TU Delft

Delft University of Technology

Document Version
Final published version

Licence
CC BY

Citation (APA)
Osali, F., Alipour, I., Gañán, C. H., Feldmann, A., & Fiebig, T. (2026). *Poster: Private, Reserved, Visible: IPv4 Bogon Addresses in DNS*. 1235-1236. Poster session presented at 2026 ACM Internet Measurement Conference, IMC 2026, Karlsruhe, Germany. https://doi.org/10.1145/3777912.3847604

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Poster: Private, Reserved, Visible: IPv4 Bogon Addresses in DNS

Fariba Osali
MPI for Informatics/Saarland University
fosali@mpi-inf.mpg.de

Iman Alipour
MPI for Informatics
ialipour@mpi-inf.mpg.de

Carlos H. Gañán
TU Delft
C.HernandezGanan@tudelft.nl

Anja Feldmann
MPI for Informatics
anja@mpi-inf.mpg.de

Tobias Fiebig
TU Wien
tobias@internet.wien

Abstract

In this poster, we provide an empirical investigation of bogon IPv4 addresses in Domain Name System (DNS). To this end, we conduct a longitudinal analysis of large-scale passive DNS data to characterize their patterns and evolution over 10 years.

CCS Concepts

• **Networks** → **Network measurement**.

Keywords

DNS, IPv4 bogon addresses

ACM Reference Format:
Fariba Osali, Iman Alipour, Carlos H. Gañán, Anja Feldmann, and Tobias Fiebig. 2026. Poster: Private, Reserved, Visible: IPv4 Bogon Addresses in DNS. In *Proceedings of the 2026 ACM Internet Measurement Conference (IMC '26), October 12–16, 2026, Karlsruhe, Germany*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3777912.3847604

1 Introduction

The DNS [9] is essential for connecting users and services on the Internet. Among other uses, DNS translates human-readable domain names to machine-readable IP addresses. In general, when DNS returns globally routable addresses, users can reach the corresponding services over the *public* Internet. However, DNS records can also contain addresses that are not globally routable. These include private, reserved, or special-purpose IPv4 and IPv6 addresses, commonly referred to as bogon addresses [3, 6, 7, 10].

Indeed, *at the protocol level*, DNS does not prohibit using non-publicly routable IP addresses. For example, nothing prevents one from setting an `A` record for `ns.example.com` to `192.168.0.1` and using that name as an `NS` record for `example.net`. As illustrated in Fig. 1, in fact, it is common practice to include addresses from internal networks—often private or reserved—in DNS records to make back-office applications accessible via a name, for example. However, when such addresses are used in `NS` records, they can make the corresponding domain unreachable to external users.

At the same time, the presence of a bogon address in publicly observable DNS does not necessarily indicate a misconfiguration.

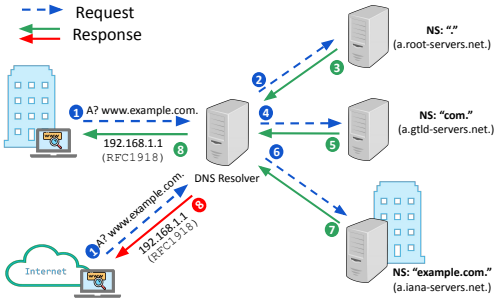


Figure 1: DNS resolution overview for an `A` query for `www.example.com`, where a bogon IP address is returned to both an internal client and a client on the public Internet.

Such records may reflect intentional operational practices or internal naming conventions. However, when records intended only for internal use become externally visible, they may indicate configuration errors, reveal information about internal addressing, or result in unreachable services.

The use of bogon IP addresses is becoming increasingly relevant as operators use private, reserved, and special-purpose address space in large-scale networks. Recent operational reports show that large-scale networks facing exhaustion of private IPv4 address space have started to consider or use reserved IPv4 space such as `240.0.0.0/4` [1, 8]. Related concerns also arise in IPv6, where the use of reserved IPv6 addresses is a notable problem for assessing the IPv6 readiness of DNS [5]. Such developments blur the boundary between addresses that are globally routable, privately useful, and invalid as public service endpoints. The operations community has called for such a study [11]. In this poster, we present the first empirical investigation of IPv4 bogon addresses within DNS. To this end, we leverage a passively collected DNS dataset spanning January 2015 to April 2025 to characterize the operational context of bogon IP addresses. We develop a multi-layered classification scheme that characterizes the likely *purpose* of bogon IP use by analyzing the names in which these bogon addresses appear.

2 Methodology and Dataset

We use the Farsight Security Information Exchange (Farsight SIE) passive DNS dataset, now part of DomainTools [4]. It offers broad coverage, global visibility into queried names, and a historical view. In this study, we rely on data from 2015-01-01 to 2025-04-30, aggregated monthly.

This work is licensed under a Creative Commons Attribution 4.0 International License.
IMC '26, Karlsruhe, Germany
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2327-8/2026/10
https://doi.org/10.1145/3777912.3847604

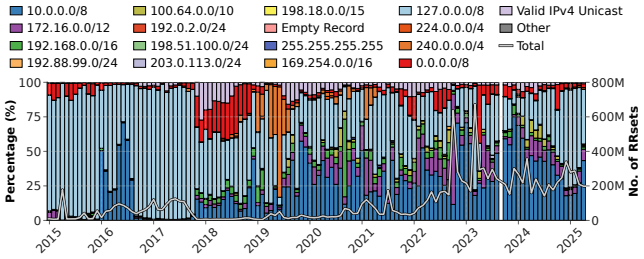


Figure 2: Temporal evolution of IPv4 address categories for A records in filtered RRsets. Stacked bars (left axis) show the relative share of each category per snapshot; the line (right axis) shows the total number of RRsets.

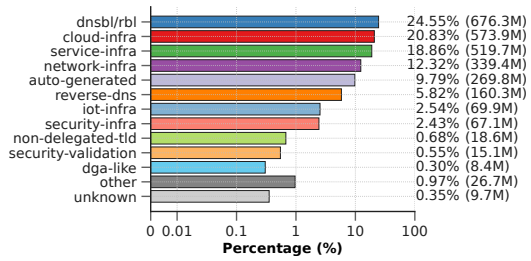


Figure 3: Distribution of primary labels for A records. The x-axis uses a symmetric logarithmic scale, with a linear region from 0 to 0.01%.

Dataset Selection and Bogon Identification: For the analysis in this work, we focus specifically on IPv4 bogon addresses observed in A records, as A records directly associate domain names with IPv4 addresses. We select A Resource Record sets (RRsets) containing at least one IPv4 address belonging to a curated set of bogon and special-purpose IPv4 prefixes. When an RRset contains both bogon and globally routable addresses, we retain the complete RRsets to preserve this context. The resulting A RRsets form the basis of the classification and quantitative analyses presented in this work.

Classification Pipeline: We develop the classification scheme through iterative manual analysis of the data. We first inspect the dataset, focus on heavy hitters, manually analyze their semantics, and generalize the observed patterns into regular expressions. We further use additional indicators and external data sources to identify dynamically or automatically generated Resource Record names (RRnames) and non-delegated Top-level Domains (TLDs). We repeat this process until each RRset receives at least one label.

Primary and Secondary Labels: The classifier first assigns one primary label to each RRset and may then attach secondary labels. Primary labels capture the likely operational context of the RRset, while secondary labels describe additional properties of the RRname, such as hash-like labels, IPv4-like tokens, or wildcard indicators.

Rule/Regex Creation Procedure: To build our rule set, we manually inspect more than two years of data to identify patterns. We group syntactically similar RRnames using character *n*-grams and TF–IDF [2, 12], and add rules for recurring dominant families. We infer group semantics from naming patterns and, where useful, the relation between RRname and Record data (Rdata).

Classifier Validation: We manually validate a sample of classified RRsets by checking whether each assigned category matches its RRname, Rdata, naming patterns, and operational context.

3 Preliminary Results

Fig. 2 shows the temporal distribution of IPv4 address categories observed in RRsets after filtering. Notable fluctuations and short-lived spikes in the total number of RRsets can be observed. Overall, the total number of RRsets exhibits substantial growth and variability, particularly after 2019. In earlier years, a larger fraction of addresses falls into reserved, special-purpose, or non-standard categories. Private address spaces (*e.g.*, `10.0.0.0/8`, `127.0.0.0/8`, `0.0.0.0/8`) contribute significantly.

Using the pipeline described in Sec. 2, we find that `dnsbl/rbl` is the most frequent primary label for A RRsets, accounting for 24.55 % of observations (Fig. 3). These names reflect abuse reporting and reputation infrastructure (blacklists and blocklists) and represent a recurring class of observations in our dataset. The next two largest groups are `cloud-infra` (20.83 %) and `service-infra` (18.86 %), capturing infrastructure naming associated with cloud providers (*e.g.*, cloud-provider and region-specific naming conventions) and service platforms (*e.g.*, API and SaaS-oriented conventions). We also observe `network-infra` (12.32 %), which captures internal network naming as well as ISP and access-network infrastructure. Finally, we could not identify a category for 9.7M names, which we classify as unknown.

4 Conclusion and Future Work

In this poster, we provide an empirical characterization of IPv4 bogon address use in DNS. Our study sheds light on *how* and *why* bogon addresses end up in public DNS. Finally, our classification provides a longitudinal view of bogon use in DNS, highlighting the prominent role of cloud and infrastructure naming schemes. As future work, we will further investigate the persistence and temporal behavior of these observations and extend our analysis to IPv6 bogons and other record types.

References

[1] Basant Amarkhed and Gopinath Balakrishnan. 2024. Leveraging Class E IPv4 Address Space to Mitigate IPv4 Exhaustion Issues in GKE. https://cloud.google.com/blog/products/containers-kubernetes/how-class-e-addresses-solve-for-ip-address-exhaustion-in-gke.
[2] William B. Cavnar and John M. Trenkle. 1994. N-Gram-Based Text Categorization. In *3rd Annual Symposium on Document Analysis and Information Retrieval.*
[3] M. Cotton, L. Vegoda, R. Bonica, and B. Haberman. 2013. *Special-Purpose IP Address Registries.* RFC 6890.
[4] DomainTools. 2026. DomainTools. https://www.domaintools.com/
[5] Tobias Fiebig and Anja Feldmann. 2025. 'How I learned to stop worrying and love IPv6': Measuring the Internet's Readiness for DNS over IPv6. In *ACM Internet Measurement Conference (IMC).*
[6] Internet Assigned Numbers Authority (IANA). 2026. IPv4 Special-Purpose Address Space.
[7] Internet Assigned Numbers Authority (IANA). 2026. IPv6 Special-Purpose Address Space.
[8] Qasim Lone. 2022. 240/4 As Seen by RIPE Atlas. https://labs.ripe.net/author/qasim-lone/2404-as-seen-by-ripe-atlas/.
[9] P. Mockapetris. 1987. *Domain names-implementation and specification.* RFC 1035.
[10] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. 1996. *Address Allocation for Private Internets.* RFC 1918.
[11] RIPE NCC. 2024. IPv6 WG Minutes–RIPE 88. https://www.ripe.net/community/wg/active-wg/ipv6/minutes/ipv6-working-group-minutes-ripe-88/.
[12] Gerard Salton and Christopher Buckley. 1988. Term-Weighting Approaches in Automatic Text Retrieval. *Information Processing & Management* (1988).