Investigating Deployment Issues of DNS Root Server Instances From a China-Wide View

Fenglu Zhang, *Member, IEEE*, Baojun Liu, *Member, IEEE*, Chaoyi Lu, Yunpeng Xing, Haixin Duan, *Member, IEEE*, Ying Liu, *Member, IEEE*, and Liyuan Chang

***Abstract*—DNS root servers are the starting point of most DNS queries. To ensure their security and stability, multiple anycast instances are operated worldwide, and new root instances have been rapidly deployed in recent years. Apart from authorized instances managed by Root Server System, some networks equip unauthorized instances to hijack queries from clients. Despite various root instances handling queries within their residing networks, few studies have focused on the deployment issues of these instances. In this article, we provide the first study to reveal the deployment issues of root instances from a nationwide view. With the support of 7,860 vantage points, we utilized a suite of methodologies to identify the deployment of unauthorized instances. 54 vantage points witnessed the evidence of unauthorized instances, and 70.4% of them further observed security issues of unauthorized instances, including DoS, unavailability of DNSSEC validation, and vulnerable DNS software. Additionally, we utilized the side-channel information of censorship mechanisms to measure the catchment area of authorized instances. We found that most authorized instances in the Chinese mainland serve with limited catchment areas due to restricted BGP policies. Through discussions with ISPs and network operators, we make recommendations to improve the deployment status of different root instances.**

***Index Terms*—Domain name system, DNS security, root server instances, Internet measurement.**

I. Introduction

DOMAIN Name System (DNS), a cornerstone of Internet infrastructure, not only enables maps from domains to IP addresses but also anchors modern security mechanisms [2], [3], [4], [5]. Recent studies have uncovered that adversaries can exploit DNS to acquire fraudulent CA certificates, thereby compromising upper-layer services [6], [7], [8], [9]. Also, historical DNS outages have been demonstrated to cause widespread interruption of internet services [10], [11]. Given these risks, ensuring the *security and stability* of DNS is significant for global Internet operations.

As the starting point of the whole domain space, the DNS root takes control of resolving the vast majority of domain names. To resist denial-of-service (DoS) attacks and improve stability, each of the 13 DNS root servers is deployed using anycast with multiple *root server instances* that act collectively behind a shared IP addresses [12]. A query toward the DNS root server is routed to one root instance, preferably the closest to its origin [13]. Recently, root instances have been deployed rapidly to provide faster and more reliable service [14]. At the time of writing, 1,627 root instances are deployed worldwide, which is 57.5% more compared to the number three years ago [15].

Depending on routing policies, multiple types of root instances are operating on the Internet. DNS Root Server System (RSS) comprises both *global* and *local* root instances. Global instances should provide service to worldwide users since their route advertisement can spread throughout the Internet [15], [16]. By contrast, local instances only serve a limited range of networks due to the limited BGP policy [13], [15], [17]. Apart from the *authorized* root instances mentioned before, network operators may also establish *unauthorized* root instances that are set up outside RSS and can hijack root queries within the networks they reside in.

Unauthorized instances can lead to severe security risks and operational issues [9], [18]. As unauthorized instances control the whole DNS namespace, they have the capability to manipulate the resolution of arbitrary domain names. By manipulating DNS resolution, they can further manipulate the subsequent traffic (e.g., web and email traffic). Unauthorized instances can also disrupt security mechanisms such as DNSSEC, which ensures data integrity and authentication during DNS resolution. Moreover, due to the lack of supervision from RSS, some unauthorized instances may suffer from poor maintenance. The operational issues, such as utilizing vulnerable software, can cause severe consequences.

Research gap: The deployment of diverse root instances significantly affects the security of domain resolutions within their serving networks. Given the significance, though, few

Manuscript received 1 February 2023; revised 14 February 2024; accepted 27 February 2024. Date of publication 5 March 2024; date of current version 13 November 2024. This research was supported in part by the National Key Research and Development Program of China under Grant 2023YFB3105600, in part by the National Natural Science Foundation of China under Grant 62102218, Grant U1836213, Grant U19B2034, and Grant 62302258, in part by the Alibaba Innovative Research Program (AIR), CCF-Tencent Rhino-Bird Young Faculty Open Research Fund under Grant CCF-Tencent RAGR20230116. The work of Haixin Duan was supported in part by the Taishan Scholars Program. An earlier version of this article was presented at Passive and Active Measurement Conference 2022 (PAM 2022) [DOI: 10.1007/978-3-030-98785-5_11]. *(Corresponding author: Baojun Liu.)*

Fenglu Zhang, Baojun Liu, Chaoyi Lu, and Yunpeng Xing are with the Institute for Network Sciences and Cyberspace, Tsinghua University, Beijing 100084, China (e-mail: zfl23@mails.tsinghua.edu.cn; lbj@tsinghua.edu.cn; luchaoyi@tsinghua.edu.cn; xyp23@mails.tsinghua.edu.cn).

Haixin Duan and Ying Liu are with the Institute for Network Sciences and Cyberspace, Tsinghua University, Beijing 100084, China, and also with Zhongguancun Laboratory, Beijing 100190, China (e-mail: duanhx@tsinghua.edu.cn; liuying@cernet.edu.cn).

Liyuan Chang is with the Innovation Department, China Telecom Network Security Technology Company Ltd., Beijing 100020, China (e-mail: changly@chinatelecom.cn).

Digital Object Identifier 10.1109/TDSC.2024.3373530

existing studies have investigated the deployment issue of these instances. Some research papers have focused on measuring the latency [19], [20], [21] and the health status [22], [23], [24] in the scope of root servers (e.g., measuring the global-wide latency to A-Root). However, there is still a lack of insights into an individual root instance (e.g., the catchment area of a specified A-Root instance). Besides, previous studies have only discovered a few unauthorized instances [9], [25] but have not explored the security issues caused by these instances.

Our study: To fill the research gap, we try to answer four research questions about deployment issues associated with two types of root instances. *Regarding unauthorized instances:*① Which network deploys them? ② What security issues do they cause? *Regarding authorized instances:*③ Why do they not serve some nearby networks? ④ How do they absorb queries from nearby recursive resolvers? We believe that seeking answers to the above questions can help improve the security and robustness of root instances.

Answering the above questions is not straightforward and requires addressing three main challenges. First, analyzing unauthorized instances is challenging with supervised learning since ground truth is not available. This challenge is partly due to DNS specifications [18], [26], [27], [28], [29] that strictly prohibit the manipulations of security mechanisms or the zone files of the DNS root servers. As a result, to our knowledge, there have been no announcements of unauthorized instance deployments, which can be utilized as ground truth. Second, we need to locate the exact root instance sharing the same anycast IP address. This is not trivial work since information such as hosting networks and peering ASes of root instances is private. Eight years ago, a study targeted B-Root, which did not activate anycast at that time [9]. However, we cannot reuse this method since all root servers have been deployed using anycast [30]. Third, no out-of-box solution provides vantage points satisfying the requirements of this study. Specifically, our vantage points require sending special DNS queries to identify manipulations that remain transparent to users. Also, they should cover a wide range of geolocations and ISPs to detect cases where an unauthorized instance exclusively manipulates resolutions within a specified network. Most proxy services and measurement platforms utilized by previous studies do not satisfy our requirements [9], [31], [32].

In this study, we utilized a suite of novel methodologies and conducted a comprehensive measurement to reveal the deployment issues of root instances. As a first step forward, we performed a case study in the Chinese mainland,[1] an understudied region with a large Internet population. To support our research, we developed a measurement platform from scratch and collected 7,860 vantage points from 17 commercial proxy vendors supporting forwarding special DNS queries. Utilizing our measurement platform, we first identified the deployment of unauthorized instances by conducting DNS traceroute and requesting server identifiers (R.Q. ①). Furthermore, we examined the integrity and timeliness of the root zone file, along with the availability of DNSSEC validation, to uncover security issues of observed unauthorized instances (R.Q. ②). Taking advantage of the side-channel information from DNS censorship mechanisms, we then measured the catchment area of authorized instances in the Chinese mainland. We shed light on the reason behind the restricted catchment area (R.Q. ③). Besides, we investigated how root instances impact the root server selection of resolvers and further absorb the queries from nearby resolvers through software analysis and measurement from the perspective of resolvers (R.Q. ④).

Our findings: Our study reveals *security* and *stability* issues arising from the current deployment of root server instances in the Chinese mainland.

We first proved that the unauthorized instances deployed in the Chinese mainland can cause severe security issues. Overall, 54 vantage points reported evidence of deploying unauthorized instances, including both stable unauthorized instances and instances exhibiting random behaviors of root manipulation. **70.4%** of them further observed the security issues of the unauthorized instances. In the network of China Telecom (Zunyi City), we observed unauthorized instances that failed to resolve all tested domain names and facilitated DoS issues. Besides, DNSSEC validation, which provided integrity checking and authentication, was unavailable to 48.1% of vantage points served by unauthorized instances. We also discovered that the unauthorized instances in CERNET [33] (China Education and Research Network) equipped vulnerable software and outdated root zone files. Even worse, such a vulnerable DNS software was released ten years ago and could be affected by 100 CVE vulnerabilities. As a point of comparison, our experiment did not observe security issues, including DoS attacks, use of vulnerable software, or outdated zones, among the 5,489 vantage points served by authorized instances.

We then uncovered that some authorized instances in the Chinese mainland served in limited catchment areas, even if they were global instances, and should serve the whole Internet. In particular, our results show that some authorized instances were not accessible from major ISP networks due to the limitation of BGP routing policies. For example, an I-Root global instance in Beijing cannot be accessed from all major commercial ISP networks in the Chinese mainland. Consequently, the majority of queries toward root servers would turn to overseas root instances despite closer domestic root instances operating in the Chinese mainland. Besides, we confirmed that deploying authorized instances in domestic contributed to lowering the latency of their corresponding root server, and algorithms of DNS software would thus prefer such a root server. As a result, the domestic instances can further absorb queries from nearby recursive resolvers. This result provides a guideline for the future deployment of authorized instances.

Through in-depth discussions with ISPs and network operators, we make recommendations for multiple parties to help improve the practical effect of root instances. We also make our artifacts publicly available, including tools, measurement data, and a technical report of analyzed DNS software.[2]

Our contributions:

[1]Due to different network policies, we excluded Hong Kong, Macao, and Taiwan from the scope of our study.

[2]Our artifacts are available at https://github.com/zhangshanfen9/ idiori

TABLE I
ROOT SERVER INSTANCES DEPLOYED IN THE CHINESE MAINLAND

Root	Global	Local	Location
A	1	0	Beijing
F	0	4	Beijing, Chongqing, Hangzhou, Xining
I	1	0	Beijing
J	2	0	Beijing, Hangzhou
K	0	3	Beijing, Guangzhou, Guiyang
L	13	0	Beijing(4), Changsha, Haikou, Shanghai, Wuhan(2), Xining(2), Zhengzhou(2)
Total	17	7	

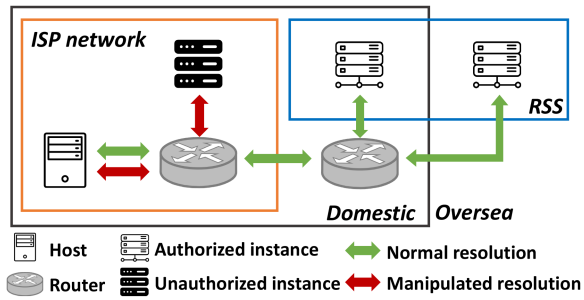

Fig. 1. Threat model of unauthorized instances.

- We provided a suite of novel methodologies and conducted an in-depth case study in China to detect the deployment issues of root server instances.
- We shed light on the security issues of unauthorized instances, including DoS, unavailability of DNSSEC validation, and vulnerable DNS software.
- We revealed that even global authorized instances can serve limited catchment areas due to the restriction of BGP policies.

II. BACKGROUND

This section provides background information about DNS root server system, unauthorized root server instances, and nation-scale DNS manipulation.

A. DNS Root Server System (RSS)

DNS Root Server and Root Instance: DNS namespace is organized as a hierarchical structure, and DNS root servers provide eventual access to the whole domain space. Specifically, a DNS resolution process typically starts when a client requests a recursive resolver to resolve a domain (e.g., *example.com.*). Following the hierarchical structure of domain name space, the recursive resolver turns to query a root server (.), the top-level domain (TLD) server (*com.*), and the authoritative server (*example.com.*), respectively [34]. Due to early payload size limits, there are only 13 root servers in the RSS [35], which are named by `A-Root` through `M-Root`. The 12 Root Server Operators (RSO) administer 13 root servers (Verisign operates two identities for historical reasons). All root servers in the RSS serve one individual copy of the DNS root zone managed by IANA [36].

To resist DoS attacks and improve the stability of DNS, the current deployment of root servers relies on anycast [12], which enables multiple *root instances* to operate behind an IP address of a root server collectively. As a result, a query toward a root server can be routed to a nearby root instance among a group of instances sharing the same IP address. At the time of writing, RSS is operating 1,627 instances in the world [15]. In the Chinese mainland, 24 instances have been deployed, and Table I shows their locations [15].

Catchment Area of a Root Instance: RSS includes both *global* and *local* root instances, which vary in accessibility across different networks [15]. Specifically, the route advertisement of a local instance is limited to nearby networks, with the catchment area restricted to the hosting ASes or the boundaries of the BGP confederation [13], [17]. However, a global instance should serve the entire Internet. This is because its route advertisement is permitted to spread throughout the Internet, and any router on the Internet could know the path to that instance [15], [16]. As presented in Table I, the Chinese mainland hosts 17 global and seven local instances.

B. Security Risks of Unauthorized Root Instances

Except for authorized instances operated by RSS, some networks may host unauthorized instances that manipulate DNS queries toward root servers. A depiction of the threat model is shown in Fig. 1. Typically, DNS queries toward a root server should be directed to an authorized instance belonging to the RSS, both in cases where the instances are domestic or overseas. However, some network (e.g., an ISP network) operators may deploy unauthorized instances within their networks. By injecting forged DNS responses or manipulating IP routing, such unauthorized instances respond to the queries toward root servers and further control the entire DNS namespace of users within their catchment area.

Unauthorized instances may pose serious security risks. ICANN RSSAC (Root Server System Advisory Committee) [18] has discussed the potential security risks of rogue root server operators. Precisely, such operators may arbitrarily manipulate the response data in DNS root zone, which controls resolutions of all domain names. Consequently, they can block access to any website by disrupting the corresponding domain resolution and even redirect the subsequent traffic (e.g., web and email traffic) [9]. Such attacks can even circumvent the security mechanisms designed for upper-layer applications. For instance, through DNS hijacking or manipulation, attackers are capable of acquiring fraudulent HTTPS certificates, as proven in previous studies [6], [7], [8]. Then, attackers can further bypass a series of defense mechanisms for web services, including the Same Origin Policy [37], [38] and HTTP Strict Transport Security (HSTS) [39], thereby simplifying attacks such as web phishing. Alternatively, rogue root server operators can also hinder the requests for DNSSEC validation [40], which ensures data integrity and authentication in a domain resolution process. Although it is unlikely for an official root server operator to engage in rogue behavior, the operator of unauthorized instances

that fall outside the restriction of RSS, is more likely to pose the above negative impact.

Also, deploying unauthorized instances violates the requirement of a unique DNS root and could cause operational issues [41]. Generally, the root server operators ensure the coordination of updates on all authorized instances. Such updates include the latest DNS software and zone file to protect against known vulnerabilities and potential outages. However, this is not the case for operators of unauthorized instances. Some unauthorized instances may be in poor maintenance and cause operational issues, such as utilizing vulnerable DNS software or failing to serve.

In this study, we conducted an in-depth analysis of unauthorized instances and confirmed the *real-world security issues posed by unauthorized instances:* The identified security issues include DoS, unavailability of DNSSEC validation, and vulnerable DNS software (Section IV).

C. Nation-Scale DNS Manipulation

Some countries enable DNS censorship and perform nation-scale DNS manipulation since DNS lacks data integrity and authentication. Previous studies have demonstrated that countries (e.g., China [42], [43], [44], [45], Russia [46], [47], India [48], Pakistan [49], and Italy [50]) deploying DNS censorship share a large proportion of the world population.

One efficient censorship approach is censoring at the choke points of the network topology. Such choke points are located at the border of a network (e.g., the network of an ISP or even a country). Once detecting a DNS request for a censored domain crossing the border, the censorship system immediately injects a forged response (e.g., with a blackhole IP address), which prevents the client from the legitimate answer [45], [51], [52]. In this study, we observed that such a censorship mechanism was widely deployed and coincidentally provided side-channel information to distinguish whether DNS queries cross the border of a network. Such a characteristic can further help us infer the catchment area of root instances (Section V).

III. Vantage Point

In this section, we introduce the challenges of collecting vantage points to support our study. Subsequently, we elaborate on our approach to addressing the challenges, which involves developing a measurement platform for vantage point collection.

Challenge: Our study targets root instances serving diverse catchment areas and explores the associated security issues. To this end, we need to tackle three challenges:

(1) Collecting vantage points with the capability of sending special DNS packets: To conduct our measurement methods (Sections IV-A and V-A), vantage points are expected to probe with special DNS packets. In particular, they should support DNS queries with a modified TTL field to conduct DNS traceroute. Moreover, the requests for server identifiers and DNSSEC require the support of `EDNS0` [53].

(2) Collecting vantage points with wide geographical coverage: Previous studies have shown that the scope of DNS manipulation can be highly related to the region [9], [54].

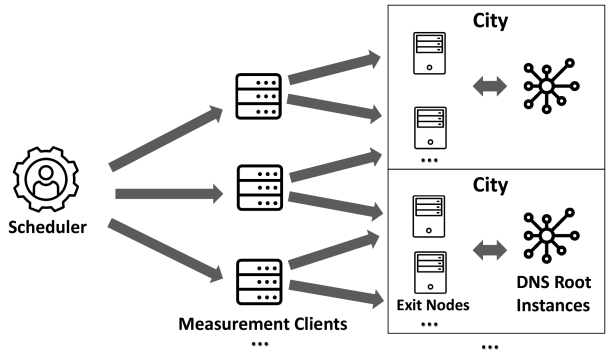

Fig. 2. Architecture of our measurement platform.

Moreover, the situation can be more complex in China since ISPs manage their networks at the provincial level, leading to potential variations in network policies across provinces. Consequently, wide geographical coverage of vantage points is necessary to uncover geo-distributed unauthorized instances.

(3) Ensuring vantage points are representative of real traffic: A main challenge in our study was collecting real DNS traffic from vantage points. Previous studies have revealed that vantage points may not be representative of real traffic. Specifically, vantage points may suffer from DNS hijacking (e.g., `NXDOMAIN` rewriting by a middlebox) [55], [56], thereby producing false-positive results. Some traffic manipulation can even affect traffic for upper-layer applications (e.g., web traffic) [57], [58], [59]. Also, the locations advertised by vendors cannot be relied on and may lead to skewed results [60], [61].

To our knowledge, no out-of-box solution fulfills the requirements of this research. Previous studies [9], [31], [32] utilized proxy networks (e.g., ProxyRack [62]) and measurement platforms (e.g., RIPE Atlas [63]) to collect vantage points. However, their vantage points suffer from poor geographical coverage in China and a lack of supporting special DNS queries. For instance, RIPE Atlas [63] maintains only 67 vantage points in China and does not provide DNS traceroute functionality.

Our solution: To tackle the above challenges, we collected vantage points from VPN vendors. Specifically, we selected 17 commercial VPN vendors in China and utilized their proxy services to collect vantage points. To set up proxy services, such vendors organize both dedicated servers in data centers and end-user devices as their exit nodes, covering all provinces and major ISPs in the Chinese mainland. Also, they allow customers to access the exit nodes through L2TP [64] and PPTP [65], which are VPN protocols supporting forwarding special DNS queries. As a result, the exit nodes of VPN services satisfied our requirements for vantage points.

We built a measurement platform with a distributed architecture from scratch to ensure efficient measurement, as shown in Fig. 2. Specifically, we equipped 50 measurement clients for our platform. Each measurement client was an independent machine running an operational system. Under the control of the scheduler, clients conducted measurement tasks simultaneously. With ethical considerations, we equipped an independent account for every measurement client to acquire a proxy service.

TABLE II
COUNT AND COVERAGE OF SELECTED VANTAGE POINTS

ISP	# VP	# BGP Prefix	# Covered Provinces
China Telecom [67]	3313	327	29
China Unicom [68]	3892	118	18
China Mobile [69]	460	30	13
Tencent Cloud [70]	78	44	6
Alibaba Cloud [71]	106	36	7
CERNET [33]	11	10	8
Total	7860	565	31/31

The proxy service provided a series of exit nodes, which helped measurement clients access to the root instances in different cities.

We also conducted dedicated validations to ensure vantage points collect real DNS traffic before our measurement. For each vantage point, we checked whether the related network was involved in unintended DNS hijacking and verified the announced geo-location. In particular, we will utilize a DNS censorship mechanism deployed on the border of a country to infer the catchment area of root instances (Section V-A). However, some entities (e.g., a domestic middlebox) may also inspect users' DNS queries and hijack the corresponding traffic, which can generate false-positive results. To filter out such unintended DNS hijacking, a measurement client first sent DNS queries with censored domains to five IP addresses that were distributed nationwide and did NOT provide a DNS service. Typically, the queries would time out. A measurement client would skip an exit node if the exit node received DNS responses, indicating the DNS hijacking of an entity. Furthermore, a measurement client launched DNS queries from each exit node to our custom DNS server, which reported the source addresses of incoming queries against the MaxMind database [66]. A measurement client would skip an exit node if the location did not match what was advertised by vendors. Our validations didn't detect traffic manipulation toward vantage points for upper-layer applications (e.g., web) since they didn't affect our measurement of DNS root instances.

In total, we obtained **7,860** exit nodes as our vantage points that were unaffected by the hijacking of domestic middleboxes and advertised the correct location. After passing validations, a measurement client executed tasks elaborated in Sections IV and V, respectively. Then, it requested switching to a new exit node. The above process was repeated until the running time reached 60 days. As shown in Table II, our vantage points covered all three major ISPs, two major cloud services, and all 31 provinces in the Chinese mainland. We also included vantage points in CERNET, which exclusively serves universities and research facilities. Each vantage point supported sending special DNS queries to arbitrary IP addresses.

IV. SECURITY ISSUES OF UNAUTHORIZED ROOT INSTANCES

In this section, we elaborate on uncovering security issues of unauthorized instances, which can affect any DNS resolutions of users within their catchment area. Our national-wide analysis targeted China, which has a large population on the Internet.

Challenge: Identifying unauthorized instances is not trivial work since three challenges need to be addressed:

(1) Absence of ground truth: The detection of unauthorized instances is incompatible with supervised learning due to the absence of ground truth. This limitation may be partially attributed to DNS specifications, which strictly forbid manipulations related to security mechanisms or the zone files of the DNS root servers [18], [26], [27], [28], [29]. To our knowledge, no deployment of unauthorized instances has been announced yet. Consequently, the development of a heuristic-based approach becomes necessary.

(2) Transparent manipulation: Network operators can transparently redirect users' traffic to unauthorized instances. Hence, we require methods of confirming identities of unauthorized instances, which may be covered by injected responses or manipulation of IP routing [9].

(3) Exclusive manipulation: An unauthorized instance can exclusively serve a specific network in a particular location (e.g., an ISP network in a province) [9], [54], [72]. Consequently, our vantage points require covering enough geolocations and networks to detect such manipulation.

Our solution: We utilized the concept of *similarity* among authorized instances to address the above addresses, with the support of nationwide covered vantage points. As official root server operators uniformly manage authorized instances, the penultimate hops and server identifiers can exhibit similar characteristics, such as configuring identifiers following the same naming convention. As a result, we can collect the normal characteristics of authorized instances through nationwide covered vantage points and find the unauthorized instances violating similar characteristics. For example, all authorized `G-root` instances follow the server identifier format of *"[id].groot"*, which adheres to the naming convention of the official root operator. Hence, we can identify the unauthorized `G-root` instance if a vantage point observes a server identifier violating the naming convention. Also, we can confirm the deployment of unauthorized `E-root` instances when a vantage point reports a penultimate hop to `E-root` since the official announcement [15] states that no `E-root` instances deployed in the Chinese mainland.

This section proposes a suite of methodologies to identify unauthorized instances and uncover potential security issues (Section IV-A). Based on the collected information, we identified the deployment of unauthorized instances (Section IV-B) and further analyzed the related security issues (Section IV-C). We also analyzed the security guarantees of authorized instances to highlight the risks posed by unauthorized instances (Section IV-D).

A. Methodology

Identifying the deployment of unauthorized instances: Our vantage points collected two characteristics from each root server and leveraged the concept of similarity to detect unauthorized instances. Specifically, a vantage point first collected the following characteristics:

TABLE III
NORMAL CASES OF PENULTIMATE HOP LOCATIONS THAT MATCH THE OFFICIAL ANNOUNCEMENT

Root	A	B	C	D	E	F	G	H	I	J	K	L	M
Loc.	Timeout 90.5% Oversea 9.5%	Timeout 98.8% Oversea 1.2%	Oversea 62.0% Timeout 38.0%	Oversea 71.5% Timeout 28.5%	Oversea 52.1% Timeout 47.9%	Beijing 49.5% Timeout 10.2% Oversea 40.5% Chongqing 2.8%	Timeout 100%	Timeout 100%	Oversea 85.8% Timeout 10.5% Shenyang 3.7%	Timeout 89.5% Oversea 10.2% Beijing 0.3%	Timeout 53.1% Guangzhou 42.2% Oversea 3.5% Beijing 1.2%	Beijing 48.5% Timeout 32.7% Wuhan 16.7% Yichang 1.2% Shanghai 0.9%	Timeout 50.3% Oversea 49.7%

(1) Penultimate hops probed by DNS traceroute: Through DNS traceroute, we collected the penultimate hop to each root server, obtaining information about the geolocation and network of a specific root instance. Unlike traditional traceroute carried out using ICMP, TCP, and UDP. DNS traceroute is a more reliable technology for displaying the path to a root instance, as it utilizes normal DNS queries to carry the probing packets. Specifically, our vantage points sent a series of DNS queries to the root server with a modified TTL field in the IP header [72]. Hence, we can display the path to unauthorized instances that may only manipulate DNS traffic rather than other traffic. Our vantage points conducted DNS traceroutes toward each root server five times. Then, we extracted the penultimate hops since they are the closest to the probed root instance (the last hop is an anycast IP address of the root server).

(2) Server identifiers: A server identifier of a root instance is configured by the operator and contains the information for identification. To differentiate nameservers activating anycast or load balancing, RFCs [73], [74] specify special queries for server identifiers, including `NSID` and `ID.SERVER`. Besides, the `VERSION.BIND` identifier provides the software version of the nameserver [74]. Some identifiers even provide the geolocation of a root instance. For example, the identifier "*b1-sin*" represents the B-root instance in Singapore. In our experiment, our vantage points requested three types of server identifiers from each root server.

We further utilized the concept of *similarity* among authorized instances to confirm the deployment of unauthorized instances. The reason is that all authorized instances are under the management of the RSS and RSO, resulting in similar characteristics. Such characteristics include the geolocation of the penultimate hop in traceroute paths and the uniform identifier naming convention. RFC 7108 also specifies the meaning of identifier naming conventions [75]. Specifically, all penultimate hops in traceroute paths to the same root instance can be in the same city, the same subnet, or even the same IP address. Also, server identifiers of a root server can be represented by a regular expression, which is a naming convention ruled by its root server operator. Hence, we can find unauthorized instances that do not adhere to the similarity of authorized instances.

Through nationwide covered vantage points, we collected normal cases of penultimate hop locations and server identifiers for further identification, as shown in Tables III and IV. We confirmed that geolocations of penultimate hops align with the official announcement [15]. For example, the official announcement states that two authorized `F-Root` instances are deployed in Beijing and Chongqing. Our result also shows 49.5% and 2.8% of penultimate hops are located in the two cities,

TABLE IV
REGULAR EXPRESSIONS OF SERVER IDENTIFIERS COLLECTED BY VANTAGE POINTS

Root	NSID	ID.SERVER	VERSION.BIND
A	\w{2,3}\.\w{2}-\w{3}\.root	rootns | nnn1-\w+-?\w+	NSD | ATLAS
B	\w{2}-\w{3}	\w{2}-\w{3}	knot 2.x
C	\w{5}.c.root-servers.org	\w{5}.c.root-servers.org	C-root
D	\w{5}.droot.maxgigapop.net	\w{5}.droot.maxgigapop.net	NSD 4
E	c01.\w{3}.eroot	c01.\w{3}.eroot	cloudflare-e-root-20190930 9.16.27 (53.4%) 9.16.33 (2.5%)
F	.+.f.root-servers.org	.+.f.root-servers.org	cloudflare-f-root-20190930 (44.1%)
G	\w{1,4}.groot	(REFUSED)	(REFUSED)
H	001.\w{3}.h.root-servers.org	001.\w{3}.h.root-servers.org	NSD 4.5.0
I	s1.\w{3}	s1.\w{3}	contact info@netnod.se
J	\w{2,3}\.\w{2}-\w{3}\.root	rootns | nnn1-\w+-?\w+	NSD | ATLAS
K	ns\d.\w{2}-\w{3}.k.ripe.net	ns\d.\w{2}-\w{3}.k.ripe.net	Knot DNS | BIND | NSD
L	cn-\w{3}-\w{2}	cn-\w{3}-\w{2}	Knot DNS 3 | NSD 4
M	M-NRT-JPNAP-\d	M-NRT-JPNAP-\d	9.16

respectively. We can also extract regular expressions of server identifiers to represent the uniform naming convention of a root server operator. Except for `A-Root` and `J-Root` administered by the same operator (Verisign), the naming conventions of other root servers are different from each other. As a result, we can confirm the deployment of unauthorized instances which violates normal cases of the measured characteristics and conducts further security analysis.

Uncovering the security issues: Once we identified an authorized instance, a vantage point then checked three metrics to expose the corresponding security issues. In particular, our vantage point collected the metrics as follows:

(1) Integrity of the zone file: We detected the manipulation of an unauthorized instance's root zone by examining its zone file's integrity. To verify the integrity of the zone file, we examined the TLD referral information provided by unauthorized instances, as the zone file defines the responses of a nameserver [76]. For this examination, we queried three gTLDs (*com*, *net*, *org*), and four ccTLDs (*us*, *uk* , *cn*, *ru*) from each root server. Subsequently, we compared these responses with the accurate ones presented in the official root zone file [77].

(2) Timeliness of the root zone file: We investigated whether an unauthorized instance updated its zone file promptly to prevent an outage caused by changing TLD nameservers. According to RSSAC requirements, root instances must sync their zone copies with the official zone file in time [24]. Each version of the root zone file is assigned a serial number, which is provided in a `SOA` record [76]. Our vantage points requested the `SOA` record from each root server and subsequently examined whether a root instance was using an outdated root zone file.

(3) Availability of DNSSEC validation: We confirmed whether unauthorized instances supported DNSSEC validation since they served as the initial step for DNSSEC validation within their catchment area [40]. From our vantage point, we sent three

Fig. 3. Heatmap displaying vantage points witnessed unauthorized instances.

TABLE V
COUNT OF VANTAGE POINTS WITNESSED EVIDENCE OF UNAUTHORIZED INSTANCES

Evidence	# VP
Traceroute Paths	28
Server Identifiers	12
Integrity of Root Zone File	25
Total	54

queries with the DO flag [78] to each root server. To mitigate the impact of network jitter, we only considered DNSSEC validation for DNS roots was not available when all queries failed. Also, we verified the authenticity of the DNSKEY and RRSIG records provided by root instances.

B. Overview of Unauthorized Root Instances

In this section, we provide an overview of observed unauthorized instances. Our identification of unauthorized instances started on Sep 4, 2022, and spanned 60 days. We witnessed evidence of unauthorized instances from 54 (0.7%) vantage points out of 7,860 vantage points. Fig. 3 illustrates the locations and numbers of these instances, while Table V provides detailed information about them. Based on these findings, we conducted three case studies:

Case study I. CERNET: We observed that *all* vantage points within CERNET reported abnormal penultimate hops and server identifiers, suggesting the presence of unauthorized instances that manipulate requests to all root servers.

We found all responsive penultimate hops to root servers are within Chinese cities and the network of CERNET, even during the probing of root servers (e.g., B-Root) without deploying authorized instances in the Chinese mainland. We present the result of DNS traceroute in Table VI. Interestingly, unauthorized instances in CERNET activated a load balancing mechanism while manipulating the requests toward different root servers. For instance, the unauthorized instance associated with

TABLE VI
RESULTS OF DNS TRACEROUTE IN CERNET

Final hop (root server)	Penultimate hop	Location and ISP	# VP
All roots	Timeout	-	8
CFG	101.4.113.53		7
ADEIJKLM	101.4.117.134		7
BFGH	101.4.113.234	Beijing, China	6
JKM	101.4.117.177	CERNET	2
ABCDEFGH	101.4.113.213		1
IJKM	101.4.116.158		1
ABDEHIJKLM	202.115.255.1	Chengdu, China CERNET	1

```
$ dig @b.root-servers.org +nsid
...
;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
; NSID: 62 31 2d 73 69 6e ("b1-sin")
...
;; ADDITIONAL SECTION:
a.root-servers.net. 518400 IN A ...
...
;; Query time: 130 msec
```

```
$ dig @b.root-servers.org +nsid
...
;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232

...
;; ADDITIONAL SECTION:
a.root-servers.net. 3600000 IN A ...
...
;; Query time: 1 msec
```

Fig. 4. After requesting the NSID identifier, our vantage points received the expected response from an authorized instance (left) and the abnormal response from an unauthorized instance in CERNET (right).

101.4.113.53 manipulated queries toward C, F, and G-Root, while the instance behind *101.4.117.177* was responsible for traffic towards J, K, and M-Root.

Our vantage points in CERNET also reported abnormal server identifiers, which violated the naming convention for server identifiers. Fig. 4 compares an expected response and an abnormal response from CERNET. The abnormal response lacked the NSID field, and the TTL of resource records in the additional section was significantly longer than in the expected case. Additionally, the query latency in CERNET was extremely low (1 ms). Such a result indicates that the response originated from a nearby unauthorized instance in domestic rather than an authorized instance overseas since no B-Root authorized instance was deployed in the Chinese mainland. While requesting ID.SERVER identifiers, we received responses containing an SOA record in the authority section, which differs from the expected response that includes a TXT record in the answer section. Moreover, the majority of vantage points in CERNET received VERSION.BIND responses with empty TXT records, while one vantage point even observed a version of vulnerable DNS software. We further analyze this security issue in Section IV-C.

While the comprehensive hijacking within an ISP may appear unexpected, validations from CERNET's operators and a prior research study affirm our observations. Through discussion, the operators of CERNET acknowledged the deployment of unauthorized instances and provided their motivation: These unauthorized instances were deployed in a period when the Chinese mainland lacked any authorized instances, serving a role in speeding up root queries within CERNET. Moreover, an earlier study [9] proved the deployment of unauthorized instances in CERNET as early as 2016. Hence, we believe these insights affirm the credibility of our result.

Case study II. Network of China Telecom (Zunyi City): Our vantage points in the network of China Telecom reported an exclusive manipulation of unauthorized instances. This manipulation exclusively targeted the vantage points at Zunyi, a city in southwest China. The observed unauthorized instances only manipulated traffic toward `B, C, F, G, K, L` and `M-Root` instead of all root servers. Moreover, they generated abnormal referrals for a particular TLD and did not resolve any other domain name. We provide a detailed analysis of this security issue in Section IV-C.

Case study III. Random behaviors of root manipulation: We observed some unauthorized instances exhibited random behaviors of root manipulation instead of consistently manipulating clients' DNS queries. Consequently, our vantage points reported only a subset of responses affected by manipulation, even when conducting the same type of measurement. We discuss the cases as follows:

(1) Penultimate hops in a LAN (local area network: In the network of Alibaba Cloud, 16 vantage points reported that the IP addresses of penultimate hops were within a private network (e.g., *10.102.50.9*), and only the paths to `J-Root` were affected. Notably, our vantage points consisted of dedicated servers in data centers or end-user devices, which were not likely the penultimate hops of authorized instances. The result indicates that an unauthorized instance within the LAN served our vantage points.

(2) Domestic penultimate hop to a root server without domestic authorized instance: While conducting DNS traceroute to `E-Root`, a vantage point in China Telecom identified the penultimate hops in the Chinese mainland (Chongqing City). Ten months after the experiment, the official announcement [15] still did not mention the presence of an `E-Root` instance in the Chinese mainland. Therefore, the deployment of an unauthorized instance is highly probable.

(3) Abnormal server identifiers: After requesting server identifiers, one vantage point in China Unicom (Zigong City) reported that responses didn't follow any naming convention summarized in Table IV. The result also suggests the deployment of unauthorized instances.

C. Security Issues of Unauthorized Instances

After confirming the deployment of unauthorized instances, we shed light on the security issues arising from unauthorized instances. Among the 54 vantage points affected by unauthorized instances, **70.4%** (38) of them further reported security issues. The measurement result confirms the presence of various security risks, including DoS, unavailability of DNSSEC validation, vulnerable software, and outdated root zone files. We discuss the issues as follows:

Denial of Service: All 25 vantage points in the network of China Telecom (Zunyi City) observed unauthorized instances failing to resolve all tested domains. Typically, while receiving a query for a TLD, a root instance will provide the referral to the TLD zone with authorized and additional sections (Fig. 5, left) [34]. However, our vantage points received responses in which the answer section contained a localhost IP address while requesting the ccTLD of China (Fig. 5, right). To make matters worse, all other queries to root servers, including queries for other tested TLDs, were time-out. Consequently, these responses can make clients fail to resolve all tested domains. A likely explanation is that the unauthorized instances were experiencing an operational issue.

```
$ dig cn @b.root-servers.net
...
; ANSWER: 0, AUTHORITY: 8, ADDITIONAL: 11
...
;; AUTHORITY SECTION:
cn.    172800  IN  NS  a.dns.cn.
...
cn.    172800  IN  NS  g.dns.cn.

;; ADDITIONAL SECTION:
a.dns.cn.  172800  IN  A 203.119.25.1
...
g.dns.cn.  172800  IN  A 66.198.183.65
```

```
$ dig cn @b.root-servers.net
...
; ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0
...

;; ANSWER SECTION:
cn.     15  IN  A   127.0.0.1
```

Fig. 5. Expected TLD referral (left) and abnormal TLD referral observed in the network of China Telecom at Zunyi (right).

We also found that it was challenging to troubleshoot the operational issue due to the failover mechanism of resolvers. Specifically, the manipulation only affected the traffic toward seven root servers, instead of impacting all 13 root servers as observed in CERNET. We observed the expected results for all testing cases when testing other root servers from our vantage points. Despite unauthorized instances failing to resolve any domain, resolvers will automatically retry querying other root servers while a root server fails to respond (Section V-D) and finally obtain the correct answers from other root candidates. However, such an operational issue increases the cost of domain resolution and hinders the load balancing mechanism of root servers.

Unavailability of DNSSEC validation: We found that 48.1% (26) of vantage points served by unauthorized instances were not accessible for DNSSEC validation. Consequently, resolvers that shared networks with these vantage points would be unable to validate any domain through DNSSEC, as the DNS root is the initial point for validating a domain. The unavailability of DNSSEC validation was not evenly distributed and only affected specific root servers. For example, the majority of vantage points were unavailable to the validation from `K` and `L-Root`, while no vantage point cannot access `E`, `H`, `I`, and `J-Root`.

Vulnerable DNS software: We discovered that an unauthorized instance was equipped with vulnerable DNS software, which could be affected by 100 CVE vulnerabilities. Specifically, a vantage point in CERNET requested `VERSION.BIND` identifier from `H-Root` and received a response indicating an outdated version (*"BIND 9.8.2 rc1"*). The identifier did not adhere to the naming convention of `H-Root` (Table IV) and revealed the software version installed on unauthorized instances. Worse, the version was released on Jan 19th, 2012. From Jan 2012 to July 2023, BIND9 has been assigned 100 CVE numbers [79] corresponding to vulnerabilities such as cache poisoning and DoS attack. An attacker could exploit these published CVE vulnerabilities to compromise the unauthorized instance and further manipulate all DNS traffic within its catchment area.

Outdated root zone files: All 12 vantage points in CERNET confirmed that unauthorized instances equipped outdated root

TABLE VII
COMPARISON OF SECURITY GUARANTEES BETWEEN AUTHORIZED AND UNAUTHORIZED INSTANCES

Security issue	5,489 VPs served by authorized instances	54 VPs served by unauthorized instances
Denial of service	0	46.3%
Unavailable DNSSEC	11.8%	48.1%
Vulnerable software	0	1.9%
Outdated zone	0	22.2%

zone files, which posed the risk of a service outage. This is because the abandoned TLD nameservers may stop their services after an exchange of servers. As a result, authorized instances must promptly load and serve the latest root zone file to avoid an outage.

We observed that unauthorized instances in CERNET experienced significantly slower updates on the zone file when compared to the authorized instances. To measure the update time of authorized instances, RSSAC defines *load time* as the time elapsed between receipt of the notification of updating until 95% of corresponding root instances are ready to serve the new zone file. On the time when we observed outdated zone files (Oct 24th, 2022), the shortest load time was only 7 seconds (`A` and `J-Root`), while the longest load time was 483 seconds (`I-Root`). However, during the same period, ten vantage points in CERNET reported a root zone published 11 days ago, while one vantage point reported a zone published seven days ago. Also, the official metric data [80] reported that 26 versions of root zone files were published within those 11 days, highlighting the importance of timely zone file updates.

D. Security Guarantees of Authorized Instances

In the previous section, we discussed the security issues associated with unauthorized instances. To highlight the risks unauthorized instances pose, we performed a global analysis to assess the security guarantees of authorized instances. This experiment started on Oct 28, 2023, and lasted for seven days. We collected data from *global-wide* vantage points. Utilizing methods mentioned in Section IV-A, we *excluded* the vantage points impacted by unauthorized instances. In total, we collected 5,489 vantage points served by authorized instances, spanning 82 countries, 120 Autonomous Systems (ASes), 161 ISPs, and 270 cities. We detail the result in Table VII.

Our analysis reveals that authorized instances provided significantly greater security guarantees compared to unauthorized instances. Of the 5,489 vantage points served by authorized instances, none were associated with security issues, including Denial of Service (DoS) attacks, use of vulnerable software, or outdated zones. Furthermore, only 11.8% of the vantage points served by authorized instances experienced DNSSEC validation unavailability, compared to 48.1% for those served by unauthorized instances. These findings align with expectations, as DNS specifications [18], [26], [27], [28], [29] require such security guarantees for the DNS root.

E. Summary

In this section, we detail the deployment of unauthorized instances in the Chinese mainland and shed light on the related security issues. To this end, we propose a set of methodologies to analyze unauthorized instances. Through our measurement, 54 vantage points reported evidence of unauthorized instances, including ones deployed in CERNET and the network of China Telecom (Zunyi City). We also confirmed security issues from **70.4%** of vantage points manipulated by unauthorized instances. Such unauthorized instances not only resulted in DoS issues but also caused the unavailability of DNSSEC validation. We also compared the security guarantees between authorized and unauthorized instances to highlight the uncovered security issues. We believe this work provides a deeper insight into unauthorized instances which have not been widely studied.

V. LIMITED CATCHMENT AREA OF AUTHORIZED ROOT INSTANCES

Unlike unauthorized instances serving limited networks, authorized instances offer services on a significantly larger scale. Global instances can even serve the entire Internet, as their route advertisement is allowed to propagate across the Internet [15], [16]. Furthermore, authorized instances have been rapidly deployed in recent times. Over the past three years, the number of authorized instances has increased by 58.6% [15]. However, few studies have examined the catchment area of authorized instances, which may be influenced by deployment issues. In this section, we propose a novel method to measure the catchment area of authorized instances. Through a case study targeting China, we uncovered the deployment issue responsible for the limited catchment area of authorized instances.

Challenge: Measuring the catchment area of root instances is not straightforward since we require addressing two challenges:

(1) Identifying root instances activating anycast: Confirming the exact instance responding to our vantage points is difficult since responses from different root instances contain the same anycast IP address. Due to information on authorized instances, such as hosting networks, not being publicly available, one may reuse the methods for identifying unauthorized instances introduced in Section IV-A. However, such methods cannot provide enough information to identify an authorized instance accurately. For instance, some penultimate hops may not respond to the requests of DNS traceroute and server identifiers due to security concerns (e.g., `G-root` instances, as shown in Tables III and IV). Additionally, some server identifiers could provide insufficient information to determine their identities (e.g., "*contact info@netnod.se*").

(2) Impact of recursive resolvers: The catchment area of authorized instances also depends on the behavior of recursive resolvers. This is because a root instance is queried only when a recursive resolver selects the corresponding root server to send the query. As a result, investigating the algorithm of recursive DNS software is necessary. However, how the instances can affect root server selection and further absorb queries from recursive resolvers is still unknown.

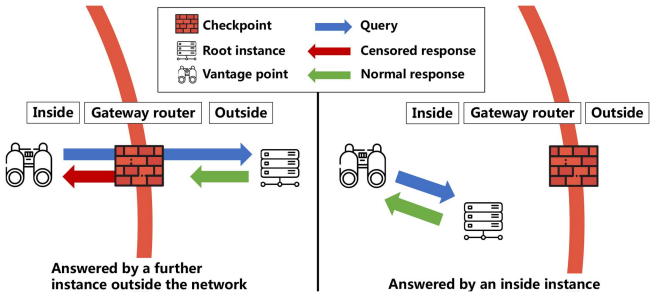

Fig. 6. Using DNS censorship to determine if root instances inside the network serve root queries.

Our solution: To address the challenges mentioned above, we propose a novel methodology that utilizes the side-channel information of DNS censorship mechanisms to measure the catchment area of root instances (Section V-A). Such mechanisms are deployed at the choke points of network topology and coincidentally provide side-channel information that can be used to determine if a query is resolved by a root instance within the censored network. Our result also proves that the method can apply to several countries implementing DNS censorship. Based on the measurement, we shed light on the deployment issues, such as limited BGP policies, that caused the limited catchment area of authorized instances in the Chinese mainland (Sections V-B and V-C). Through code reviewing and dynamic debugging, we also conducted the first study about how the catchment area of domestic instances affects root server selection from the perspective of mainstream recursive software (Section V-D).

A. Methodology

Measuring the catchment area of authorized instances: We utilized the side-channel information of DNS censorship mechanisms to infer the catchment area of authorized instances. To answer the question: *"what is the catchment area of authorized instances?"*, we try to answer another question: *"are DNS queries to root servers resolved inside or outside a network?"*. If resolved inside, we can infer that the catchment area of the responding authorized instances, whose geolocation is publicly available [15], covers the networks of corresponding vantage points. To this end, we utilized a type of DNS censorship mechanism. By censoring at the choke points of network topology, the mechanisms can efficiently censor all DNS queries routed to other networks, irrespective of the requested resolver. Once detecting DNS queries associated with censored domains, the censorship system injects forged responses immediately before the arrival of authentic responses to the users. Consequently, users will accept the forged responses and discard the authentic responses. Such a censorship mechanism coincidentally provides side-channel information to determine the geolocation of an authorized instance.

Fig. 6 elaborates on our approach for determining whether root queries are answered by authorized instances within a network. We sent five DNS queries for censored domains from each vantage point inside a network activating DNS censorship. To ensure that the queries reached the root servers instead of being answered from the cache (e.g., of middleboxes), we prefixed the domains in queries with a nonce value (e.g., *[nonce].censored.com*). Note that we have excluded the affected vantage points in a validation process (Section III) to eliminate the impact of censorship targeting high-level domains (e.g., *[any].censored.com*). As a result, the root query passes through the gateway router for an external root instance, if we receive a censored response (Fig. 6, left). Conversely, an authorized instance resolves the root query within the network, if we capture a normal response with TLD referrals (Fig. 6, right). Mechanisms for performance purposes (e.g., CDN) do not affect our experiment since we directly send queries to DNS roots, which return TLD referrals instead of IP addresses of application servers.

We confirmed the applicability of our methodology in measuring countries that enable DNS censorship. To conduct our method, the censorship system should examine DNS queries from vantage points to root servers instead of only checking the queries to specific resolvers. We verified the condition in six countries whose censorship mechanisms have been widely studied by previous research. As shown in Table VIII, previous research has proven that five out of the six countries would censor all DNS queries regardless of which resolver was requested. As a result, the censorship systems of the five countries also examine the queries from our vantage points toward root servers, satisfying the requirement of our methodology.

In this study, we apply our methodology to China and provide a nationwide case study. To this end, we measured the catchment area of authorized instances in the Chinese mainland (Section V-B) and further revealed the reason behind the restricted catchment area (Section V-C).

Analyzing impact of domestic instances on recursive resolvers: We have designed the method for measuring the catchment area of root instances from the perspective of our vantage points. However, the catchment area of root instances also depends on the selection algorithm of recursive resolvers. This is because an instance is queried only when a recursive resolver selects the corresponding root server. As a result, a perspective of mainstream recursive software is necessary.

To this end, we reviewed the source code of four mainstream DNS software and analyzed their root selection algorithms. In particular, We selected recent versions of four popular open-source recursive DNS implementations: BIND 9 [84] (9.18.9), Unbound [85] (1.17.0), Knot Resolver [86] (5.5.3) and PowerDNS Recursor [87] (4.7.3) as our targets. We first set up a docker container [88] running Ubuntu 22.04 and connected it to a GDB remote debugger [89]. Then, we compiled each DNS software from its source code and launched it within the container. Using the GDB debugger, we reviewed the root selection algorithms by tracing their execution paths. Having identified the factors that affect root server selection, we conducted a measurement to uncover the real-world impact of these factors (Section V-D).

TABLE VIII
DNS CENSORSHIP MECHANISMS ACROSS DIFFERENT COUNTRIES

Country	China	Italy	Russia	Pakistan	Iran	Greece
Scope of censorship	Chinese mainland	ISP	ISP	ISP	ISP	ISP's resolvers
Censored resolvers	All	All	All	All	All	ISP's resolvers
Forged response	forge IP address	NXDOMAIN, timeout, etc.	blockpage, timeout	NXDOMAIN	forge IP address	blockpage
# Authorized instances	24	22	18	5	2	7
Applicable	✔	✔	✔	✔	✔	✘
Reference	[42], [44], [45], [51] [32], [43], [81], [82]	[50] [32], [82]	[46], [47] [32], [82]	[49] [32], [82]	[52] [32], [82]	[83] [32], [82]

TABLE IX
RATIO OF QUERIES THAT TRIGGERED NORMAL RESPONSES, WHICH INDICATED THAT THE QUERIES WERE SERVED BY DOMESTIC INSTANCES

Root	China Telecom	China Unicom	China Mobile	Tencent Cloud	Alibaba Cloud	CERNET
A	100.00%	1.60%	94.63%	100.00%	67.05%	100.00%
B	1.42%	1.27%	0.77%	0.26%	0.76%	100.00%
C	1.76%	1.59%	0.31%	0.00%	0.38%	100.00%
D	1.66%	1.18%	0.56%	0.79%	0.00%	100.00%
E	1.67%	0.87%	9.82%	0.26%	1.52%	100.00%
F	1.78%	100.00%	100.00%	0.26%	15.62%	100.00%
G	1.70%	1.03%	1.33%	0.26%	0.76%	100.00%
H	1.95%	2.31%	0.61%	0.53%	0.38%	100.00%
I	2.48%	1.38%	97.75%	0.53%	65.90%	100.00%
J	1.62%	91.90%	98.77%	0.79%	78.67%	100.00%
K	100.00%	100.00%	0.31%	100.00%	41.90%	100.00%
L	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
M	2.81%	1.24%	0.41%	0.53%	0.57%	100.00%
Total	24.53%	31.10%	38.87%	23.40%	28.73%	100.00%

Root servers with domestic instances that were deployed in the chinese mainland are marked with lighter backgrounds.

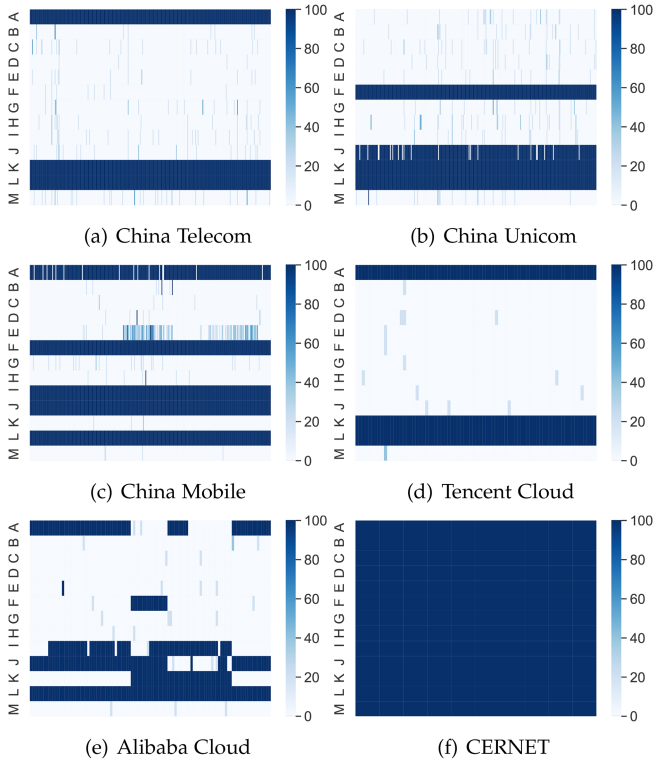

Fig. 7. Ratio of queries that triggered normal responses at each vantage point. Darker cells indicate that more queries from the corresponding vantage points were resolved by domestic instances.

B. Networks Served by Domestic Instances

In this section, we measured the catchment area of domestic instances and confirmed the networks served by these instances. With the support of 7,860 vantage points, we started our experiment on Sep 4, 2022, and the experiment spanned 60 days. Table IX illustrates the ratio of root queries that receive normal responses (i.e., served by domestic instances) for each ISP network, while Fig. 7 provides a closer view of the ratios reported by each vantage point. We present our analyses as follows:

CERNET: We discovered that *all* root queries from CERNET vantage points were domestically answered, as expected since we have confirmed the unauthorized instances in CERNET manipulating queries toward all root servers (Section IV-B).

Major commercial ISPs: We observed that the catchment area of authorized instances can *partially* cover major commercial ISPs in the Chinese mainland. The domestic instances of A, F, J, K, and L-Root had nationwide catchment area for vantage points in at least one of the three major ISPs, as they answered over 90% of root queries from these networks (highlighted in bold in Table IX, e.g., Telecom to A-Root). Zooming into individual vantage points, as illustrated in Fig. 7, we found that the catchment area of authorized instances within a given ISP network exhibited minor differences across geolocations. We further analyze the reason behind the limited catchment area of authorized instances in Section V-C.

Cloud services: We observed that Tencent Cloud was served by domestic A, K, and L-Root instances, while the case of Alibaba Cloud was more complex, as shown in Fig. 7. Through discussions with the vendors, we confirmed that no authorized instance was deployed in their networks, and the accessibility of their networks to root instances depended on the ISP networks that they connected to. As a result, the reason behind our result could be that Tencent Cloud ASes exclusively peered to China Telecom ASes. In contrast, multiple Alibaba Cloud ASes peered to different ISP ASes.

Change in the catchment area: We traced the change in the catchment area of authorized instances by comparing an experiment conducted two years ago. Using the same methodology, we conducted the experiment in Dec 2020, and 16 authorized instances (0, 4, 1, 2, 3, and 6 instances for A, F, I, J, K, and L-Root respectively) were deployed in the Chinese mainland at that time. Interestingly, the catchment areas of F, I, and J-Root *changed without a deployment of new instances:* For example, in 2022, domestic F-Root instances served our vantage points in Unicom and Mobile, while in 2020, only vantage points in Telecom could be served. Besides, both I and J-Root domestic

TABLE X
SUCCESS RATE OF DNS CENSORSHIP

ISP	Success Rate	ISP	Success Rate
China Telecom	98.46%	Tencent Cloud	99.47%
China Unicom	98.23%	Alibaba Cloud	99.39%
China Mobile	99.17%	CERNET	99.57%
Total			**99.05%**

instances served vantage points in Mobile in 2022 but not in 2020. We also found the newly deployed `A-Root` instance efficiently served almost all vantage points in Telecom and Mobile.

Normal responses from a root without domestic instance: We confirmed that censorship failures and unauthorized instances contributed to normal responses from a root server without corresponding authorized instances deployed in domestic (e.g., normal responses from `B-Root`). To verify the impact of censorship failure, we conducted a separate experiment. Utilizing our measurement platform, we sent DNS queries for censored domains (e.g., *[nonce].google.com*) to a self-built DNS server located outside the Chinese mainland. Typically, all queries should pass the DNS censorship system and trigger censored responses. However, as shown in Table X, the DNS censorship system achieved an overall success rate of 99.05%, explaining why DNS queries to root servers without domestic instances receive normal (uncensored) responses. Additionally, in Section IV-B, we have confirmed the deployment of unauthorized instances, which can also generate responses in domestic.

C. Reasons Behind the Limited Catchment Area

From some ISP networks to root servers that *did* deploy domestic instances, the proportion of receiving normal responses remained remarkably low, as shown in Table IX. Worse, some authorized instances with limited catchment area were even *global* instances, which were supposed to serve users on the whole Internet [15], [16] (e.g., Unicom to the `A-Root` instance in the Chinese mainland). The results suggest that domestic instances almost failed to serve these networks. To reveal the reason, we provide a detailed analysis below:

Domestic `I-Root` instance: We observed that the *global* `I-Root` instance in Beijing did not serve any vantage point in the three major commercial ISPs, due to the exclusive service in CSTNET (China Science and Technology Network). To locate the hosting network of this instance, we performed traceroute measurements toward `I-root` in Dec 2020 and found the penultimate hops belonged to CSTNET, which served research institutions and hi-tech enterprises in China. To find out the exact instance serving CSTNET, we employed seven volunteers as vantage points in CSTNET. We asked them to request server identifiers from `I-Root`, and the responses were "*s1.bei*". We confirmed with Netnod (the operator of `I-Root`) that the string represented the instance in Beijing, which was presented in the official announcement. As a result, the domestic `I-root` instance exclusively served CSTNET in Dec 2020, possibly due to the lack of peering between major commercial ISPs and Netnod. However, we observed that all responses from `I-root` were censored responses, indicating they passed through the international gateway. Such a result indicates the root instance was physically deployed within the Chinese mainland but outside the scope of DNS censorship. This could result from a security incident in 2010 when the `I-root` instance provided incorrect responses to oversea users due to DNS censorship [90].

We also found an unannounced `I-root` instance contributed to the domestic resolution in a major commercial ISP, instead of the existing one in CSTNET extending its catchment area. Specifically, we discovered that vantage points in Mobile (one of the three major commercial ISPs) were served by a domestic `I-root` instance in Oct 2022. This was in contrast to the official announcement [15], [91], which stated that no new `I-Root` instance had been deployed in the Chinese mainland since Dec 2020. To confirm if the `I-Root` instance in CSTNET had extended its catchment area, we examined the corresponding penultimate hops and server identifiers. Surprisingly, we found that the `I-Root` instance was deployed in a Chinese city named Shenyang with a server identifier as "*s1.she*". The official announcement [15] has not updated the information of the new root instance until Feb 2023. We also analyzed the catchment area of the `I-Root` instance in Beijing. To this end, we hired ten volunteers in CSTNET and repeated the experiment conducted two years ago. We found that the instance still exclusively served vantage points within CSTNET.

Restricted BGP policy: Similarly, we located the hosting networks of other domestic instances through DNS traceroute (Section IV-A). For example, vantage points in Telecom network were not served by domestic `F` and `J-Root` instances since these instances were located in networks of Unicom and Mobile. The Mobile vantage points could not access domestic `K-Root` instances, which were found in Telecom and Unicom networks. Some domestic instances (e.g., `A-Root`) could not be located since their penultimate hop of the traceroute path did not respond with an ICMP message containing their IP addresses.

We conclude that the reason behind the limited catchment area is the restricted BGP policy by discussing with ISPs and root server operators. In particular, the BGP announcement of authorized instances in the Chinese mainland primarily originated from ISP networks rather than Internet exchange points. However, both global and local root instances hosted by an ISP were typically unshared with other networks due to the limitation of BGP routing policies. Meanwhile, major commercial ISPs did not directly peer with each other.[3] Consequently, the catchment area of local `F` and `K-Root` instances did not cover networks of other ISPs. Some domestic ISPs even can not access the global instances deployed in the Chinese mainland (e.g., Unicom to `A-Root`).

Routing adjustment: An adjustment of the routing policy can also limit the catchment area of authorized instances. As presented in Table IX, vantage points in Unicom network reported only 91.90% of queries to `J-Root` were resolved in domestic, which was lower than other cases where we confirmed

[3]We attempted to verify this through inspecting BGP routing information in RouteViews [92]. However, the dataset had little coverage of ASes in China.

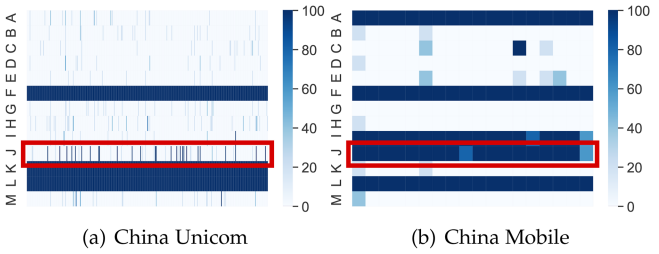

Fig. 8. From Sep 13th, 12:58 pm to Sep 17th, 12:06 pm (UTC), the ratio of queries that triggered normal responses at each vantage point in the networks of China Unicom and China Mobile.

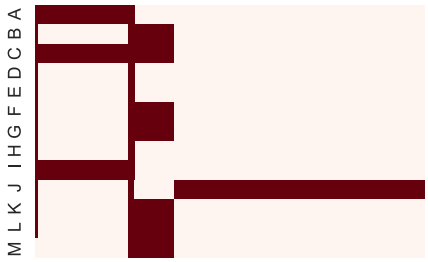

Fig. 9. Vantage points cannot access root servers. A red cell indicates that the root server (X-axis) was unavailable to the vantage point (Y-axis).

ISPs hosting domestic instances (e.g., Unicom to `F-Root`). We further confirmed a temporary service interruption from 12:58 pm on Sep 13th to 12:06 pm on Sep 17th (UTC). As shown in Fig. 8, we zoomed on the result during the time frame. The majority of Unicom vantage points were not served by domestic `J-Root` instances, whereas the majority of Mobile vantage points remained unaffected. As a result, we infer that this issue resulted from an adjustment in the routing policy rather than domestic root instances being out of service.

D. Impact on Root Selection of Recursive Resolvers

In Section V-B, we determined the catchment area of authorized instances by directly sending DNS queries to each root server. However, a root instance is queried only when a *recursive resolver* selects the corresponding root server to send the query. As a result, we took the perspective of mainstream recursive software and investigated how the catchment area of authorized instances affects the root selection of a resolver. To this end, we first identified the factors (e.g., the latency of a root server) that impacted the root selection of resolvers through software analysis. Based on the measurement result, we then analyzed how authorized instances impacted these factors and further affected the root server selection of resolvers.

Root server selection algorithm: By conducting source code analysis and dynamic debugging, we confirmed that four mainstream recursive DNS software avoided selecting the root server failing to respond and preferred the root server responding with lower latency.[4] For example, BIND9 and Knot Resolver significantly preferred the root server with the smallest Round Trip Time (RTT) in the majority of cases. As a result, the *accessibility* and *latency* of a root server were the key factors that affected the root server selection of a recursive resolver. Such implementations followed the suggestion of DNS standards: recursive resolvers should "find the best server to ask" [34].

Measurement of the factors affecting root server selection: We further measured the accessibility and latency of each root server from our vantage points, which were the factors affecting whether the corresponding root instances were selected and queried. To this end, our vantage points sent five DNS queries to each root server. To ensure that the queries arrived at the root servers instead of being answered by middleboxes or censorship mechanisms, we queried a non-censored domain and prefixed it with a nonce value in each query (i.e., *[nonce].non-censored.com*). We labeled a root server unavailable to a vantage point only if all queries were time-out. Since a root server is non-recursive (i.e., it does not query other servers), we can determine its latency by sending DNS queries directly and recording the arrival time of its responses.

We observed that the majority of vantage points accessed all root servers successfully, except for 129 vantage points whose all queries were time-out, as shown in Fig. 9. Interestingly, only some root servers were inaccessible by specific vantage points, while others remained unaffected. Consequently, recursive resolvers sharing the networks with the above vantage points would avoid sending queries to unresponsive root servers and instead request to other candidates.

We also discovered that authorized instances in domestic contributed to reducing the overall delay of the corresponding root servers and attracted root queries from recursive resolvers. Fig. 10 illustrates the delay experienced by vantage points within each ISP when querying the 13 root servers. Specifically, queries in CERNET to any root server were answered within 30 ms due to the deployment of unauthorized instances, which have been identified in Section IV-B. For the other ISPs we examined, we observed that networks with a high ratio (>90%) of queries resolved by domestic instances experienced considerably low delays (lower than 100 ms in most cases) when querying the corresponding root servers. In contrast, vantage points encountered high delays (several hundred ms) when querying root instances located oversea. This was due to either no domestic instance serving their network (e.g., Unicom to `A-Root`) or no domestic instances being deployed (e.g., `B-Root`).

E. Summary

In this section, we uncover the deployment issues of authorized instances in the Chinese mainland through a comprehensive measurement. We propose a novel methodology to measure the catchment area of authorized instances by leveraging the side-channel information of DNS censorship. Our results demonstrate that the catchment area of authorized instances in the Chinese mainland can be limited, even when they were global instances that should serve the entire Internet. We also revealed that the restricted routing policy was the reason behind

[4] As part of our contribution, we provide the pseudo-code and detailed explanations of all root selection algorithms in the shared repository.

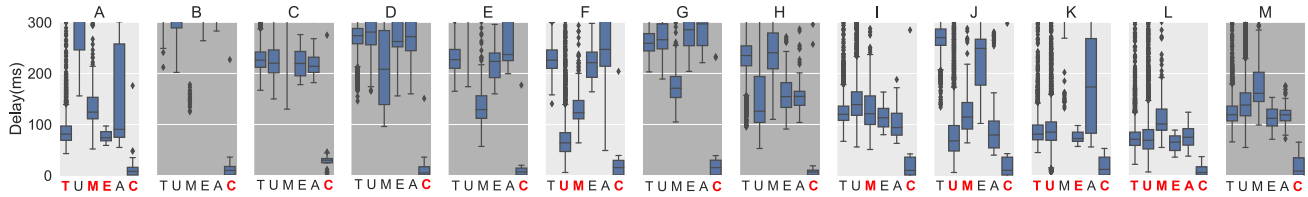

Fig. 10. Latency of each root server measured from different vantage points (T: Telecom, U: Unicom, M: Mobile, E: Tencent, A: Alibaba, C: CERNET). Root servers with domestic instances are marked with lighter backgrounds. The bold ISP indicates a high ratio (>90%) of queries resolved domestically.

the limited catchment area through an analysis of mainstream DNS software. Additionally, we confirmed that the authorized instances in the Chinese mainland can help absorb queries from domestic recursive resolvers. We believe this study is helpful for the future deployment of root instances to cover the currently unserved networks.

VI. Discussion

A. Ethics

To avoid negative impacts, we carefully designed our methodology and addressed two ethical aspects with the guide of ethical principles [93], [94].

One primary ethical consideration is avoiding causing negative impacts on proxy vendors from censorship mechanisms, such as triggering alerts. In this study, we did not investigate or challenge DNS censorship mechanisms through vantage points. Instead, we only leveraged the known characteristics of censorship mechanisms (i.e., injection of DNS responses), which have been studied in numerous works [32], [42], [44], [45], [51], [82]. Specifically, we only queried non-existent subdomains under the censored domains (i.e., *[nonce].censored.com*) and did not make any connection to the IP address in a censored response. Several studies [32], [82] focusing on censorship mechanisms have also concluded that such DNS queries for censored domains posed negligible harm or judicial risks to the operators of vantage points.

Another ethical consideration is to avoid overloading proxy services. To this end, we equipped each measurement client in our platform with an independent account and ensured that we paid for all redirected traffic (Section III). In total, we signed up 50 accounts and paid 1242.9$for the proxy services. During the measurement, we requested DNS resolution and switched exit nodes, which fell within the business scope of the proxy services. To comply with service regulations, we also strictly limited the rate of switching exit nodes and sending queries on each account. While utilizing an account, we switched an exit node every 10 minutes on average.

B. Recommendations

Through discussions with ISPs and network operators, we make recommendations to improve the deployment status of DNS root instances.

Our recommendations for fixing deployment issues of unauthorized instances include the following: (1) We recommend NOT establishing unauthorized instances since we have proven they can cause serious security issues. (2) However, our research indicates that individuals may deploy unauthorized instances with benign intentions. That is, while authorized instances administered by RSS and RSO fail to adequately serve a network, operators could depend on self-initiated unauthorized instances to enhance network performance. For the above situation, we recommend considering methods that follow RFC recommendations to improve access to the RSS (e.g., running a local root copy [95]) instead of relying on unauthorized instances. (3) Maintainers of unauthorized instances should ensure a proper operation and promptly troubleshoot operational issues. This can be achieved by implementing real-time monitoring and designing a series of test cases for unauthorized instances. (4) Maintainers of unauthorized instances should promptly update the DNS software and zone files from the official source. Patches of the latest DNS software can defend against known vulnerabilities, and timely updates to the zone file can help avoid the risk of an outage. (5) Maintainers of unauthorized instances should enable all security mechanisms for root instances, such as DNSSEC validation, according to the best practice adopted by authorized instances.

We also make the following suggestions for improving the catchment area of authorized instances: (1) For networks not covered by nearby instances' catchment areas, we recommend establishing BGP peering with the root server networks, subject to political and commercial considerations. (2) Root server operators and ISPs may take these areas into prior consideration for future deployment of authorized instances. (3) To inform operators about whether their networks can be served, we recommend making BGP peering information between ISPs and root servers transparent (e.g., disclosing which networks host a root server or peer with them). (4) For DNS community, while the status of root servers has been extensively monitored, systems that measure root servers from the resolvers' perspective still need to be developed.

C. Limitations

In Section IV, we utilized the naming conventions of authorized instances as one of the metrics for detecting unauthorized instances. We did not exhaustively verify the assumption that all authorized instances adhere to these naming conventions. It is also possible that an attacker could circumvent our detection mechanism by mimicking these naming conventions. However, apart from naming conventions, we also leveraged the result of DNS traceroute, which is challenging for attackers to forge. Besides, authorized instances are constrained by

the guidelines of RSS, RSO, and RFC documentation [75], which typically makes authorized instances exhibit consistent naming conventions. Consequently, we consider that our methodology has reduced the occurrence of false positives and negatives.

The accuracy of geolocation data may also affect our results. We mitigate the impact from three aspects. First, we utilized the Maxmind database [66] for our research. Maxmind is renowned and has been extensively used in studies focusing on DNS security (e.g., [96], [97]). Second, we ensured the use of the latest data available during our experiments to follow changes in geolocation relationships. Third, we verified the geolocation of our vantage points against both the announcements from proxy vendors and the Maxmind database, discarding any vantage point with inconsistent geolocation labels. We believe the above considerations have mitigated the impact of non-deterministic geolocation.

VII. Related Work

DNS manipulation: DNS protocol is an attractive target of attackers since it lacks authentication and integrity validation. Jones et al. [9] conducted a measurement of unauthorized instances by leveraging query latency and UDP traceroute. This method worked because the study targeted `B-Root`, which wasn't activated anycast eight years ago [30]. Such a study also echoed our finding about unauthorized instances in CERNET but did not further investigate their characteristics and security issues. Besides, a series of studies explored DNS manipulation through vantage points. Some of them took open resolvers as vantage points [54], [98], while others utilized proxy networks [31], [32] or measurement platforms [9]. Analyzing historical data (e.g., passive DNS data) was also an effective method to detect DNS hijacking [6], [99]. In this study, we conducted an extensive measurement and analyzed both the deployment and security issues of unauthorized instances in the Chinese mainland.

DNS censorship: As a type of DNS manipulation, DNS censorship has been investigated in different scopes. Some researchers provided a global scope of DNS censorship through worldwide measurement [32], [54], [82]. Numerous studies also shed light on country-specific DNS censorship, including China [42], [44], [45], [51], Russia [46], [47], Italy [50], Pakistan [49], Iran [52], India [48], and Greece [83].

Prior studies have investigated the implementation and characteristics of DNS censorship mechanisms. Unlike these works, our study did not focus on dissecting or exploring these features. Instead, we leveraged the features to measure the catchment area of root instances. To our knowledge, this approach is novel and has not been previously explored.

Performance of DNS roots: To understand the performance of the RSS, researchers have devoted to investigating the impact of uneven distribution of root instances on end-user query latency [19], [21] and evaluating effects of anycast through examining DNS traffic and BGP data [13], [100], [101], [102]. However, little has been done to understand the catchment area of anycast instances behind root servers or how their deployment and operation can be improved in the future.

NS Selection of Recursive Software: A series of works have examined how common DNS resolvers select and query authoritative servers (nameservers instead of root servers) [103], [104], [105]. Almost all have answered this question by designing simulation experiments or inspecting outgoing DNS queries. They concluded the majority of implementations preferred authoritative servers with the lowest latency, while others chose randomly. However, the reasons remain unrevealed, as few provide source code analysis.

VIII. Conclusion

This paper investigates the deployment issues of root server instances in the Chinese mainland from a nationwide perspective. To conduct this study, we developed a measurement platform to collect vantage points. We first focused on unauthorized instances to identify their deployment and associated security issues. In total, 54 vantage points witnessed the evidence of unauthorized instances. Furthermore, 70.4% of these vantage points confirmed security issues of unauthorized instances, including DoS, unavailability of DNSSEC validation, and vulnerable DNS software. Then, we examined the catchment area of authorized instances in the Chinese mainland by leveraging side-channel information provided by DNS censorship. We discovered that some authorized instances were inaccessible from major ISP networks due to the limitation of BGP routing policies. We also evaluated the impact of domestic instances on root server selection in recursive resolvers through software analysis and measurement. Finally, we provided recommendations for different entities. We believe multiple parties can take action to improve the security and operational status of the DNS root instances in China.

Acknowledgment

The authors appreciate all editors and reviewers for their constructive comments.

References

[1] F. Zhang, C. Lu, B. Liu, H. Duan, and Y. Liu, "Measuring the practical effect of DNS root server instances: A China-wide case study," in *Proc. Int. Conf. Passive Act. Meas.*, O. Hohlfeld, G. Moura, and C. Pelsser, Eds., Cham: Springer International Publishing, 2022, pp. 247–263.

[2] P. E. Hoffman and J. Schlyter, "The DNS-Based authentication of named entities (DANE) transport layer security (TLS) protocol: TLSA," RFC 6698, Aug. 2012. [Online]. Available: https://www.rfc-editor.org/info/rfc6698

[3] V. Dukhovni and W. Hardaker, "The DNS-Based authentication of named entities (DANE) protocol: Updates and operational guidance," RFC 7671, Oct. 2015. [Online]. Available: https://www.rfc-editor.org/info/rfc7671

[4] M. Kucherawy, D. Crocker, and T. Hansen, "DomainKeys identified mail (DKIM) signatures," RFC 6376, Sep. 2011. [Online]. Available: https://www.rfc-editor.org/info/rfc6376

[5] S. Kitterman, "Sender policy framework (SPF) for authorizing use of domains in email, Version 1," RFC 7208, Apr. 2014. [Online]. Available: https://www.rfc-editor.org/info/rfc7208

[6] G. Akiwate et al., "Retroactive identification of targeted DNS infrastructure hijacking," in *Proc. 22nd ACM Internet Meas. Conf.*, New York, NY, USA, 2022, pp. 14–32. [Online]. Available: https://doi.org/10.1145/3517745.3561425

[7] T. Dai, P. Jeitner, H. Shulman, and M. Waidner, "The hijackers guide to the galaxy: Off-path taking over internet resources," in *Proc. 30th USENIX Secur. Symp.*, 2021, pp. 3147–3164. [Online]. Available: https://www.usenix.org/conference/usenixsecurity21/presentation/dai

[8] T. Dai, H. Shulman, and M. Waidner, "Let's downgrade let's encrypt," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, New York, NY, USA, 2021, pp. 1421–1440, doi: 10.1145/3460120.3484815.

[9] B. Jones, N. Feamster, V. Paxson, N. Weaver, and M. Allman, "Detecting DNS root manipulation," in *Proc. Int. Conf. Passive Act. Netw. Meas.*, Springer, 2016, pp. 276–288.

[10] DDoS attacks on dyn, 2016. [Online]. Available: https://en.wikipedia.org/wiki/DDoS_attacks_on_Dyn

[11] Akamai Edge DNS goes down, takes a chunk of the internet with it, 2021. [Online]. Available: https://www.theregister.com/2021/07/22/akamai_edge_dns_outage/

[12] T. Hardie, "Distributing authoritative name servers via shared unicast addresses," RFC 3258, Apr. 2002. [Online]. Available: https://rfc-editor.org/rfc/rfc3258.txt

[13] Z. Liu et al., "Two days in the life of the DNS anycast root servers," in *Proc. Int. Conf. Passive Act. Netw. Meas.*, Springer, 2007, pp. 125–134.

[14] RSSAC, "RSSAC057: Requirements for measurements of the local perspective on the root server system," 2021.

[15] "Root server technical operations association," 2022. [Online]. Available: https://root-servers.org/

[16] K. E. Lindqvist and J. Abley, "Operation of anycast services," RFC 4786, Dec. 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4786

[17] J. Abley, "Hierarchical anycast for global service distribution," 2003.

[18] RSSAC, "RSSAC056: RSSAC advisory on rogue DNS root server operators," 2021.

[19] J. Liang, J. Jiang, H. Duan, K. Li, and J. Wu, "Measuring query latency of top level DNS servers," in *Proc. Int. Conf. Passive Act. Netw. Meas.*, Springer, 2013, pp. 145–154.

[20] ISC, "Atlas data viewer," 2022. [Online]. Available: https://atlas-vis.isc.org/

[21] T. Koch, E. Katz-Bassett, J. Heidemann, M. Calder, C. Ardi, and K. Li, "Anycast in context: A tale of two systems," in *Proc. ACM SIGCOMM Conf.*, 2021, pp. 398–417.

[22] Measuring the health of the domain name system, 2010. [Online]. Available: https://www.icann.org/en/system/files/files/dns-ssr-symposium-report-1--03feb10-en.pdf

[23] ICANN, "The ITHI (identifier technologies health indicators) project," 2022. [Online]. Available: https://ithi.research.icann.org/about.html

[24] RSSAC, "RSSAC002: RSSAC advisory on measurements of the root server system," 2015.

[25] X. Fan, J. Heidemann, and R. Govindan, "Evaluating anycast in the domain name system," in *Proc. IEEE Conf. Comput. Commun.*, 2013, pp. 1681–1689.

[26] M. Blanchet and L.-J. Liman, "DNS root name service protocol and deployment requirements," RFC 7720, Dec. 2015. [Online]. Available: https://www.rfc-editor.org/info/rfc7720

[27] D. Karrenberg, M. Kosters, R. Plzak, and R. Bush, "Root name server operational requirements," RFC 2870, Jun. 2000. [Online]. Available: https://www.rfc-editor.org/info/rfc2870

[28] RSSAC, "RSSAC055: Principles guiding the operation of the public root server system," 2021.

[29] RSSAC, "RSSAC020: RSSAC statement on client side reliability of root DNS data," 2021.

[30] Univ. of Southern California, "B-root begins anycast in may," 2017. [Online]. Available: https://b.root-servers.org/news/2017/04/17/anycast.html

[31] B. Liu et al., "Who is answering my queries: Understanding and characterizing interception of the DNS resolution path," in *Proc. 27th USENIX Secur. Symp.*, 2018, pp. 1113–1128.

[32] A. A. Niaki et al., "ICLab: A global, longitudinal internet censorship measurement platform," in *Proc. IEEE Symp. Secur. Privacy*, 2020, pp. 135–151.

[33] Cernet, 2022. [Online]. Available: https://www.edu.cn/

[34] Domain names - Concepts and facilities, RFC 1034, Nov. 1987. [Online]. Available: https://rfc-editor.org/rfc/rfc1034.txt

[35] D. Conrad, "Brief overview of the root server system," 2020. [Online]. Available: https://www.icann.org/en/system/files/files/octo-010--06may20-en.pdf

[36] IANA, "Root zone management," 2022. [Online]. Available: https://www.iana.org/domains/root

[37] S. Khodayari and G. Pellegrino, "The state of the samesite: Studying the usage, effectiveness, and adequacy of samesite cookies," in *Proc. IEEE Symp. Secur. Privacy*, 2022, pp. 1590–1607.

[38] M. Bugliesi, S. Calzavara, R. Focardi, and W. Khan, "CookiExt: Patching the browser against session hijacking attacks," *J. Comput. Secur.*, vol. 23, no. 4, pp. 509–537, Sep. 2015, doi: 10.3233/JCS-150529.

[39] M. Kranch and J. Bonneau, "Upgrading HTTPS in mid-air," in *Proc. Thz Netw. Distrib. Syst. Secur. Symp.*, 2015.

[40] D. E. E. 3rd, "Domain name system security extensions," RFC 2535, Mar. 1999. [Online]. Available: https://www.rfc-editor.org/info/rfc2535

[41] IAB, "IAB technical comment on the unique DNS root," RFC 2826, May 2000. [Online]. Available: https://www.rfc-editor.org/info/rfc2826

[42] N. P. Hoang et al., "How great is the great firewall? Measuring China's DNS censorship," in *Proc. 30th USENIX Secur. Symp.*, 2021, pp. 3381–3398. [Online]. Available: https://www.usenix.org/conference/usenixsecurity21/presentation/hoang

[43] R. Ensafi, P. Winter, A. Mueen, and J. R. Crandall, "Analyzing the great firewall of China over space and time," in *Proc. Privacy Enhancing Technol.*, vol. 2015, no. 1, pp. 61–76, 2015.

[44] Anonymous, "Towards a comprehensive picture of the great firewall's DNS censorship," in *Proc. 4th USENIX Workshop Free Open Commun. Internet*, San Diego, CA, 2014. [Online]. Available: https://www.usenix.org/conference/foci14/workshop-program/presentation/anonymous

[45] G. Lowe, P. Winters, and M. L. Marcus, "The great DNS wall of China," MS, New York University, vol. 21, p. 1, 2007.

[46] D. Xue, B. Mixon-Baca, A. Ablove, B. Kujath, J. R. Crandall, and R. Ensafi, "TSPU: Russia's decentralized censorship system," in *Proc. 22nd ACM Internet Meas. Conf.*, 2022, pp. 179–194.

[47] R. Ramesh et al., "Decentralized control: A case study of Russia," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2020.

[48] T. K. Yadav, A. Sinha, D. Gosain, P. K. Sharma, and S. Chakravarty, "Where the light gets in: Analyzing web censorship mechanisms in India," in *Proc. Internet Meas. Conf.*, New York, NY, USA, 2018, pp. 252–264, doi: 10.1145/3278532.3278555.

[49] Z. Nabi, "The anatomy of web censorship in Pakistan," in *Proc. 3rd USENIX Workshop Free Open Commun. Internet*, 2013.

[50] G. Aceto, A. Montieri, and A. Pescapé, "Internet censorship in Italy: An analysis of 3G/4G networks," in *Proc. IEEE Int. Conf. Commun.*, 2017, pp. 1–6.

[51] J. R. Crandall, D. Zinn, M. Byrd, E. T. Barr, and R. East, "ConceptDoppler: A weather tracker for internet censorship," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2007, pp. 352–365.

[52] S. Aryan, H. Aryan, and J. A. Halderman, "Internet censorship in Iran: A first look," in *Proc. 3rd USENIX Workshop Free Open Commun. Internet*, 2013.

[53] J. da Silva Damas, M. Graff, and P. A. Vixie, "Extension mechanisms for DNS (EDNS(0))," RFC 6891, Apr. 2013. [Online]. Available: https://www.rfc-editor.org/info/rfc6891

[54] P. Pearce et al., "Global measurement of DNS manipulation," in *Proc. 26th USENIX Secur. Symp.*, Vancouver, BC, 2017, pp. 307–323. [Online]. Available: https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/pearce

[55] N. Weaver, C. Kreibich, and V. Paxson, "Redirecting DNS for ads and profit," in *Proc. USENIX Workshop Free Open Commun. Internet*, 2011, pp. 2–3.

[56] S. Huang, F. Cuadrado, and S. Uhlig, "Middleboxes in the Internet: A HTTP perspective," in *Proc. Netw. Traffic Meas. Anal. Conf.*, 2017, pp. 1–9.

[57] J. Jueckstock et al., "Towards realistic and reproducibleweb crawl measurements," in *Proc. Web Conf.*, New York, NY, USA, 2021, pp. 80–91, doi: 10.1145/3442381.3450050.

[58] N. Demir, M. Große-Kampmann, T. Urban, C. Wressnegger, T. Holz, and N. Pohlmann, "Reproducibility and replicability of web measurement studies," in *Proc. ACM Web Conf.*, New York, NY, USA, 2022, pp. 533–544, doi: 10.1145/3485447.3512214.

[59] S. Roth, S. Calzavara, M. Wilhelm, A. Rabitti, and B. Stock, "The security lottery: Measuring client-side web security inconsistencies," in *Proc. 31st USENIX Secur. Symp.*, Boston, MA, 2022, pp. 2047–2064. [Online]. Available: https://www.usenix.org/conference/usenixsecurity22/presentation/roth

[60] M. T. Khan, J. DeBlasio, G. M. Voelker, A. C. Snoeren, C. Kanich, and N. Vallina-Rodriguez, "An empirical analysis of the commercial VPN ecosystem," in *Proc. Internet Meas. Conf.*, New York, NY, USA, 2018, pp. 443–456, doi: 10.1145/3278532.3278570.

[61] Z. Weinberg, S. Cho, N. Christin, V. Sekar, and P. Gill, "How to catch when proxies lie: Verifying the physical locations of network proxies with active geolocation," in *Proc. Internet Meas. Conf.*, Boston, MA, USA, 2018, pp. 203–217. [Online]. Available: https://dl.acm.org/citation.cfm?id=3278551
[62] ProxyRack, "Proxyrack: Buy proxies HTTP, UDP, SOCKS proxy," 2022. [Online]. Available: https://www.proxyrack.com/
[63] R. N. Staff, "Ripe atlas: A global Internet measurement network," *Internet Protocol J.*, vol. 18, no. 3, pp. 2–26, 2015.
[64] A. J. Valencia, G. Zorn, W. Palter, G.-S. Pall, M. Townsley, and A. Rubens, "Layer two tunneling protocol "L2TP"," RFC 2661, Aug. 1999. [Online]. Available: https://www.rfc-editor.org/info/rfc2661
[65] G. Zorn, G.-S. Pall, and K. Hamzeh, "Point-to-point tunneling protocol (PPTP)," RFC 2637, Jul. 1999. [Online]. Available: https://www.rfc-editor.org/info/rfc2637
[66] MaxMind, "IP geolocation and online fraud prevention," 2022. [Online]. Available: https://www.maxmind.com/en/home
[67] China telecom limited, 2022. [Online]. Available: https://www.chinatelecom-h.com/en/global/home.php
[68] China unicom limited, 2022. [Online]. Available: http://www.chinaunicom.com.cn/
[69] China mobile limited, 2022. [Online]. Available: https://www.chinamobileltd.com/en/global/home.php
[70] Alibaba cloud, 2022. [Online]. Available: https://www.alibabacloud.com/
[71] Tencent cloud, 2022. [Online]. Available: https://www.tencentcloud.com/
[72] A. Bhaskar and P. Pearce, "Many roads lead to rome: How packet headers influence DNS censorship measurement," in *Proc. 31st USENIX Secur. Symp.*, 2022, pp. 449–464.
[73] R. Austein, "DNS name server identifier (NSID) option," RFC 5001, Aug. 2007. [Online]. Available: https://rfc-editor.org/rfc/rfc5001.txt
[74] 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.txt
[75] J. Abley and T. Manderson, "A summary of various mechanisms deployed at L-root for the identification of anycast nodes," RFC 7108, Jan. 2014. [Online]. Available: https://www.rfc-editor.org/info/rfc7108
[76] Domain names - Implementation and specification, RFC 1035, Nov. 1987. [Online]. Available: https://www.rfc-editor.org/info/rfc1035
[77] IANA, "Root files," 2022. [Online]. Available: https://www.iana.org/domains/root/files
[78] D. R. Conrad, "Indicating resolver support of DNSSEC," RFC 3225, Dec. 2001. [Online]. Available: https://www.rfc-editor.org/info/rfc3225
[79] The MITRE Corporation, "CVE," 2024. [Online]. Available: https://cve.mitre.org/
[80] Verisign, "Root server system metrics," 2024. [Online]. Available: https://a.root-servers.org/rssac-metrics/raw
[81] P. Zhu et al., "Characterizing transnational internet performance and the great bottleneck of China," in *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 4, no. 1, Jun. 2020, Art. no. 13, doi: 10.1145/3379479.
[82] R. Sundara Raman, P. Shenoy, K. Kohls, and R. Ensafi, "Censored planet: An internet-wide, longitudinal censorship observatory," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, New York, NY, USA, 2020, pp. 49–66, doi: 10.1145/3372297.3417883.
[83] V. Ververis, G. Kargiotakis, A. Filastò, B. Fabian, and A. Alexandros, "Understanding internet censorship policy: The case of Greece," in *Proc. 5th USENIX Workshop Free Open Commun. Internet*, Washington, D.C., 2015. [Online]. Available: https://www.usenix.org/conference/foci15/workshop-program/presentation/ververis
[84] ISC, "BIND 9," 2022. [Online]. Available: https://www.isc.org/bind/
[85] NLnet Labs, "Unbound," 2022. [Online]. Available: https://www.nlnetlabs.nl/projects/unbound/about/
[86] CZ.NIC, "Knot resolver," 2022. [Online]. Available: https://www.knot-resolver.cz/
[87] PowerDNS, "PowerDNS recursor," 2022. [Online]. Available: https://www.powerdns.com/recursor.html
[88] Docker: Empowering app development for developers, 2022. [Online]. Available: https://www.docker.com/
[89] GDB: The GNU project debugger - GNU.org, 2022. [Online]. Available: https://www.gnu.org/software/gdb/
[90] I. van Beijnum, "Misadventure at root I in China," 2010. [Online]. Available: https://web.archive.org/web/20110622092029/http://arstechnica.com/tech-policy/news/2010/03/china-censorship-leaks-outside-great-firewall-via-root-server.ars
[91] Netnod, "Host an I-root," 2022. [Online]. Available: https://www.netnod.se/host-an-i-root
[92] University of Oregon, "Route views project," 2022. [Online]. Available: http://www.routeviews.org/routeviews/
[93] C. Partridge and M. Allman, "Ethical considerations in network measurement papers," *Commun. ACM*, vol. 59, no. 10, pp. 58–64, Sep. 2016, doi: 10.1145/2896816.
[94] E. Kenneally and D. Dittrich, "The menlo report: Ethical principles guiding information and communication technology research," Available at SSRN 2445102, 2012.
[95] W. A. Kumari and P. E. Hoffman, "Running a root server local to a resolver," RFC 8806, Jun. 2020. [Online]. Available: https://rfc-editor.org/rfc/rfc8806.txt
[96] A. Hilton, C. Deccio, and J. Davis, "Fourteen years in the life: A root server's perspective on DNS resolver security," in *Proc. 32nd USENIX Secur. Symp.*, Anaheim, CA, 2023, pp. 3171–3186. [Online]. Available: https://www.usenix.org/conference/usenixsecurity23/presentation/hilton
[97] R. Chhabra, P. Murley, D. Kumar, M. Bailey, and G. Wang, "Measuring DNS-over-HTTPS performance around the world," in *Proc. 21st ACM Internet Meas. Conf.*, New York, NY, USA, 2021, pp. 351–365, doi: 10.1145/3487552.3487849.
[98] D. Dagon, C. Lee, W. Lee, and N. Provos, "Corrupted DNS resolution paths: The rise of a malicious resolution authority," in *Proc. 15th Netw. Distrib. Syst. Secur. Symp.*, San Diego, CA, 2008. [Online]. Available: http://www.citi.umich.edu/u/provos/papers/ndss08_dns.pdf
[99] R. Houser, S. Hao, Z. Li, D. Liu, C. Cotton, and H. Wang, "A comprehensive measurement-based investigation of DNS hijacking," in *Proc. 40th Int. Symp. Reliable Distrib. Syst.*, 2021, pp. 210–221.
[100] S. Castro, D. Wessels, M. Fomenkov, and K. Claffy, "A day at the root of the internet," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 5, pp. 41–46, 2008.
[101] G. C. M. Moura et al., "Anycast vs. DDoS: Evaluating the Nov. 2015 Root DNS Event," in *Proc. ACM Internet Meas. Conf.*, Santa Monica, CA, USA, 2016, pp. 255–270. [Online]. Available: http://dl.acm.org/citation.cfm?id=2987446
[102] M. Wicaksana, "Ipv4 vs ipv6 anycast catchment: A root dns study," Aug. 2016. [Online]. Available: http://essay.utwente.nl/70921/
[103] Y. Yu, D. Wessels, M. Larson, and L. Zhang, "Authority server selection in DNS caching resolvers," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 2, pp. 80–86, 2012.
[104] M. Müller, G. C. Moura, R. de O. Schmidt, and J. Heidemann, "Recursives in the wild: Engineering authoritative DNS servers," in *Proc. Internet Meas. Conf.*, 2017, pp. 489–495.
[105] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig, "Comparing DNS resolvers in the wild," in *Proc. 10th ACM SIGCOMM Conf. Internet Meas.*, 2010, pp. 15–21.

Fenglu Zhang (Member, IEEE) received the BE degree from Sun Yat-sen University, China, in 2020. He is currently working toward the PhD degree with the Institute for Network Sciences and Cyberspace, Tsinghua University, China. In addition, he holds a position as a visiting scholar with UC, Irvine. His research focuses on network security and Internet measurement. As the first author, he contributed to a study that received the Distinguished Paper Award with ACM Conference on Computer and Communications Security (ACM CCS), in 2023.

Baojun Liu (Member, IEEE) received the PhD degree from Tsinghua University, in 2020. He is an assistant professor with Tsinghua University. He has been a visiting research scholar with the International Computer Science Institute (ICSI), UC Berkeley. His research covers a range of topics in network security, Internet measurement, and data analysis. He is currently a caucus member of ICANN Root Server System Advisory Committee (RSSAC).

Chaoyi Lu received the PhD degree from Tsinghua University, in 2022. He is a postdoctoral researcher with Tsinghua University. His research interests include network security and Internet measurement. One of his studies was honored with the IRTF Applied Networking Research Prize (ANRP), in 2020.

Yunpeng Xing received the BE degree from Beihang University, in 2023. He is currently working toward the master's degree with the Institute for Network Sciences and Cyberspace, Tsinghua University, China. His research interests include DNS Security and Internet Measurement.

Haixin Duan (Member, IEEE) received the PhD degree in computer science from Tsinghua University, and then became a faculty member with Tsinghua University. He has been a visiting scholar with UC Berkeley and a senior scientist with International Computer Science Institute (ICSI). He focuses his research on network security, including security of network protocols (DNS, Web, HTTP, and HTTPS). Most of his papers are published in the top security conferences (Oakland S&P, USENIX Security, CCS, and NDSS).

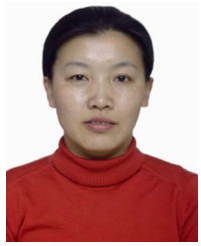

Ying Liu (Member, IEEE) received the MS degree in computer science and the PhD degree in applied mathematics from Xidian University, China, in 1998 and 2001, respectively. She is currently a professor with Tsinghua University, China. Her major research interests include network architecture design, next-generation Internet architecture, routing algorithms, and network protocols.

Liyuan Chang is a senior researcher with China Telecom. His research interest lies in Cyberspace security, including network security, system security, application security, and cloud security.