debugging anycast, it might be useful or necessary to learn the identity of the responding name server. The DNS Name Server Identifier (NSID) Option, specified in RFC 5001 [1], is a DNS extension that standardizes such a mechanism. It allows name servers to include a name server identifier in their DNS responses, if requested by a client. This NSID value can then be used to disambiguate the anycast name servers.

This paper contributes a usage evaluation of the NSID DNS extension. The analysis addresses the question of how widely it is supported by the most popular name servers. It also evaluates how the NSID values are chosen and whether they follow certain patterns, which is of interest because NSID values are specified as arbitrary byte strings. An NSID dataset obtained from a DNS scan of the Alexa ranked domains is the basis for the usage evaluation.

In Section 2 of this paper, background information about NSID and its requirement EDNS (Extension Mechanisms for DNS) is given. Afterwards, related concepts and work are outlined in Section 3. In Section 4, a DNS scan dataset is evaluated and the usage of EDNS and NSID is analyzed, with a focus on NSID. Finally, a conclusion is drawn and future work is suggested in Section 5.

2. Background

Should name server administrators choose to support the NSID Option, they first have to provide EDNS support, on which NSID is built.

2.1. EDNS

The "Extension Mechanisms for DNS" (EDNS) [2] extends DNS [3] in multiple ways. It increases the maximum DNS message payload size over UDP from 512 B to 65 535 B. It also extends the number of possible return codes and flags.

Classical DNS messages over UDP have a fixed maximum payload size of 512 B. DNS over TCP could be used to circumvent this size limit, which however would be inefficient for single DNS request-response exchanges, as a TCP handshake would have to be performed. EDNS was therefore created to allow for efficient DNS messages over UDP with extended limits in a backward-compatible manner.

EDNS in its version 0 defines a new (pseudo) resource record called "OPT". It contains meta information, but no actual DNS data. DNS clients that support EDNS can include an OPT resource record in the "additional data" section of their request. A DNS server with EDNS support would then process it accordingly and add a corresponding OPT record to its response.

The format of the OPT resource record is shown in Figure 1. Some resource record fields have a different meaning compared to regular DNS records. The `NAME` field always has value 0 to indicate the root domain. A type value of 41 has been assigned to the OPT resource record, to which the `TYPE` field is set. In the reinterpreted `CLASS` field, the requestor specifies the maximum UDP response payload size it is able to receive. The extended return code, EDNS version number and extended flags are embedded in the reinterpreted `TTL` field. The last resource record field `RDATA` contains a list of "options" in the form of attribute-value pairs. The size of the `RDATA` field is found in the preceding field `RDLEN`.

Figure 2 shows the format of an option. They consist of the fields `CODE`, `LENGTH` and `DATA`. The length field specifies the size of the option data. If a DNS client includes an option in its OPT resource record and the DNS server understands it, a corresponding option will be included in the response. Unsupported option codes would be ignored instead.

One example for an extension to DNS that builds on EDNS is the DNS security extension DNSSEC [4]. It

doi: 10.2313/NET-2021-05-1_08

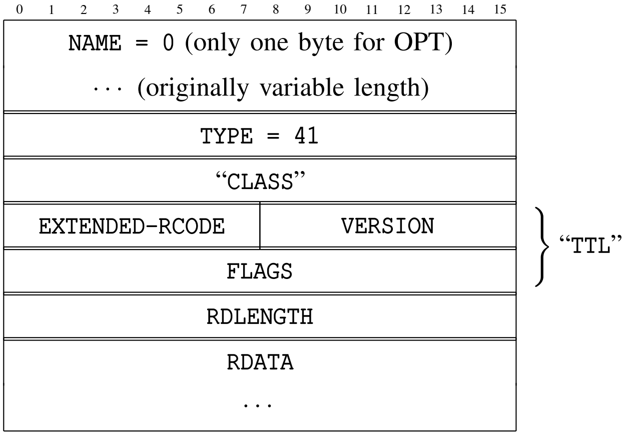

Figure 1: OPT Resource Record Format

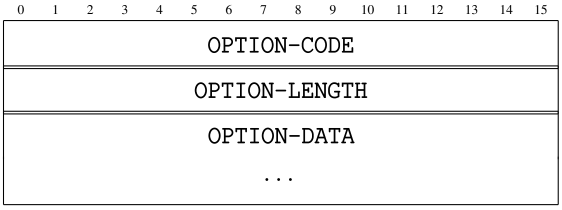

Figure 2: OPT Option Format

introduces new resource records that hold security-related information and also allocates a new EDNS flag in the FLAGS section of the OPT resource record.

2.2. NSID

With DNS anycast, multiple name servers can share a single IP address. In such a scenario, the IP address is not enough to tell which of the name servers responded to a query. In order to reliably learn the identity of the responding name server, it would have to include a server identifier directly in the response.

The DNS Name Server Identifier (NSID) Option [1] is a DNS extension that realizes this. With an NSID in the DNS response, clients can disambiguate the name servers with shared IP addresses. This can be useful for debugging DNS problems as well as for internet research.

NSID is realized as an EDNS Option. It is assigned the option code 3. DNS clients can request an NSID by including an EDNS Option with code 3 and empty option data in their request. Should the DNS server support NSID, it includes its identifier in the response, embedded in the OPTION-DATA field.

The meaning of NSID values in this option data field is however undefined. NSIDs are specified as raw byte strings or sequences of hexadecimal digits. It is up to the name server administrators to decide on the meaning of NSID values for their servers. Several suggestions are given in the RFC for the NSID meaning. Possible choices are a "real" host name or IP address, a static identifier derived from the name or IP address, a dynamically generated identifier or an encrypted identifier. The administrators can therefore decide whether the NSID should be meaningful for everyone or only for a specific group of people, such as themselves. Should the value be meaningful for everyone, an appropriate encoding must also be chosen, for example UTF-8.

NSID is a hop-by-hop DNS extension, i.e. requests for NSID values are not recursively forwarded by resolvers. Instead, clients learn the NSID value of the DNS server or recursive resolver that they directly addressed with their NSID request.

3. Related work

Work related to EDNS and NSID is rather sparse. In 2020, Stipovic [5] analyzed the RFC compliance of EDNS implementations in popular DNS server software such as BIND.

Before the introduction of NSID in 2007, there was another, non-standard mechanism to query for a name server identifier [6]. A BIND name server could be configured to return the server host name when queried for a TXT resource record of the special domain "HOSTNAME.BIND.". Usually the CLASS of DNS requests is "IN" for "Internet". For such an identifier query however, the "Chaosnet" (CH) class was used instead. Similar "CHAOS" queries also allowed clients to request the BIND server version.

Fan et al. [7] made use of such CHAOS queries in combination with traceroutes in 2013 to evaluate DNS anycast. They enumerated as many DNS anycast nodes as they could find. They did however intentionally not use the NSID extension in their scans, due to the lack of standardized NSID values, the lack of recursive queries and too little NSID deployment at that time.

Li et al. [8] also used CHAOS queries in 2018 instead of NSID in their analysis of internet anycast. No specific reason was given this time, but they mentioned that such queries were commonly used to analyze anycast.

4. Evaluation

The Alexa Top Sites [9] domain list provided the basis for a DNS scan with the goal of creating an NSID dataset. The list from 23rd Nov 2020 contained 532 839 entries. Based on these domains, an exhaustive DNS scan was performed one day later. All name servers authoritative for these domains and higher-level domains such as top-level domains were scanned, resulting in 14 782 146 executed queries in total. The query parameters and the server responses have been recorded.

Not all rows of the DNS scan dataset are considered in the NSID evaluation. Queries that resulted in an error response are for example filtered out. The analysis is further being limited to queries for domains from the Alexa Top Sites list. All queries with a NAME that is not on the Alexa list are therefore filtered out. The same applies to queries with a TYPE other than "A" or "AAAA". The remaining 2 885 428 rows of the dataset (19.5 %) have been included in the following analysis.

4.1. EDNS and NSID usage

89 003 unique name server IP addresses have been found to be directly responsible for the Alexa domains. 87 739 (98.6 %) of these NS IPs supported EDNS, as

TABLE 1: Analysis of EDNS and NSID usage

Description	#IPs	Rel. %
All NS IPs	89 003	100.0
Consistent EDNS support	87 285	98.1
No EDNS support	1264	1.4
Inconsistent EDNS support	454	0.5
NS IPs with EDNS support	87 739	100.0
Consistent NSID support	28 100	32.0
No NSID support	59 203	67.5
Inconsistent NSID support	436	0.5
NS IPs with NSID support	28 536	100.0
Consistent NSID value	25 123	88.0
Varying NSID value	3413	12.0
Description	**#Domains**	**Rel. %**
Alexa Top Sites	474 254	100.0
Any NSID support	312 223	65.8
No NSID support	162 031	34.2

TABLE 2: Analysis of NSID support of top TLDs of the Alexa Top Sites

TLD	#Domains	%Supp.	TLD	#Domains	%Supp.
com	255 759	65.8	ua	7050	65.4
ru	32 946	60.0	au	4814	99.9
net	19 113	67.9	tr	4705	63.8
org	16 421	69.4	co	4512	76.7
ir	15 441	31.7	uk	4012	85.0
in	7577	80.4	gr	3913	50.5

their inclusion of OPT records in DNS responses show. A small amount of servers (454 IPs) however only provided inconsistent EDNS support, i.e. for some queries they did support it and for other queries they did not. Out of all unique name server IPs, 28 536 (32.1 %) supported NSID in addition to EDNS and included name server identifiers in their responses. Similarly, there was a small portion (436 NS IPs) with only inconsistent NSID support.

Most name server IPs that did return an identifier only identified themselves with that single NSID. This applies to 25 123 (88.0 %) of the NS IPs that sometimes or always supported NSID. The other 3413 (12.0 %) IPs have replied with different NSIDs for multiple queries. Such varying NSIDs are an indicator for the presence of anycast, in which case multiple queries to the same IP would have been answered by different name servers. It is however no proof for anycast, as the RFC specification of NSID also allows for dynamic identifiers. In that case, the same name server would respond to possibly every query with a different NSID, whose meaning might only be apparent to the server's administrators.

While only about a third of name server IPs supported NSID, two-thirds of Alexa domains actually did provide at least some NSID support. 312 223 (65.8 %) of the Alexa Top Sites name servers responded with an NSID at least once.

Table 1 summarizes the EDNS and NSID usage findings. In Table 2, the NSID support percentages of the Alexa domains are shown, grouped by the top-level domains with most occurrences. Similar to above, a domain is counted as one that supports NSID if and only if one of its name servers responded with an NSID at least once.

TABLE 3: Analysis of NSID values

Description	#NSIDs	Rel. %
Decodable NSIDs	33 864	80.9
Valid domain name	13 792	40.7
IPv4-like	350	1.0
IPv6-like	0	0.0
Contains IATA airport code	10 253	30.3
Contains “ns” or “dns”	7482	22.1
Hyphenated alphanumeric	14 089	41.6
Specific 32-hex-char pattern	4828	14.3
Non-decodable NSIDs	7973	19.1

TABLE 4: Analysis of NSID domains

Description	#NSIDs	Rel. %
All NSID domains	8090	100.0
Resolved to an IP	6687	82.7
Mapped to original NS IP	5214	78.0
Mapped to multiple IPs	999	14.9
Mapped to IPv4 only	5752	86.0
Mapped to IPv4 and IPv6	932	13.9
Mapped to IPv6 only	3	0.0

4.2. NSID values

In the DNS scan dataset, 41 837 unique (NSID, NS IP)-pairs could be observed (with 10 473 unique NSIDs). The NSID values and their possible meaning are evaluated in the following. NSIDs are counted multiple times if and only if different IPs announce the same NSID. Table 3 gives a summary of the findings.

4.2.1. Decodable NSIDs. 33 864 (80.9 %) of the NSIDs could be decoded to valid UTF-8 strings. These NSIDs were between 1 B and 74 B long. They have been subjected to further automated analysis that tries to match them to certain regular expressions. Subsequent percentage numbers are given relative to the number of UTF-8-decoded NSIDs.

Domain names. In a check of each NSID against a regular expression for valid domain names, 13 792 (40.7 %) of NSIDs fully matched the domain name syntax and 14 148 (41.8 %) NSIDs partially matched, i.e. had a domain string embedded in the NSID. In order to find out whether the fully matching domain strings actually resolve to an IP address, a follow-up DNS scan has been performed on 22nd Dec 2020 on the 8090 unique NSID domain strings. 6687 (82.7 %) of NSID domains indeed resolved to an IP address. A majority of these (5214 or 78.0 %) also mapped back to their original name server IP. Some of the resolved NSID domains pointed to multiple IP addresses (999 or 14.9 %). In most cases, the resolved IPs were IPv4 addresses (5752). In 932 cases, the NSID domain resolved to both an IPv4 and IPv6 address. In three cases, it resolved to an IPv6 address only. A summary of the domain string statistics is given in Table 4.

IP addresses. Next to domain names, name server administrators could also choose to set an IP address as identifier. A check of each NSID against a regular expression for valid IP addresses however revealed that almost none were doing so. Only 15 NSIDs fully matched the syntax of IPv4 addresses and zero that of IPv6 addresses. 34 (0.1 %) of NSIDs partially matched the IPv4 syntax, often in form of the IP address being embedded in a domain name, with still zero IPv6 matches. In the presence of domain names, it might be beneficent to format IP addresses without dots within them. After modifying the IP regular expressions by replacing dots and colons with hyphens, more matching NSIDs have been found. 316 (0.9 %) of NSIDs contained an IPv4 address with hyphens, often as part of a domain name. IPv6 addresses did not seem to appear in any NSID, independent of colon or hyphen as separator.

Airport codes. Sometimes server administrators choose to include a regional identifier in domain names. NSID values could also contain such regional identifiers. Airport codes [10] are one type of location identifier. A check against the three-lettered IATA airport codes revealed that 10 253 (30.3 %) of NSIDs seemed to contain such an airport code. To reduce the amount of false positives, this check has been conducted with some additional conditions applied to the regular expression. The airport code had to appear as whole word and before any dot, limiting its appearance to any first domain label. Additionally, a number of codes have been blacklisted due to them also being technical abbreviations, for example "cdn", "srv" and "vps".

Miscellaneous. Some more arbitrary regular expression checks have also been performed. With 7482 (22.1 %) NSIDs, quite a few identifiers contained the string "ns" or "dns" as whole word in lower- or uppercase. 144 (0.4 %) identifiers were just integer numbers. 715 (2.1 %) NSIDs purely consisted of alphabet letters, i.e. a to z in lower- or uppercase. Many NSIDs (14 089 or 41.6 %) consisted of only alphabet letters, digits and hyphens. A small number of IDs were found to be so called "globally unique identifiers" (87 IDs or 0.3 %). Some NSIDs followed a very specific pattern, consisting of 32 hexadecimal digits, followed by two spaces and a dash. This was the case for 4828 NSIDs (14.3 %).

4.2.2. Non-decodable NSIDs. 7973 of the NSID values (19.1 %) could not be decoded to valid UTF-8 strings. These NSIDs were between 2 B and 48 B long. Almost all of the non-decodable NSIDs were duplicates (99.5 %). Only 42 values were unique.

The 7973 NSID values have then be subjected to another decoding attempt. This time however, invalid bytes have been ignored in the decoding process. 84.1 % of the bytes could be decoded to partial NSID UTF-8 strings this way.

A manual review of these partial strings led to more insights. In all except three cases, the string started with the sequence of non-printable ASCII characters `NUL SOH CAN`. In most cases (6458 or 81.0 %), this sequence was only extended by another `NUL`. Sometimes, the sequence continued with `NUL ETX NUL` instead.

Only 1494 partially decoded NSIDs contained printable ASCII characters. There was one repeating ASCII pattern in some of the NSIDs, containing the word "proxy", a number and presumably a location abbreviation. One example for this is "proxy-121-defra.hivecast-121-defra", where "defra" seems to mean Frankfurt, Germany. In a similar find, "nlams" apparently means Amsterdam, Netherlands, which solidifies the interpretation as location code. Such proxy patterns were always preceded by a sequence of non-printable characters. The similarities of the proxy text patterns suggest that these NSIDs belong to the same service.

It did not become apparent what the meaning of the other, non-UTF-8 bytes was. As the RFC defines NSID values as arbitrary byte strings, the name server administrators could have chosen a more obscure meaning here.

5. Conclusion and future work

This paper analyzed the usage of the NSID DNS extension by the name servers of the Alexa domains. About a third of the name server IP addresses did support NSID, while about two-thirds of domains supported it at least sometimes. Most NSID values could successfully be decoded to UTF-8 strings. Many of these name server identifiers have been found to be domain names, most of which even resolved back to their original name server IP address.

In future work, a closer look could be taken at the IP addresses and domain names of the name servers that are using NSID in order to learn more about *who* is using it, in addition to *if* and *how*. Next to name servers responsible for the Alexa domains, the use of NSID by other name servers could also be investigated. These could be root servers, TLD servers or other, less popular name servers. Another idea for future research is to evaluate the usage of anycast with the help of multiple NSID scans performed from geographically distinct locations.

References

[1] R. Austein, "DNS Name Server Identifier (NSID) Option," RFC 5001, Aug. 2007. [Online]. Available: https://rfc-editor.org/rfc/rfc5001.html

[2] J. L. S. Damas, M. Graff, and P. A. Vixie, "Extension Mechanisms for DNS (EDNS(0))," RFC 6891, Apr. 2013. [Online]. Available: https://rfc-editor.org/rfc/rfc6891.html

[3] P. Mockapetris, "Domain Names - Implementation and Specification," RFC 1035, Nov. 1987. [Online]. Available: https://rfc-editor.org/rfc/rfc1035.html

[4] S. Rose, M. Larson, D. Massey, R. Austein, and R. Arends, "DNS Security Introduction and Requirements," RFC 4033, Mar. 2005. [Online]. Available: https://rfc-editor.org/rfc/rfc4033.html

[5] I. Stipovic, "Analysis of an Extension Dynamic Name Service – A Discussion on DNS Compliance with RFC 6891," 2020.

[6] D. R. Conrad and S. Woolf, "Requirements for a Mechanism Identifying a Name Server Instance," RFC 4892, Jun. 2007. [Online]. Available: https://rfc-editor.org/rfc/rfc4892.html

[7] X. Fan, J. Heidemann, and R. Govindan, "Evaluating Anycast in the Domain Name System," in *2013 Proceedings IEEE INFOCOM*, 2013, pp. 1681–1689.

[8] Z. Li, D. Levin, N. Spring, and B. Bhattacharjee, "Internet Anycast: Performance, Problems, & Potential," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 59–73. [Online]. Available: https://doi.org/10.1145/3230543.3230547

[9] Alexa Internet, Inc., "Alexa Top Sites," Nov. 2020. [Online]. Available: https://www.alexa.com/topsites

[10] Fubra Limited, "World Airport Codes," Dec. 2020. [Online]. Available: https://www.world-airport-codes.com/