RIPEn at Home – Surveying Internal Domain Names using RIPE Atlas

Elizabeth Boswell
University of Glasgow

Colin Perkins
University of Glasgow

***Abstract*—Internal domain names are domain names that are resolved locally and not by the global DNS. Name collisions occur if an internal name is resolved in the global DNS, e.g. if queries are accidentally sent to a public resolver. This can lead to security issues. While previous studies of name collisions used passive measurement data, we use active measurements on RIPE Atlas to survey the use of internal names in home networks. We discover 3092 names, used by 4305 probes, of which 34.51% are at risk of collision if their top-level domain is delegated.**

I. Introduction

Many home networks let users refer to the gateway using an internal domain name that the gateway resolves to its own local address [1]. These internal names often use top-level domains (TLDs) that don't exist in the public DNS. If the TLD doesn't exist, any query for the name that is inadvertently sent to the public DNS will fail [2] [3]. However, this practice can cause issues if the TLD is delegated, as a recent case has shown.

AVM FRITZ!Box home gateways, one of the most popular home gateways in Germany, use internal names under the `box` TLD (e.g. `fritz.box`) for the gateway's configuration page and other features. The `box` TLD was added to the DNS root in August 2023 [4], and advertised to the general public on 18 January 2024 [5] [6]. AVM did not appear to register `fritz.box` and other related names, and for several weeks in January and February 2024, several such names were owned by likely domain speculators. This is a security risk, as queries for `fritz.box` could accidentally be sent to the public DNS, e.g. when using a public resolver. The public `fritz.box` domain could spoof the home gateway, e.g. to steal login credentials, misguide users to install malicious software, or otherwise interfere with the home network.

There is no comprehensive survey of which internal names are used by home gateways, and which of these names are vulnerable to name collision. In this paper we use RIPE Atlas to survey internal names used in the probes' local networks. We develop a way to find internal names, and then determine which names are vulnerable to name collision, and which names could be vulnerable if the TLD is delegated.

We find 3092 internal names used by 4305 RIPE Atlas probes. Of these, 2.13% are currently vulnerable to collision (e.g. unregistered subdomains of existing TLDs), and 34.51% use an undelegated TLD and could be vulnerable if it is delegated.

Internal domain names were studied extensively due to the introduction of new generic top-level domains (gTLDs), starting in October 2013 [7]. As of November 2023, there are 1241 new gTLDs [8]. There was a concern that if common internal TLDs are added to the DNS, these names could resolve in two different ways depending on where the query is sent. Studies from the time [1] [9] [10], which find a large variety of internal names, use root server data and other passive measurement sources. We perform active client-side measurements, as they can capture internal names that don't frequently appear in root server logs because the queries are usually answered by the local resolver. We focus on internal names used by home gateways, so we use the RIPE Atlas measurement network [11], which has many vantage points in home networks. We also provide an update on the usage of internal names, over 10 years after the introduction of new gTLDs.

We structure the remainder of this paper as follows. In section II we present background information on the DNS, name collisions, and RIPE Atlas. Our method for detecting internal names on RIPE Atlas is described in section III, the results are discussed in section IV. Related work is discussed in section V and we conclude in section VI.

II. Background

The Domain Name System (DNS) [12] is a globally distributed system mapping domain names to IP addresses and other data. It is organised hierarchically, different parts of the namespace are controlled by different entities. The root of the DNS is managed by the Internet Corporation for Assigned Names and Numbers (ICANN) [13] through the Internet Assigned Numbers Authority (IANA) [14]. Absent any caching, regular DNS queries are first sent to a DNS root server, which refers the DNS resolver to the nameserver responsible for the name's TLD (e.g. `com`). The TLD nameserver refers the resolver to a nameserver lower in the hierarchy, and so on, until a nameserver responds with the queried data.

Some networks use internal domain names to refer to local devices such as the gateway. Instead of being sent to the DNS, a local nameserver responds to queries for these names [1].

The use of internal names can lead to *name collisions*, where an internal name also exists in the global DNS and is inadvertently resolved in the global DNS [15]. If the two responses differ, e.g. if `example.net` is resolved to `192.168.1.1` locally but to `203.0.113.2` by the global DNS, and if the internal and global name are not controlled by the same entity, the global name can spoof the local resource [16].

978-3-903176-64-5

RIPE Atlas [11] is a global Internet measurement network, consisting of ∼12,000 *probes*, which are custom measurement devices or virtual machines located in various networks. Probes are used as vantage points for measurements such as traceroute or DNS queries. The large number of probes, many of which are located in home networks, make it a suitable choice for client-side measurements of internal names.

III. Methodology

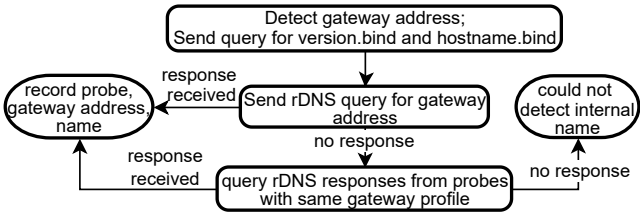

Figure 1. Procedure for detecting internal names.

We don't know beforehand which internal names a RIPE Atlas probe uses, and the large number of possible names makes it unfeasible to exhaustively query them. Instead, we detect internal names using traceroutes and DNS based fingerprinting. Note that we couldn't detect multicast DNS names [17], as RIPE Atlas doesn't support mDNS queries.

We use `CHAOS TXT` queries for `hostname.bind` and `version.bind` for gateway fingerprinting. Some BIND resolvers respond to these queries with their hostname and version of BIND [18]. Many resolvers don't respond, but their response codes (e.g. `NXDOMAIN`, `SERVFAIL`) can differ. Two gateways with the same responses to the BIND queries and the same local address, determined using the same method (see below), have the same *gateway profile*. We assume they are more likely to be the same gateway model and use the same internal name.

Our method, as shown in Figure 1, is:

1. Detect gateway address: Home gateways often integrate a NAT44, so we assume the gateway is a NAT44 or behind it[1]. We estimate the *local* address of the probe's home gateway (if present). We do this in two different ways, so steps 2-4 are performed twice. The two methods are:

a) Traceroute We assume the gateway is the last private address[2] in an IPv4 traceroute starting at the probe[3].

b) Local resolver We assume the gateway has the same address as the probe's DNS resolver, provided it has a private, non-loopback address.

The next steps use the probes with gateways found in step 1.

2. BIND queries: Each probe sends `CHAOS TXT` queries for `version.bind` and `hostname.bind` to its default resolver(s).

3. rDNS queries: Each probe sends a reverse DNS (rDNS) query for its gateway address to its default resolver(s). Any response is an internal name[4].

4. Gateway profile fingerprinting: If a probe didn't receive a response to its rDNS query, we gather the names from rDNS responses received by probes with the same gateway profile. The probe sends `A` record queries for all such names that were received by two different probes. If the response contains a local address and is different from the response from the global DNS, the name is an internal name. This step doesn't discover any new names, but it finds more probes using internal names.

[1]this will also detect non-residential NAT gateways, but more fine-grained detection of home gateways is out of scope for this work.

[2]defined as being in the IANA Special-Purpose Address Registry [19] [20]

[3]Probes could be behind a carrier-grade NAT, so a better choice might be the last address in the probes' private address range, but this could be inaccurate as many probes' local networks appear to use several private prefixes.

Table I
Results of the internal name detection (GW = Gateway).

	Probe GW found	rDNS resp. probes	GW profile probes	% of probes with GW	Names found
TR GW	7441	2573	63	35.43%	1344
Res. GW	6045	3872	100	65.71%	2574

IV. Results

A. Internal name detection

We performed the internal name detection on all available IPv4 probes in early 2024. The results are shown in Table I. We found gateway addresses for 7441 probes using traceroute (method a) and 6045 probes using DNS (method b). However, only 2573 probes received a response for their rDNS query for the traceroute gateway address, while 3872 received a response for the local resolver address. Note that the gateway IP addresses often differ: none of the addresses match for 2104/5021 probes for which we could determine a gateway address with both methods. The gateway profile fingerprinting step added another 102 probes. In total, we found 3092 internal names, used by 4305 probes (50.86% of probes tested).

Figure 2 shows the top 10 internal full, second+top-level and top-level domains, by number of probes using them. All top 10 full domain names appear to be related to the FRITZ!Box. This is likely due to its popularity in Europe (where many RIPE Atlas probes are located), and because a single rDNS query to a FRITZ!Box often returns multiple names. Five of the top 10 second+top-level domains are also likely FRITZ!Box-related, alongside other names such as `pi.hole` (PiHole ad blocker).

The most common TLDs are `box`, `lan` and `nas`. In fifth place is `home`, which was found to be used internally so frequently [1] that ICANN has indefinitely delayed its delegation due to the collision risk [21]. Some common public TLDs (`com`, `net`, `org`) are also used. This is partly because some RIPE Atlas users use their own domain name internally. In fact, 1146 names (37.06%) only occur once and might be unique to the probe's network. This is likely more common on RIPE Atlas than average, due to its more technical user base.

B. Current Collision Risk

1766 names (57.12%) have a TLD in the public DNS. We check how many of these names are at risk of name collision,

[4]We don't verify that the response comes from the gateway - this is challenging because some gateways appear to spoof the source address.

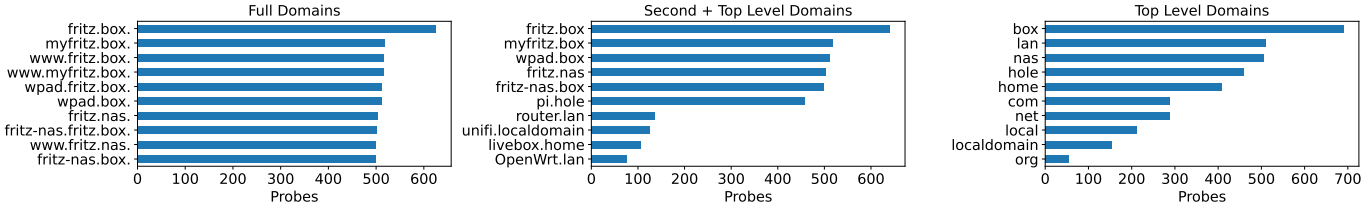

Figure 2. Top 10 internal full, second level and top-level domains.

i.e. are unregistered but registrable. We only consider potential name collisions and not ongoing name collisions: it is difficult to determine whether the internal and global name are owned by the same party (which would be an intentional name collision).

Subdomains of public suffixes [22] [23] are registrable by individuals. We thus extract the subdomain of the public suffix of the name and check if it resolves by querying for an SOA record. We don't check if the full domain name resolves because a domain owner might use a subdomain internally, e.g. the owner of `example.co.uk` might use `gw.example.co.uk` internally without adding it to the DNS. This name wouldn't resolve, but can't be registered; only the domain owner can add new subdomains. In this case we would check if `example.co.uk` (subdomain of the public suffix `co.uk`) resolves.

Out of the 1766 names with a public TLD, 1687 (95.53%) have a resolvable public suffix subdomain. 66 names (3.74%, 2.13% of *all* names) don't resolve and could be registered. We couldn't assess the collision risk for the remaining 13 names.

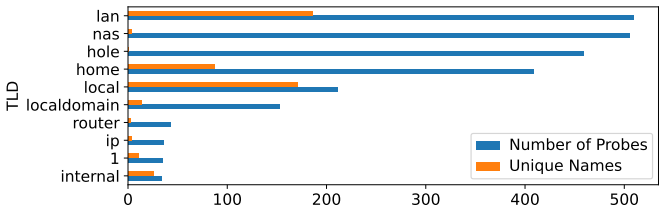

Figure 3. Top 10 undelegated TLDs.

C. Undelegated TLDs

1326 names (42.88%) use a TLD that's not in the public DNS. Of these internal names, 1067 (34.51% of all names) are *not* subdomains of special-use domain names [24], and are thus at risk of collision if their TLD is added to the DNS.

Figure 3 shows the top 10 non-delegated TLDs (including special-use names), by number of probes that use them. After `lan`, the most common TLD is `nas`, mostly from `fritz.nas` and `www.fritz.nas`. It is likely related to the Network Attached Storage feature of the FRITZ!Box. This TLD doesn't appear in the top results of past studies of invalid TLDs reaching root servers [2] [3] [1] or recursive resolvers [25]. This could be because these studies are over 10 years old, because RIPE Atlas has overproportionately many FRITZ!Boxes, or because these queries don't often reach the root servers. Regardless, this is another potential name collision for FRITZ!Box gateways.

The TLDs `lan`, `home`, `local`, `localdomain`, `router` and `internal` do appear in the results of some or all aforementioned studies [2] [3] [1] [25], suggesting more persistent use. The 26 names using `internal` and the 96 names using `home` are at lower risk of collision; ICANN has proposed reserving `internal` for internal use [26], and indefinitely delayed delegation of `home` [21]. However, `home` isn't a special-use domain name [24] designated for internal use. Special-use names aren't widely used: only 24 probes use the special-use alternative to `home` (`home.arpa` [27]), and only one top 10 TLD (`local`, for multicast DNS [17]) is a special-use name.

V. Related Work

Some past studies of root server and resolver logs discuss queries for internal names. ICANN commissioned [1] [10] to determine the collision risk of new gTLDs. Interisle Consulting Group [1] find a "substantial" collision potential, especially for `home` and `corp`. JAS Global Advisors [10] present a "controlled interruption" approach for safe delegation. Verisign [9] analyse which proposed new gTLDs are used internally, and quantify the risk of delegation. These studies focus on the collision risk of new gTLDs; to the best of our knowledge, ours is the first study of the collision risk of names used by home gateways. Other studies of root server [2] [3] and resolver logs [25] also find queries for invalid TLDs, including TLDs found by us.

Chen et al. [16] evaluate security risks of client-side name collisions by searching root server logs for common internal domain names, and analysing the services using these names.

Our study involves a form of gateway fingerprinting – other studies of home gateway fingerprinting use web interfaces [28] or port scanning [29]. Randall et al. [30] use `CHAOS TXT` queries to detect DNS interception by home gateways.

VI. Conclusions

We detect internal domain names used by RIPE Atlas probes and determine their name collision risk. We discover 3092 names, used by 4305 probes. Of these, 66 names are at risk of collision, 1067 names could be at risk if their TLD is delegated. Individuals hosting a RIPE Atlas probe are likely more technical than the average Internet user. It is thus unclear how representative our results are of internal names in other home networks. However, the large number of names with undelegated TLDs shows that name collisions in home networks warrant further study. In future work, we will increase the number of probes found through gateway fingerprinting, and explore ways to achieve more representative results.

References

[1] Interisle Consulting Group, "Name Collision in the DNS," Name Collision Study Report Version 1.5, Aug. 2013.

[2] D. Wessels and M. Fomenkov, "Wow, That's a Lot of Packets," in *Passive and Active Network Measurement Workshop (PAM)*, Apr. 2003.

[3] S. Castro, D. Wessels, M. Fomenkov, and K. Claffy, "A day at the root of the internet," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 5, pp. 41–46, Sep. 2008.

[4] "Box | ICANN New gTLDs," https://newgtlds.icann.org/en/program-status/sunrise-claims-periods/box.

[5] "Introducing .Box – A New Era in Domain Names Powered by 3DNS," https://3dns.box/blog/posts/introducing-box-tld/.

[6] Chainwire, "Introducing .box – The World's First Blockchain Native, DNS Routable Domain," https://decrypt.co/213372/introducing-box-the-worlds-first-blockchain-native-dns-routable-domain, Jan. 2024.

[7] ICANN, "About the Program | ICANN New gTLDs," https://newgtlds.icann.org/en/about/program.

[8] ——, "Program Statistics | ICANN New gTLDs," https://newgtlds.icann.org/en/program-status/statistics.

[9] Verisign Labs, "New gTLD Security, Stability, Resiliency Update: Exploratory Consumer Impact Analysis," Verisign Labs Technical Report 1130008 Version 1.1, Aug. 2013.

[10] JAS Global Advisors, "Mitigating the Risk of DNS Namespace Collisions - A Study on Namespace Collisions in the Global Internet DNS Namespace and a Framework for Risk Mitigation," Final Report, Oct. 2015.

[11] "RIPE Atlas," https://www.ripe.net/analyse/internet-measurements/atlas.

[12] P. Mockapetris, "Domain names - concepts and facilities," Internet Engineering Task Force, Request for Comments RFC 1034, Nov. 1987.

[13] ICANN, "Acronyms and Terms," https://www.icann.org/en/icann-acronyms-and-terms/internet-assigned-numbers-authority-en.

[14] "Root Zone Management," https://www.iana.org/domains/root.

[15] ICANN, "Name Collision Occurrence Management Framework," Jul. 2014.

[16] Q. A. Chen, M. Thomas, E. Osterweil, Y. Cao, J. You, and Z. M. Mao, "Client-side Name Collision Vulnerability in the New gTLD Era: A Systematic Study," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '17. New York, NY, USA: Association for Computing Machinery, Oct. 2017, pp. 941–956.

[17] S. Cheshire and M. Krochmal, "Multicast DNS," Internet Engineering Task Force, Request for Comments RFC 6762, Feb. 2013.

[18] Internet Systems Consortium, "BIND 9 Configuration Reference," https://bind9.readthedocs.io/en/latest/reference.html#built-in-server-information-zones, Apr. 2024.

[19] IANA, "IANA IPv4 Special-Purpose Address Registry," https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml.

[20] ——, "IANA IPv6 Special-Purpose Address Registry," https://www.iana.org/assignments/iana-ipv6-special-registry/iana-ipv6-special-registry.xhtml.

[21] ICANN, "Approved Board Resolutions | Regular Meeting of the ICANN Board," https://www.icann.org/en/board-activities-and-meetings/materials/approved-board-resolutions-regular-meeting-of-the-icann-board-04-02-2018-en#2.c.

[22] "Public Suffix List," https://publicsuffix.org/.

[23] S. McQuistin, P. Snyder, C. Perkins, H. Haddadi, and G. Tyson, "A First Look at the Privacy Harms of the Public Suffix List," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, ser. IMC '23. New York, NY, USA: Association for Computing Machinery, Oct. 2023, pp. 383–390.

[24] IANA, "Special-Use Domain Names," https://www.iana.org/assignments/special-use-domain-names/special-use-domain-names.xhtml.

[25] H. Gao, V. Yegneswaran, Y. Chen, P. Porras, S. Ghosh, J. Jiang, and H. Duan, "An empirical reexamination of global DNS behavior," in *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: Association for Computing Machinery, Aug. 2013, pp. 267–278.

[26] "ICANN Seeks Feedback on Proposed Top-Level Domain String for Private Use," https://www.icann.org/en/announcements/details/icann-seeks-feedback-on-proposed-top-level-domain-string-for-private-use-24-01-2024-en.

[27] P. Pfister and T. Lemon, "Special-Use Domain 'home.arpa.'," Internet Engineering Task Force, Request for Comments RFC 8375, May 2018.

[28] M. Niemietz and J. Schwenk, "Owning Your Home Network: Router Security Revisited," in *Web 2.0 Security and Privacy Workshop*, San Jose, CA, May 2015.

[29] T. Papastergiou, R. Perdisci, and M. Antonakakis, "Returning to Port: Efficient Detection of Home Router Devices," in *2022 IEEE Conference on Communications and Network Security (CNS)*, Oct. 2022, pp. 172–180.

[30] A. Randall, E. Liu, R. Padmanabhan, G. Akiwate, G. M. Voelker, S. Savage, and A. Schulman, "Home is where the hijacking is: Understanding DNS interception by residential routers," in *Proceedings of the 21st ACM Internet Measurement Conference*, ser. IMC '21. New York, NY, USA: Association for Computing Machinery, Nov. 2021, pp. 390–397.