SAC101v2
SSAC Advisory Regarding Access to Domain Name Registration Data

An Advisory from the ICANN Security and Stability Advisory Committee (SSAC)
12 December 2018

Preface

This is an advisory to the ICANN Board, the ICANN Organization staff, the ICANN community, and, more broadly, the Internet community from the ICANN Security and Stability Advisory Committee (SSAC) about access to domain name registration data and Registration Data Directory Services (RDDS). Per the ICANN Bylaws, one of SSAC's roles is "To make policy recommendations to the ICANN community and Board."

The SSAC focuses on matters relating to the security and integrity of the Internet's naming and address allocation systems. This includes operational matters (e.g., pertaining to the correct and reliable operation of the root zone publication system), administrative matters (e.g., pertaining to address allocation and Internet number assignment), and registration matters (e.g., pertaining to registry and registrar services). SSAC engages in ongoing threat assessment and risk analysis of the Internet naming and address allocation services to assess where the principal threats to stability and security lie, and advises the ICANN community accordingly. The SSAC has no authority to regulate, enforce, or adjudicate. Those functions belong to other parties, and the advice offered here should be evaluated on its merits.

Version 2 of SAC101 was published to reflect evolving circumstances related to ICANN's Temporary Specification for gTLD Registration Data, and the ongoing Expedited Policy Development Process (EPDP) on the Temporary Specification for gTLD Registration Data. Version 1 of SAC101 has been retired and version 2 is authoritative.

Table of Contents

Executive Summary

Reliable, consistent, and predictable access to domain name registration data (via *Registration Data Directory Services*, or RDDS) is essential for a variety of legitimate purposes, especially the identification and mitigation of various types of Internet abuse and technical problems. In recent years, access to the data by those who have a legitimate need for it has been diminished, and availability is more constrained and more restricted than ever. This has happened for two main reasons: new legal and policy developments, and an operational practice known as *rate limiting*. These developments represent a shift in the abuse detection and mitigation landscape. The ability of security practitioners and law enforcement to detect and mitigate cybercrime and DNS abuse has been negatively affected, and the current situation is imposing greater operational and administrative burdens on those defenders. This in turn impairs the general usability and trustworthiness of the Domain Name System (DNS) and the Internet. This document describes the results of SSAC's extended deliberation on RDDS access issues and offers recommendations for how to move forward.

One of the fundamental unresolved policy issues is to determine which parties are allowed to see what registration data, under what circumstances. Sometimes the answer is shaped by local laws. While legal obligations are a reality and must be complied with, SSAC is concerned that access to registration data has been diminished far further than legal obligations require, and further than is prudent for responsible stewardship of the namespace. SSAC believes that ICANN has an obligation to ensure the continued availability of gTLD registration data to the greatest extent possible, and does not believe that ICANN's new Temporary Specification for gTLD Registration Data delivers on that need.[1] This document concludes that for the gTLD space under ICANN's remit, ICANN must urgently pursue solutions to several long-standing policy and technical problems described below. The policy-making process must avoid past process failures.

This paper also describes how rate-limiting affects legitimate access to registration data. Rate-limiting is designed to limit the amount of data a requestor can obtain, and/or how quickly the requestor can obtain it. The practice of rate-limiting is a result of unresolved issues, notably the need to define the legitimate uses of gTLD registration data, ways of recognizing legitimate users and restricting access by abusers, and the implementation of technical means to deliver differentiated sets of data to differently permissioned users. This advisory provides background about what rate-limiting is, how it is implemented, the policy and technical issues involved, and the obstacles such practices present to activities essential to the security and stability of today's Internet. We highlight examples of both legitimate and unprincipled reasons for rate-limiting access to RDDS data, and review the current contractual obligations of RDDS operators, gaps in which contribute to the problem.

The recommendations included at the end of this document, and summarized here, are designed to help carry out ICANN's core mission, which is to "coordinate the development and implementation of policies for which uniform or coordinated resolution is reasonably necessary

[1] See "Temporary Specification for gTLD Registration Data", https://www.icann.org/resources/pages/gtld-registration-data-specs-en

to facilitate the openness, interoperability, resilience, security and/or stability of the DNS."[2] In making these recommendations, we have endeavored to balance the various legitimate public and private interests at stake, including privacy, security, and accountability. *Please see Section 7 for the full recommendations.*

Recommendation 1: The ICANN Board, ICANN Organization, and ICANN community must solve long-deferred problems regarding domain registration data and access to it. SSAC recommends that the ICANN Board oversee the creation and execution of a plan that accomplishes the following interconnected tasks in a coordinated fashion, with timely deadlines. The creation and execution of this plan should be a top priority of the ICANN Board, ICANN Organization, and ICANN community.

- **A. ICANN policy-making should result in a domain registration data policy, including statements of purposes for the collection and publication of the data.**
- **B. The ICANN Board and the ICANN Organization should require contracted parties to migrate from using the WHOIS protocol to using the RDAP protocol.**
- **C. The remaining thin gTLD registries should be required to move to thick status, per the Thick WHOIS Consensus Policy and Board Resolution 2014.02.07.08.**
- **D. The ICANN Board should support the creation of an accredited RDDS access program, with the ICANN Organization ensuring the creation, support of, and oversight of the supporting technical access mechanism.**

Recommendation 2: The ICANN Board should direct the ICANN Organization to work with the ICANN Community to: A) develop policy with clearly defined uniform purposes for RDDS rate-limiting and corresponding service level agreement requirements and B) clarify current expectations for the use of rate limiting under existing policy and agreements.

Recommendation 3: The ICANN Board and EPDP policy-makers should ensure that security practitioners and law enforcement authorities have access to domain name contact data, via RDDS, to the full extent allowed by applicable law.

Recommendation 4: The initiation of charges for RDS access, or any significant future changes in fees for RDDS access, must include a formal assessment of user impacts and the security and stability impacts, and be conducted as part of a formal Policy Development Process (PDP).

Recommendation 5: The SSAC reiterates Recommendation 2 from SAC061: "The ICANN Board should ensure that a formal security risk assessment of the registration data policy be conducted as an input into the Policy Development Process. A separate security risk assessment should also be conducted regarding the implementation of the policy." These assessments should be incorporated in PDP plans at the GNSO.

[2] Bylaws for Internet Corporation for Assigned Names and Numbers, Section 1.1 Mission. https://www.icann.org/resources/pages/governance/bylaws-en/#article1

Recommendation 6: The ICANN Board should direct the ICANN Organization to work to ensure that all methods of access to RDDS data provide an equivalent response to the same query.

Recommendation 7: The ICANN Board should direct the ICANN Organization to work to ensure that RDDS access is provided in a measurable and enforceable framework, which can be understood by all parties.

2 Background

Access to Top-Level Domain (TLD) name registration data in a reliable, consistent, and predictable fashion is essential for a variety of legitimate purposes. For this reason, ICANN historically maintained a policy of providing "timely, unrestricted and public access" to that data. The importance of the data for security and stability purposes has been a central concern of SSAC, and has been a frequent topic of SSAC's attention.[3]

Domain name registration data is provided through RDDS operated by TLD registrars and registry operators. Currently these are WHOIS servers operating on TCP port 43 and via web-based interface, and in the future there will be Registration Data Access Protocol (RDAP) servers. These services provide data that has a variety of legitimate uses, and the services and protocols were designed to provide machine access and to accommodate automated queries to accommodate those legitimate uses. Some parties also access the data for illegitimate purposes, such as to harvest contact data for spamming, and sometimes such users employ automated and high-volume queries as well.

In order to address the abuse, many registry operators and registrars restrict the number and frequency of queries that users of RDDS can make, a practice known as *rate-limiting*. Rate-limiting is designed to limit the total amount of data a requestor can obtain, and/or limit how quickly the requestor can obtain it. Rate-limiting is deployed to protect against denial-of-service of the RDDS and illegitimate access to the data. But rate-limiting tends to be deployed in ways that apply the same rules to all users, not discriminating between legitimate and illegitimate uses.[4] This indiscriminate rate-limiting denies data to those who have a legitimate need for it, and who are allowed to access the data under policy. In practice, this makes some legitimate users suffer due to the actions of abusers.

As described in more detail below, rate-limiting is widespread, and most gTLD domains are sponsored by registrars and registry operators that rate-limit. Industry practices are non-uniform.

---

[3] Over the last twelve years the SSAC has published twelve reports on the subject, all available at https://www.icann.org/groups/ssac/documents
"SAC014: Information Gathering Using Domain Name Registration Records"
"SAC023: Is the WHOIS Service a Source for email Addresses for Spammers?"
"SAC027: SSAC Comment to GNSO regarding WHOIS Studies"
"SAC033: Domain Name Registration Records and Directory Services"
"SAC038: Registrar Abuse Contacts"
"SAC051: SSAC Report on WHOIS Terminology and Structure"
"SAC054: SSAC Report on the Domain Name Registration Data Model"
"SAC055: WHOIS: Blind Men And An Elephant [SSAC Comment on the WHOIS Review Team Final Report]"
"SAC058: SSAC Report on Domain Name Registration Data Validation"
"SAC061: SSAC Comment on ICANN's Initial Report from the Expert Working Group on gTLD Directory Services"
"SAC081: SSAC Response to Request for Input on Next Generation gTLD RDS to Replace WHOIS Policy Development Process (PDP)"
"SAC087: SSAC Response to the GNSO Policy Development Process (PDP) Working Group on Next Generation gTLD Registration Directory Services – Second Outreach"

[4] For example, a prohibition on all bulk searching does not discriminate between legitimate and illegitimate uses. See "What is WhoIs Misuse and What are WhoIs Anti-Harvesting techniques?", https://blog.resellerclub.com/what-is-whois-misuse-and-what-are-whois-anti-harvesting-techniques/

The basic accessibility of the data has been a long-standing problem. Rate-limiting began to be implemented around 2005. By 2011, ICANN's Compliance Department found that, "A notable number of registrars were limiting public access to WHOIS data in an extreme manner."[5] In 2010, the GNSO's Registration Abuse Policy Working Group found that:

> .. the basic accessibility of WHOIS [service] has an inherent relationship to domain registration process abuses, and is a key issue related to the malicious use of domain names. It appears that WHOIS data is not always accessible on a guaranteed or enforceable basis, is not always provided by registrars in a reliable, consistent, or predictable fashion, and that users sometimes receive different WHOIS results depending on where or how they perform the lookup. These issues interfere with registration processes, registrant decision-making, and with the ability of parties across the Internet to solve a variety of problems.[6]

The SSAC finds that rate-limiting practices continue to pose problems in all those areas, as described in more detail below.

Rate-limiting is separate from, but related to, the policy issue of which parties are allowed to see what registration data. Over the last year the ICANN community has become much more aware of national data protection laws, notably the European Union's General Data Protection Regulation (GDPR). In order to abide by local laws, SSAC assumes that gTLD registration data will become available in a gated or gated model. Such a model will make a limited set of data available to the public via anonymous access, while more data (specifically more personally identifiable data) will be available to authorized users. Gated access for identified legitimate users has been discussed for many years, and the RDAP protocol was designed in part to enable such access.[7]

On 17 May 2018, the ICANN Board adopted the Temporary Specification for gTLD Registration Data.[8] The stated goal of this specification is to establish temporary requirements for how ICANN and its contracted parties will continue to comply with existing ICANN contractual requirements and with community developed policies as they relate to WHOIS, while also complying with the GDPR. SSAC is not convinced that the Temporary Specification meets the ICANN Organization's stated goal to "ensure the continued availability of the WHOIS system to the greatest extent possible while maintaining the security and stability of the Internet's system of unique identifiers."[9]

---

[5] See "ICANN Contractual Compliance Whois Access Audit Report Sept 2010 – Feb 2011", https://www.icann.org/en/system/files/newsletters/update-whois-access-audit-report-port43-06apr11-en.pdf, page 6

[6] See "Registration Abuse Policies Working Group Final Report", page 6, https://gnso.icann.org/en/issues/rap/rap-wg-final-report-29may10-en.pdf

[7] See "Final Report from the Expert Working Group on gTLD Directory Services: A Next-Generation Registration Directory Service (RDS)", https://www.icann.org/en/system/files/files/final-report-06jun14-en.pdf

[8] See "Temporary Specification for gTLD Registration Data", https://www.icann.org/resources/pages/gtld-registration-data-specs-en

[9] See "ICANN Board Approves Temporary Specification for gTLD Registration Data", https://www.icann.org/news/announcement-2018-05-17-en

The policy and technical problems described above are symptoms of long-delayed issues that ICANN needs to solve now, including:

- The lack of an overall registration data policy that states the purposes and acceptable uses of the data, and allows balanced compliance with privacy law. SSAC advised that ICANN address this problem in 2012.[10]
- The deferred rollout of RDAP,[11] which will provide a technical means to authenticate users and provide gated access.
- The delayed implementation of the Thick WHOIS Policy PDP,[12] which is necessary to bring more uniformity to data storage and accessibility.

Except where noted, comments in this paper are relevant to access via both WHOIS port 43 and RDAP.

3 Uses of Domain Registration Data for Security and Stability Purposes

Access to domain name registration records is vital for a variety of legitimate uses. Specifically, TLD RDDS are critical for the stability and security of the Internet because they allow for the identification and mitigation of malicious activity, and the correction of problems that negatively affect services and users online. Such uses related to security and stability include but are not limited to:

- investigation of cybercrime and fraud,
- mitigation of DNS abuse,
- detection and correction of technical problems and security breaches related to domain names, name servers, and the services that depend upon them,
- maintenance of protective systems, such as Reputation Block Lists (RBLs) and domain reputation scoring mechanisms to combat spam and other threats,
- awarding digital certificates,
- ensuring the deliverability of valid email, and
- research on topics such as DNS traffic, botnets, distributed denial of service (DDoS) attacks, and Internet adoption and use.

We use the term "security practitioners" to designate those who have a responsibility to perform the above types of functions. These include network operators; those running commercially available products and services; registry operators who run anti-abuse programs in their TLDs; law enforcement personnel and other investigators; and academic researchers.

---

[10] See "SAC055: WHOIS: Blind Men And An Elephant [SSAC Comment on the WHOIS Review Team Final Report]"

[11] See "Registration Data Access Protocol", https://www.icann.org/rdap and Reconsideration Request from gTLD Registries Stakeholder Group (RySG), https://www.icann.org/en/system/files/files/reconsideration-16-10-rysg-request-redacted-09aug16-en.pdf

[12] See "Thick WHOIS Transition Policy for .COM, .NET and .JOBS", https://www.icann.org/resources/pages/thick-whois-transition-policy-2017-02-01-en

The above purposes often require fast, automated access to domain registration data. The data is used by systems that must react quickly to security incidents (such as the maintenance of reputation lists used in firewalls), or because large numbers of records are needed due to the sheer volume of security incidents that occur each day. These scoring systems need to cross-reference and find patterns across domain records, and research often requires large data sets to establish statistical significance.[13] For these purposes web-based RDDS access is not practical, thus fast, reliable access via WHOIS, and in the future access via RDAP, is vital.

Similarly, rate-limiting is a common problem for most security and Internet measurement research efforts that perform online data gathering.[14]

For similar reasons, the ICANN Governmental Advisory Committee (GAC) formally advised the ICANN Board in March 2018 to ensure that rate-limiting does not impair such access, stating that ICANN should "Ensure that limitations in terms of query volume envisaged under an accreditation program balance realistic investigatory cross-referencing needs."[15]

Reiterating SAC055, the SSAC believes that law enforcement and security practitioners have a legitimate need to access the real identity of the responsible party(ies) for a domain name.[16] Such access must comply with legal requirements. The GDPR, for example, contains provisions to allow such access and balance the various legitimate public and private interests at stake, including privacy, security, and accountability. Processing for security purposes is allowed under GDPR Recitals 47, 49 and 50, which allow uses including but not limited to "preventing fraud," "ensuring network and information security," the ability to resist "unlawful or malicious

---

[13] See the following representative references concerning legitimate uses for security and stability, and the need for automated volumetric requests.
"Final Report from the Expert Working Group on gTLD Directory Services: A Next-Generation Registration Directory Service (RDS)", https://www.icann.org/en/system/files/files/final-report-06jun14-en.pdf
"Statement by the EC3 Advisory Group on Internet Security", https://www.icann.org/en/system/files/files/gdpr-comments-ec3-icann-proposed-compliance-models-02apr18-en.pdf
"National Crime Agency feedback on ICANN's Proposed Interim Models for Compliance with the EU's General Data Protection Regulation", https://www.icann.org/en/system/files/files/gdpr-comments-nca-icann-proposed-compliance-models-29jan18-en.pdf
"Reputation Block Lists: Protecting Users Everywhere", https://www.icann.org/news/blog/reputation-block-lists-protecting-users-everywhere
Anti-Phishing Working Group: "Advisory on Utilization of Whois Data For Phishing Site Take Down", http://docs.apwg.org/reports/apwg-ipc_Advisory_WhoisDataForPhishingSiteTakeDown200803.pdf
Anti-Phishing Working Group: "WHOIS Tiered Access and Accreditation Program", https://www.icann.org/en/system/files/files/gdpr-comments-apwg-accreditation-access-non-public-whois-data-05apr18-en.pdf

[14] See Liu, Suqi, Ian D. Foster, Stefan Savage, Geoffrey M. Voelker and Lawrence K. Saul: "Who is .com?: Learning to Parse WHOIS Records." Proceedings of the 2015 Internet Measurement Conference, pages 369-380. Section 4.1 and references. https://ian.ucsd.edu/papers/imc15-whois.pdf . Also see Lauinger, Tobias et al: "Game of Registrars: An Empirical Analysis of Post-Expiration Domain Name Takeovers." Proceedings of the 26th USENIX Security Symposium, pages 865-880. https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-lauinger.pdf

[15] See "GAC Communiqué – San Juan, Puerto Rico, 15 March 2018", https://www.icann.org/en/system/files/correspondence/gac-to-icann-15mar18-en.pdf

[16] See "SAC055: WHOIS: Blind Men And An Elephant [SSAC Comment on the WHOIS Review Team Final Report]", pages 8-9.

actions", reporting possible "criminal acts or threats to public security" to authorities, and uses "in the public interest."[17] Articles 40 to 43 of the GDPR describe the creation of codes of conduct and accreditation mechanisms to assure that private data can be accessed properly for such purposes by qualified parties.[18]

In July 2018 the European Data Protection Board wrote to ICANN Org and affirmed that "the personal data processed in the context of WHOIS can be made available to third parties who have a legitimate interest in having access to the data, provided that appropriate safeguards are in place to ensure that the disclosure is proportionate and limited to that which is necessary and the other requirements of the GDPR are met..."[19]

On 28 June 2018, the ICANN GAC issued the following Consensus Advice to ICANN: "The GAC advises the ICANN Board to: Take all steps necessary to ensure the development and implementation of a unified access model that addresses accreditation, authentication, access and accountability, and applies to all contracted parties, as quickly as possible."[20]

4 Purposes of Domain Registration Data Rate-Limiting

SSAC members researched the purposes and implementation of rate-limiting by communicating with registry and registrar operators via in-person interviews, telephone conversations, and email correspondence. The SSAC also researched the public record.

SSAC concludes that rate-limiting on public-facing RDDS is imposed for two legitimate reasons:

1. To protect the service from an abusive volume of queries that might overwhelm the service or impact its ability to meet contractual service levels. Registry operators and registrars are responsible for meeting Service Level Agreements (SLAs) for RDDS responsiveness and uptime.[21] The purpose of SLAs is to ensure that users can reach the service and receive responses in a timely fashion.
2. To prevent abuse of the data by limiting access to it. A frequently cited example is the need to keep domain contact data out of the hands of spammers, who mine the data for email addresses.[22]

[17] See GDPR Recitals, https://gdpr-info.eu/recitals/

[18] See GDPR text, http://eur-lex.europa.eu/legal-content/en/TXT/?uri=CELEX:32016R0679

[19] See "Letter to Göran Marby from the European Data Protection Board", https://www.icann.org/en/system/files/correspondence/jelinek-to-marby-05jul18-en.pdf

[20] GAC Communiqué – Panama City, Panama, https://gac.icann.org/contentMigrated/icann62-panama-communique

[21] See "Base Registry Agreement, Specification 10: Registry Performance Specifications", https://newgtlds.icann.org/sites/default/files/agreements/agreement-approved-31jul17-en.html and "2013 Registrar Accreditation Agreement, Registration Data Directory Service (WHOIS) Specification", https://www.icann.org/resources/pages/approved-with-specs-2013-09-17-en#whois

[22] See "SAC023: Is the WHOIS Service a Source for email Addresses for Spammers?", and "Guidance for Registrars on Use of the Whois Service", https://registrar-console.centralnic.com/pub/whois_guidance and "Donuts Whois Access Policy", https://donuts.domains/about/policies/whois-access-policy/

Rate-limiting can also be imposed for problematic reasons, which negatively impact legitimate users. These include but are not limited to:

1. To shield registration records from compliance scrutiny.
2. Operating an under-provisioned service. Serving registration data is a core responsibility of registry operators and registrars, and they must provision their systems to handle a reasonable load.
3. Anti-competitive reasons, especially those that impact third parties. For example, some registrars have rate-limited access from other registrars, leading to concerns that the rate-limiting was imposed to impede domain transfers and impact registrant rights.[23]

4.1 How Rate-Limiting is Implemented, and Impact on Security

Most commonly, server operators set their systems to impose limits based on how many queries come from the same IP address (from a requestor or source), usually within a defined time period. To restrain the queries coming from a particular user or network, server operators often count queries coming from a IPv4 /24 prefix or IPv6 /48 prefix as coming from the same source.

The service operator will then reduce or suspend access if a source violates one or more parameters, usually based on queries per time period. **These parameters limit the total amount of data a requestor can obtain, and/or how quickly the requestor can obtain it.**

For example, a server operator can:

1. deny access if a requestor makes more than 200 queries in a 24-hour period, and/or
2. deny access if a requestor makes more than five queries in a minute.

If an operator has both limits above in place, a requestor can completely lose access after making just five queries, if the queries are made too quickly.

When a requestor exceeds a limit, the server operator may impose a range of consequences. Some operators will refuse queries until the query rate falls below the limit. Some will block queries from the offending source completely for a period of time, such as for 24 hours, and then enable access afterwards. Some will permanently refuse queries from the offending source, blacklisting the IP address or the larger IP range within which the offending IP address is contained.

Rate-limiting can take other forms, such as the denial of certain data fields, or denying data for certain registered names.

- In January 2018 a large registrar stopped showing domain contact data in its port 43 output. The only way for users to obtain contact data was to query the registrar's web-

[23] See "GoDaddy accused of dragging feet on domain transfers, says there's nothing unusual going on", https://www.theverge.com/2011/12/26/2662250/godaddy-namecheap-whois and "GoDaddy Responds To Namecheap Accusations, Removes 'Normal'; Rate Limiting Block", https://techcrunch.com/2011/12/26/godaddy-responds-to-namecheap-accusations-removes-normal-rate-limiting-block/

based WHOIS lookup page, which is protected by CAPTCHA[24] and therefore allows only non-automated lookups.[25] The registrar imposed similar restrictions as far back as 2011.[26] In this case the registrar effectively rate-limited access to contact data by removing it from port 43 output and cutting off automatable access.

- In SSAC testing, a large registrar failed to provide WHOIS for domains in Redemption Grace Period. Its port 43 WHOIS server returned a "domain does not exist" message while for the same domains its web-based WHOIS returned: "For legal reasons, the Whois information for the requested domain cannot be provided." In this case the registrar did not provide data that registrants are supposed to have access to, especially so they can redeem their domains if desired. Based on the web-based WHOIS message, the registrar chose to deny access to this data. This is an example of unreliable and inconsistent service that the Thick WHOIS Policy was designed to address.[27]
- In SSAC testing, a different large registrar's WHOIS server sometimes returned the string "no match" when queried for .COM domains that were registered, in normal lifetime, and in the zone file. This incorrect response denied access to data that should have been available. This is an example of unreliable, inconsistent, and unpredictable service that the Thick WHOIS Policy was designed to address.

Rate-limits set by operators vary widely. At the low end, some registry operators allow only two queries a minute. Several large registry operators allow ten queries per minute. Some do not impose limits. For details, see *Section 6: Measuring Rate Limits*, *Appendix A: Examples of Stated Rate Limits*, and *Appendix B: SSAC Rate Limiting Measurements*.

Some registry operators impose a total limit across multiple TLDs, a practice the SSAC terms "pooling." An example is CentralNIC, which operates a number of TLDs itself and provides back-end services for others (e.g., .XYZ), and serves domain registration data for many of those TLDs from a shared system or server. CentralNIC imposes a maximum total query rate of 240 queries per hour per /24 IPv4 range for all TLDs it operates.[28] CentralNIC provides service for approximately 30 gTLDs, and so the average rate limit per-TLD is less than 10 queries per hour. A requestor can exhaust its queries on one TLD and be unable to query domains in another TLD CentralNIC operates.

Some registry operators and registrars offer privileged access so that parties can obtain data via port 43 at a rate higher than unprivileged users. This is usually accomplished by whitelisting the IP addresses of approved users at the public facing WHOIS server. Other operators offer private

[24] See "CAPTCHA", https://en.wikipedia.org/wiki/CAPTCHA

[25] See "GoDaddy to start masking some Whois data through Port 43", https://domainnamewire.com/2018/01/12/godaddy-start-masking-whois-data-port-43/ and "GoDaddy and DomainTools scrap over Whois access", http://domainincite.com/22510-godaddy-and-domaintools-scrap-over-whois-access

[26] See "Go Daddy Explains Whois Blocking", https://domainnamewire.com/2011/01/04/go-daddy-explains-whois-blocking/

[27] See "Thick WHOIS Transition Policy for .COM, .NET and .JOBS", pages 19-20, 22-25, 35-37 https://www.icann.org/resources/pages/thick-whois-transition-policy-2017-02-01-en

[28] See "Guidance for Registrars on Use of the Whois Service", https://registrar-console.centralnic.com/pub/whois_guidance

RDDS with access controlled some other way. Some operators make this opportunity known publicly,[29] while others offer it on a private or negotiated basis. This practice can assist legitimate users such as security providers and researchers, but the practice is not uniformly available. These trusted relationships must be pursued on a case-by-case basis and are time-consuming to achieve across the hundreds of TLD registries and registrars. In all cases the terms of the arrangement are at the discretion of the server operator.

Registrars often whitelist each others' IP addresses so that they can perform WHOIS queries to each others' servers, at rates higher than those available to other parties. This allows the registrars to obtain the contact email addresses needed to confirm registrar-to-registrar transfer requests as required by ICANN's Inter-Registrar Transfer Policy.[30] This practice is facilitated by the list of registrar IP addresses made available in ICANN's Registrar database.[31] This practice is an example of a set of legitimate users being allowed predictable access to the service.

The effectiveness of rate-limiting by IP depends in part on how determined and how well resourced the parties are. Because many server operators do not publish their rate-limits, users often *burn* IP addresses when they exceed a limit, which means the IP address can no longer be used for queries. Rate-limiting leads some requesters to make queries from a diverse range of IPs. This leads to an escalating arms race where those making queries must guess at the limits imposed by various server operators and make queries from many different IP ranges at diverse providers. Sometimes IP ranges that have been permanently burned and blacklisted are reassigned by ISPs or cloud providers to new customers, who then find themselves unable to make queries through no fault of their own.[32]

At the rate-limits currently being imposed in the industry, most users of RDDS can only observe a fraction of the activity taking place in many TLDs and registrar portfolios. This prevents security practitioners from finding and monitoring abusive domains.[33]

---

[29] *id.*

[30] See "Domain Name Transfers", https://www.icann.org/resources/pages/registrars/transfers-en

[31] See "Service Abuse", http://whois.blacknight.com/abuse.php and "ICANN Registrar Database", https://radar.icann.org/

[32] This happened to SSAC itself, when it attempted to use an IP range obtained from a hosting provider and ICANN staff discovered the IP range was already blacklisted by a WHOIS server operator.

[33] Some SSAC members have encountered this problem in their own professional work. See the following additional references from law enforcement and industry.
"Domain Name Abuse: How Cheap New Domain Names Fuel The eCrime Economy," RSA 2015, slides 15, 21, 30, 38. https://www.rsaconference.com/writable/presentations/file_upload/hta-r02-domain-name-abuse-how-cheap-new-domain-names-fuel-the-ecrime-economy_final.pdf.
"The Indispensable Role of Whois for Global Cybersecurity: Statement by the EC3 Advisory Group on Internet Security", https://www.icann.org/en/system/files/files/gdpr-statement-ec3-europol-icann-proposed-compliance-models-25jan18-en.pdf
"National Crime Agency feedback on ICANN's Proposed Interim Models for Compliance with the EU's General Data Protection Regulation", https://www.icann.org/en/system/files/files/gdpr-comments-nca-icann-proposed-compliance-models-29jan18-en.pdf

For example, Nominet imposes a limit of 1,000 WHOIS queries in a rolling twenty-four hour period at its public-facing .UK WHOIS server.[34] This is far fewer than the number of domains added to and removed from the .UK zone file each day. Thousands more domains change name servers each day, and there are 12 million domains in the .UK registry.[35]

ICANN's Domain Abuse Activity Reporting project (DAAR)[36] has been hampered by the WHOIS rate-limiting imposed by thick registry operators, which has prevented the project from finding out which registrars sponsor many domain names being listed for malicious activity such as phishing and malware.

5 RDDS Policy

5.1 Current ICANN Contractual Obligations Regarding Rate-Limiting

RDDS operators currently have wide practical latitude to set their own rate-limiting practices, to treat users differently, and are not required to distinguish between legitimate and abusive uses or users. As a result, service level metrics do not guarantee reliable service for legitimate users. There may also be weaknesses in how ICANN's compliance monitoring takes place.

ICANN's contracts state that registrars and registry operators must:

> … operate a WHOIS service available via port 43 in accordance with RFC 3912, and a web-based Directory Service at <whois.nic.TLD> providing free public query-based access to at least the following elements in the following format.[37]

And in providing this query-based public access to registration data, the Registrar Accreditation Agreement states:

> In providing query-based public access to registration data as required by Subsections 3.3.1 and 3.3.4, Registrar shall not impose terms and conditions on use of the data provided, except as permitted by any Specification or Policy established by ICANN. Unless and until ICANN establishes a different Consensus Policy, Registrar shall permit use of data it provides in response to queries for any lawful purposes *except to: ... (b) enable high volume, automated, electronic processes that send queries or data to the systems of any Registry Operator or ICANN-accredited registrar*.[38] [emphasis added]

[34] "These limits are not per IP address, they are per user." See "Nominet Acceptable Use Policy", https://registrars.nominet.uk/registration-and-domain-management/acceptable-use-policy

[35] See ".uk Register Statistics 2018", https://www.nominet.uk/news/reports-statistics/uk-register-statistics-2018/

[36] See "Domain Abuse Activity Reporting", https://www.icann.org/octo-ssr/daar

[37] See "Base Registry Agreement – Updated 31 July 2017: Specification 4, Registration Data Publication Services" paragraph 1, https://newgtlds.icann.org/sites/default/files/agreements/agreement-approved-31jul17-en.html and "2013 Registrar Accreditation Agreement, Registration Data Directory Service (WHOIS) Specification," paragraph 1, https://www.icann.org/resources/pages/approved-with-specs-2013-09-17-en#whois

[38] See "2013 Registrar Accreditation Agreement" section 3.3.5, https://www.icann.org/resources/pages/approved-with-specs-2013-09-17-en

This language is problematic because RDDS systems are correctly designed to provide "high volume, automated, electronic processes that send queries," and some high volume, automated queries are made for beneficial and lawful purposes. In the future, the language should be modified to distinguish between legitimate and abusive uses (or users) of the service and to not inhibit beneficial or lawful uses.

The above contractual language also allows, "*use* of data it provides in response to queries for any lawful purposes." [emphasis added]

The contracts also require that RDDS (port 43 WHOIS) service must respond with contact data:

> RDDS availability: Refers to the ability of all the RDDS services [sic] for the Registrar to respond to queries from an Internet user *with appropriate data* from the relevant registrar system.[39]
>
> Queries shall be about existing objects in the Registry System and *the responses must contain the corresponding information otherwise the query will be considered unanswered*.[40] [emphasis added]

At least one large registrar has interpreted the contractual language above to mean that it is allowed to serve only some of the required data elements via WHOIS port 43 as long as it serves all the required data elements via Web-based lookup.[41] This practice eliminates the ability of legitimate users to access important data via WHOIS port 43 or RDAP.

The contracts state that RDDS must not experience downtime of more than 864 minutes (14.4 hours) per month. The contracts also state that the round trip for a query, from the time the query is sent to the time the query is received back from the RDDS server, must take no more than four seconds.[42]

ICANN's Compliance Department operates an RDDS server audit system that monitors access to registrars' and registries' port 43 WHOIS servers.[43] This system operates within ICANN's IP

---

[39] See "2013 Registrar Accreditation Agreement" section 2.2.1, https://www.icann.org/resources/pages/approved-with-specs-2013-09-17-en

[40] See "gTLD Base Registry Agreement, Specification 10" section 4.6, https://newgtlds.icann.org/sites/default/files/agreements/agreement-approved-31jul17-en.html

[41] See "GoDaddy to start masking some Whois data through Port 43", https://domainnamewire.com/2018/01/12/godaddy-start-masking-whois-data-port-43/ and "Letter to Göran Marby et al regarding Critical RAA Violation; GoDaddy Port 43 Whois Masking & Throttling Programs", https://www.icann.org/en/system/files/correspondence/winterfeldt-to-chalaby-et-al-10mar18-en.pdf and "Letter to Brian J. Winterfeldt regarding Critical RAA Violation; GoDaddy Port 43 Whois Masking & Throttling Programs", https://www.icann.org/en/system/files/correspondence/icann-to-winterfeldt-05apr18-en.pdf

[42] See "2013 Registrar Accreditation Agreement, Registration Data Directory Service (Whois) Specification" section 2.2, https://www.icann.org/resources/pages/approved-with-specs-2013-09-17-en#whois

[43] See "ICANN Contractual Compliance WHOIS Access Audit Report Sept 2010 - Feb 2011", https://www.icann.org/en/system/files/newsletters/update-whois-access-audit-report-port43-06apr11-en.pdf and "WHOIS Access Monitoring Report (Port 43) Evaluation Period: July 2011 – February 2012", https://www.icann.org/en/system/files/newsletters/update-whois-access-audit-report-port43-30apr12-en.pdf

address space, but sometimes makes queries from outside of ICANN's IP address space. There are two weaknesses to this system:

1. Some registries and registrars whitelist ICANN's IP address space. For details, see *Section 6: Measuring Rate Limits*. This means that compliance monitoring from those IPs is not rate-limited and the system cannot detect if those RDDS operators are rate-limiting other users, or if they are responding to other users at all.
2. The Registrar Accreditation Agreement (RAA) says that ICANN will query a registrar's RDDS every 5 minutes, a rate of only 12 queries per hour. This means that the system cannot observe rate-liming imposed on more than 12 queries per hour.

Our conclusions are:

A. The contracts allow an RDDS server to be functionally down and non-responsive for one user but not for another. ("Down" meaning unresponsive to queries, or not providing all of the data required by contract.)
B. Rate-limiting is not adequately addressed in ICANN's contracts.

5.2 Rate-Limiting in a Gated Access System

In the future, it is likely that gTLDs will have RDDS provided under a gated, accredited access program. This will make a limited set of data available to the public via anonymous access, while more data (specifically more personally identifiable data) will be available to authorized users. The GDPR provides for such a program under Articles 40 to 43, involving "codes of conduct" and certification programs. On 5 July 2018, the European Data Protection Board advised ICANN to, "consider how all the requirements included in Chapter IV GDPR for Codes of Conduct and Accreditation shall be met to ensure that the envisaged mechanisms are fully compatible with the GDPR."[44]

In a gated access model, the rationales for imposing rate-limiting are greatly reduced; abusive parties will not have access to the gated service, approved parties will state their purposes (perhaps even for every query they make), the queries made by approved parties can be logged, and the approved parties will be responsible for their activities through law and possibly access agreements.

In such a situation, access should be driven by the legitimate needs of the authorized parties. If queries are legitimate, as defined by law and ICANN policy, then those queries are by definition allowable and it is the role of the server operator to serve them. There may need to be a process by which a server operator can demonstrate that a user's queries are excessive -- in violation of law or ICANN policy, or dangerous to the system. But generally, the SSAC recommends that users under gated access must be able to gain operational access to the registration data that policy says they are authorized to access, and must not be rate-limited unless the user poses a demonstrable threat to a properly resourced system.

[44] See "Letter to Göran Marby from the European Data Protection Board", https://www.icann.org/en/system/files/correspondence/jelinek-to-marby-05jul18-en.pdf

We note that query needs vary greatly from one security actor to another. And any given security actor's query volume needs can vary greatly from day to day, depending on what situations they are dealing with. This is another reason why server operators are not well equipped to determine the legitimate query needs of their users and impose rate limits.

5.3 Temporary Specification of May 2018

On 17 May 2018, the ICANN Board adopted the Temporary Specification for gTLD Registration Data.[45] The goal of this specification was to establish temporary requirements for how ICANN and its contracted parties will continue to comply with existing ICANN contractual requirements and with community-developed policies as they relate to WHOIS, while also complying with the GDPR. The specification contains several major policy shifts that have an impact on security.

SSAC is not convinced that the Temporary Specification meets ICANN Organization's stated goal to "ensure the continued availability of the WHOIS system to the greatest extent possible." Three reasons are:

1. The new policy allows RDDS operators complete freedom to choose when to redact domain contact data from publication, whether or not a domain contact is protected by GDPR or by any other local privacy law. The result has been blanket redactions, hiding more data than is legally called for. A more balanced and justified approach is needed.
2. The specification effectively sets aside the ICANN Procedure for Handling WHOIS Conflicts with Privacy Law, which was put in place in 2008 and revised by community process in 2016.[46] That Procedure was designed to "allow gTLD registry/registrars to demonstrate when they are prevented by local laws from fully complying with the provisions of ICANN contracts regarding personal data in WHOIS," thereby providing transparency about redaction decisions and justifying applications of privacy law. Use of this procedure will become more important in the years to come as privacy laws other than the GDPR are considered.
3. The Temporary Specification does not provide policy or technical mechanisms for law enforcement and security practitioners to retrieve data on a predictable and reliable basis.

These issues will require close examination in the next rounds of policy-making. RDDS access must comply with the law, but access should not be less timely, more restricted, and less public than law requires.

In regards to point #1 above, interpretations of the GDPR currently vary widely across the domain name industry, with different parties disagreeing about what data can be published in what circumstances. Not all of those interpretations can be correct. It is the responsibility of the ICANN Board, ICANN Organization, and the ICANN community to establish under GDPR what data fields can be legally published in public output and via gated access, and ensure uniform publication of those data fields with appropriate consent notifications.

---

[45] See https://www.icann.org/resources/pages/gtld-registration-data-specs-en and "ICANN Board Approves Temporary Specification for gTLD Registration Data", https://www.icann.org/news/announcement-2018-05-17-en

[46] See "Revised ICANN Procedure For Handling WHOIS Conflicts with Privacy Law", https://www.icann.org/resources/pages/whois-privacy-conflicts-procedure-2008-01-17-en

We note that under GDPR, RIPE-NCC continues to publish the contact data of both legal and natural persons in its thick WHOIS database, confident that it is complying with the law. RIPE-NCC has explained its purposes and rationales clearly, and notes that "Facilitating coordination between network operators (network problem resolution, outage notification etc.) is the one that justifies the publication of personal data in the RIPE-NCC Database."[47]

We also note that as of this writing, most ccTLD operators in the European Union continue to publish some (and sometimes all) contact data fields for domains registered by legal persons.[48] Some continue to publish some personal data for natural person registrants in public WHOIS output. These include:

- .EU: publishes email addresses of natural persons.[49]
- .DK: publishes names, addresses, and telephone numbers of natural persons, as per Danish law.[50]
- .NO: publishes email addresses of natural persons.[51]
- .MT: publishes all contact data of natural persons by default, with opt-out available.[52]

5.4 RDDS as a Critical Service, and Free Versus Fee Access

The data provided through RDDS is essential for the stability, security, and trustworthiness of the namespace. RDDS is a critical service provided as a public resource, for a wide variety of legitimate uses. As such, registrars and registry operators have always been obligated by contract to provide free public query access to RDDS. This view was reflected in the Affirmation of Commitments (May 2009-January 2017), under which ICANN pledged to maintain a policy of "timely, unrestricted and public access" and to regularly "assess the extent to which WHOIS policy is effective and its implementation meets the legitimate needs of law enforcement and promotes consumer trust."[53]

The view that RDDS is a core service rather than an ancillary or value-added service has been reflected in the definition of registry services used in the registry contracts and the Registry

[47] See "How We're Implementing the GDPR: The RIPE Database", https://labs.ripe.net/Members/Athina/how-we-re-implementing-the-gdpr-the-ripe-database and "How We're Implementing the GDPR: Legal Grounds for Lawful Personal Data Processing and the RIPE Database", https://labs.ripe.net/Members/Athina/gdpr-legal-grounds-for-lawful-personal-data-processing-and-the-ripe-database

[48] See "How all 33 European ccTLDs are handling GDPR", http://domainincite.com/23053-how-all-33-european-cctlds-are-handling-gdpr

[49] See "WHOIS Policy", https://eurid.eu/d/205797/whois_policy_en.pdf

[50] See "The Danish WHOIS database will remain illuminated", https://www.difo.dk/en/news/danish-whois-database-will-remain-illuminated

[51] See "Domain registration directory service", https://www.norid.no/en/personvern/domeneoppslag/

[52] See "Network Information Centre Malta approved terms and conditions", https://www.nic.org.mt/2LD/approved-terms

[53] See "AFFIRMATION OF COMMITMENTS BY THE UNITED STATES DEPARTMENT OF COMMERCE AND THE INTERNET CORPORATION FOR ASSIGNED NAMES AND NUMBERS", https://www.icann.org/resources/pages/affirmation-of-commitments-2009-09-30-en

Services Evaluation Policy for many years.[54] These define "registry services" as those services that are:

A. [..] operations of the registry critical to the following tasks: the receipt of data from registrars concerning registrations of domain names and name servers; provision to registrars of status information relating to the zone servers for the TLD; dissemination of TLD zone files; operation of the registry zone servers; and dissemination of contact and other information concerning domain name server registrations in the TLD as required by the Registry Agreement [..];
B. other products or services that the registry operator is required to provide because of the establishment of a Consensus Policy (as defined above);
C. any other products or services that only a registry operator is capable of providing, by reason of its designation as the registry operator.

In other words, the basic, "critical" registry services are operation of a Shared Registry System (SRS) itself, DNS resolution, and RDDS. This view was carried through in the recent new gTLD round, when applicants were evaluated on their ability to deliver five "critical registry functions": Shared Registry System (SRS), running the EPP protocol, WHOIS (RDDS), DNS, and data escrow.[55]

Since it is a core service, registry operators and registrars have always incorporated the cost of providing RDDS into their budgets as a basic cost of doing business. The cost of providing RDDS, DNS resolution, and the like have always been implicitly incorporated into the wholesale costs that registry operators charge their registrars for domain names, and the retail fees that registrars charge their registrants. In the new gTLD Program the expectation was made explicit, because all new gTLD applicants were required to provide detailed projected cash outflows (expenses) for the five critical registry functions.[56]

The GDPR clearly imposes new obligations on registrars and registry operators, and the need to provide gated access may be seen as a consequence of the law. Some have suggested that the cost of RDDS access, or just the cost of gated access for authorized users, should be paid for by those who use it.[57,58] This raises issues that go to the very nature of and assumptions about the service. Gated access with associated fees would be a major change to the core principles in place for many years.

---

[54] See "Registry Services Evaluation Policy", https://www.icann.org/resources/pages/registries/rsep/policy-en

[55] See "New gTLD Applicant Guidebook", https://newgtlds.icann.org/en/applicants/agb/guidebook-full-04jun12-en.pdf

[56] See "New gTLD Applicant Guidebook," Section IIb – Breakout of Critical Registry Function Operating Cash Outflows

[57] See "CoCCA to charge trademark owners for Whois access", http://domainincite.com/22860-cocca-to-charge-trademark-owners-for-whois-access and "Final Report from the Expert Working Group on gTLD Directory Services: A Next-Generation Registration Directory Service (RDS)" page 117, https://www.icann.org/en/system/files/files/final-report-06jun14-en.pdf

[58] "[Accred-Model] Philly Special 2.0", https://mm.icann.org/pipermail/accred-model/2018-May/000099.html

Such proposals raise issues that have impact on the security and stability of the domain name system. Some of these issues are:

- Are RDDS public resources maintained in the public interest, or is that view no longer valid?
- Should RDDS remain core services incorporated into the basic cost of a domain name?
- How will charges affect security actors? What may be the impact on the end-users protected by security practitioners?
- Would some types of users be charged while others are not? For example, would law enforcement or not-for-profit organizations be exempted, and how can such choices be justified?
- How would occasional users of the service be affected?
- What expectations are there for registrars and registry operators to bear responsibility for their abusive registrants, who security practitioners are tracking?
- Is it appropriate for the cost of malicious activity to be shifted from the attackers to the victims and defenders?

The cost of the queries required to locate and mitigate domain abuse could be prohibitively expensive and very difficult to manage operationally. Any decision to change the status quo away from free RDDS access to charging for RDDS would be a highly consequential decision, raises security and stability issues, and must be justified via open and robust community process.

There are three existing processes at ICANN by which registry operators could have their contracts modified to allow fee-based RDDS service; and two processes relevant to registrar contracts:

1. Policy Development Process (PDP). This is the most thorough and transparent process. It offers participation by any and all parties, including affected users.
2. The Registry Services Evaluation Process (RSEP). In the RSEP process a proposed registry service can be approved solely by ICANN Staff, and without a review of the service's security and stability implications by a Registry Services Technical Evaluation Panel (RSTEP). Even if an RSTEP evaluation takes place, the panel is mainly composed to evaluate technical aspects and is not situated to evaluate policy issues and social impacts. The RSEP process mostly takes place out of public view and does not allow participation by affected users.
3. Contractual negotiation. While there may be opportunities for public comment periods about proposed contract amendments, the negotiations take place out of community view with ICANN staff deciding priorities and tradeoffs. Security and stability risk assessments may not take place, and the interests of affected users are not directly represented.

SSAC recommends that the initiation of charges for RDDS access, or any significant future changes in fees for RDDS access, must include a formal assessment of user impacts and the security and stability impacts, and must be conducted as part of a formal Policy Development Process (PDP).

5.5 Uniform Access and Output

Since 2012, the ICANN community has made significant efforts to make the displayed output of registration data more uniform and predictable.[59] However, this work has been upended in the wake of the GDPR.

Uniform output is important for a variety of reasons. Uniform output ensures that important data fields are always displayed in server output (even if the data appearing in those fields is redacted or anonymized), and makes parsing easier when ingesting and analyzing multiple domain records. Important use cases include anti-abuse work, where it is vital to reach out to domain contacts when their domain hosting has been compromised by malefactors.

As noted in 5.3 above, interpretations of the GDPR currently vary widely across the domain name industry, and as a result are now producing non-uniform and unpredictable outputs for gTLD domain records.

An important community effort to help bring reliability and uniform service was the Thick RDDS (WHOIS) Transition Policy for .COM, .NET and .JOBS. An ICANN consensus policy, it was adopted in 2014 with implementation details posted in 2017. The main recommendation was that, "The provision of thick Whois services, with a consistent labelling and display as per the model outlined in specification 3 of the 2013 RAA, should become a requirement for all gTLD registries, both existing and future."[60]

The policy came about because the thick registry model holds advantages for security, stability, accessibility, accuracy, and registrant protection. The working group's final report stated:

> The thin model is thus criticized for introducing variability among Whois services, which can be problematic for legitimate forms of automation. It is this problem that prompted the IRTP B Working Group to recommend requiring thick Whois across incumbent registries –- in order to improve security, stability and reliability of the domain transfer process. A thick Whois model also offers attractive archival and restoration properties [escrow].... A thick Whois model also reduces the degree of variability in display formats. Furthermore, a thick registry is better positioned to take measures to analyze and improve data quality since it has all the data at hand.[61]

[59] These efforts include the standardized output specifications in the 2013 Registrar Accreditation Agreement, the Base Registry Agreement, "SAC054: SSAC Report on the Domain Name Registration Data Model", "Registration Data Access Protocol gTLD Profile", https://www.icann.org/resources/pages/rdap-gtld-profile-2016-07-26-en "Additional WHOIS Information Policy", https://www.icann.org/resources/pages/policy-awip-2014-07-02-en and "Registry Registration Data Directory Services Consistent Labeling and Display Policy", https://www.icann.org/resources/pages/rdds-labeling-policy-2017-02-01-en

[60] See "Thick WHOIS Transition Policy for .COM, .NET and .JOBS" Recommendation 1, https://www.icann.org/resources/pages/thick-whois-transition-policy-2017-02-01-en

[61] "Final Report on the Thick Whois Policy Development Process" pp. 11-12, https://gnso.icann.org/sites/default/files/filefield_42383/thick-final-21oct13-en.pdf

SSAC believes that thick registries hold those advantages over thin registries, and that the Thick WHOIS Transition Policy for .COM, .NET and .JOBS should be carried out unless there is some insuperable legal reason that will prevent it. SSAC has not yet seen compelling arguments that GDPR categorically prevents the cross-jurisdictional transfer of contact data from registrars to thick registries. Instead, there may merely be notification arrangements to data subjects that need to be put into place now in order to enable GDPR compliance.[62]

We note that ccTLD registries in the European Union remain thick with no apparent intention of becoming thin, and that they continue to accept registrations from outside their national borders and the borders of the EU.

5.6 Migration from WHOIS to RDAP

ICANN must replace the outdated WHOIS protocol with the RDAP protocol, which offers user authentication and other features that simplify gated access. The deployment of RDAP, combined with appropriate policies and operational procedures, offers a tool to solve some problems that rate-limiting causes for legitimate users.

6 Measuring Rate Limits

To understand the pervasiveness of rate-limiting and specific practices deployed, SSAC relied on an external study and also performed its own study.

The "WHOIS Misuse Study" was commissioned by ICANN, conducted by Carnegie Mellon University's Cylab in 2011-2013, and was published in 2014.[63] It offered a survey to 111 registrars and registry operators, to which a smaller number responded. Thirty-nine percent (9 of 23) reported implementing port 43 rate limiting, 52% (12) temporarily blacklisted users who exceeded rate limits for five to thirty minutes, and 30% (7) reported that they use permanent IP/domain blacklisting when necessary. The authors of the study also made port 43 queries themselves at 16 registrars in 2011, and reported that registrars allowed, on average, 83 WHOIS queries before blocking additional requests. The authors found a registry operator that began blocking after as few as four queries.

In 2017, the SSAC designed a study which was executed by SSAC support staff. These tests were executed during the months of November 2017 and March 2018. The test sent queries to the ten largest registrars and 100 largest registries by zone size at the time of testing. Zone size was determined by counting NS records in zone files available via the Centralized Zone Data

[62] On 25 October 2018, the ICANN Board resolved to further defer the implementation of the Thick WHOIS Consensus Policy. The Board Resolution's rationale section notes, "The deferral will allow additional time for Verisign, registrars and ICANN to reach agreement on the amendments needed to the registry-registrar agreements to implement the Policy. This deferred enforcement period will also allow the Expedited Policy Development Process Team to determine if the Temporary Specification for gTLD Registration Data should become an ICANN Consensus Policy while complying with the GDPR and other relevant privacy and data protection law.", https://www.icann.org/resources/board-material/resolutions-2018-10-25-en#1.d

[63] See "WHOIS Misuse Study Final Report" pages 40, 42-43,48, https://whois.icann.org/sites/default/files/files/misuse-study-final-13mar14-en.pdf and "Empirically Measuring WHOIS Misuse", https://pdfs.semanticscholar.org/d198/d4c179db374f2d7a4f829a2fdd196fe4b947.pdf

Service (CZDS). If a thick TLD's zone file was not available via CZDS that registry's WHOIS server was not tested.

The SSAC study executed port 43 queries at rates starting at one query per hour and slowly increasing in intervals to a maximum of 240 queries per hour. The test cases are described in *Appendix B*. Some queries were made from IP addresses within ICANN's corporate IP space, and others were made from three rented servers in public IP space, one each in North America, Europe, and Asia. The queries from the three rented servers were treated similarly by each WHOIS server queried. Support staff did not observe geographic bias.

To test registry WHOIS servers, support staff randomly chose domains to query from the gTLD zone files. To test registrar servers, support staff randomly chose domains from .COM and .NET sponsored by the registrars under test. Support staff generated a new set of domain names for each test case.

A query was counted as successful if the WHOIS server returned domain data; there was no attempt to validate the correctness of the returned data. A query was counted as failed if:

a) a query to a thick registry or to a registrar failed to return any registrant contact data, or
b) if the WHOIS server did not respond at all, or
c) a response was received indicating the use of rate limiting (such as "QUERY LIMIT EXCEEDED").

As noted in section 3.1, some operators served multiple TLDs from one system (i.e., pooling).

6.1 Testing Results

Detailed testing results are listed in *Appendix B*. The major findings of the SSAC designed study are:[64]

A. Access to most gTLD domain data is rate-limited.
    a. Four of the five largest registrars of .COM and .NET domains (GoDaddy, Tucows, HiChina, and Network Solutions) all rate-limited in staff's testing. Other large registrars such as Endurance/Resellerclub rate-limit as well. These registrars represent the majority of the .COM/.NET market.[65]
    b. Most operators of thick gTLD registries – including Afilias, CentralNIC, Donuts, Famous Four, Neustar, Nominet, and Public Interest Registry– impose rate limits.
B. The lowest limit was at the .REVIEW registry (operated by Famous Four), which allowed only two queries per hour.
C. The actual rate limits encountered during testing were sometimes more restrictive than the limits published by operators. For example, CentralNIC started rate-limiting at 165 queries per hour, which is less than the 240 per hour CentralNIC states is the limit for an

[64] It was not possible to determine precise rate-limits from the testing, because the precise time period over which the server had received too many queries could not be known. For example, when a server began to register failures due to rate-limiting, it was not possible to know if that was because too many queries had been received in the last hour or the last 48 hours. We assume a time period of one hour in these findings.

[65] See "Monthly Registry Reports", https://www.icann.org/resources/pages/registry-reports

anonymous/unrecognized IP address.[66] Tucows began rate-limiting at eight queries per hour, which is far below its stated "default limit" of "one (1) lookup per second".[67]

D. Some registrars and registry operators impose rate limits on the public while imposing no limits on ICANN. For example, one registrar allowed 240 queries per hour from ICANN's IP range but allowed only eight queries per hour when we conducted queries from rented servers in public IP space in the US, Europe, and Asia. As noted above, this has implications for SLA monitoring. See Appendix B for statistics.

Here are a few typical examples that will help readers interpret the statistics in *Appendix B*:

- Public Domain Registry (whois.myorderbox.com): Began to rate-limit at Level 4 (15 queries per hour; 150 total queries each performed 240 seconds apart). At that rate 89 queries succeeded and 59 failed from North America. At Level 5 (30 queries per hour; 300 total queries each performed 120 seconds apart), 98 queries succeeded and 201 queries failed from North America.
- Tucows (whois.opensrs.net): began to rate-limit at Level 3, (eight queries per hour, 80 total queries each one 450 seconds apart). At that rate, 54 queries succeeded and 26 failed. At Level 4 (15 queries per hour; 150 total queries each performed 240 seconds apart), 33 queries succeeded and 117 failed.

7 Recommendations

Recommendation 1: The ICANN Board, ICANN Organization, and ICANN community must solve long-deferred problems regarding domain registration data and access to it. SSAC recommends that the ICANN Board oversee the creation and execution of a plan that accomplishes the following interconnected tasks in a coordinated fashion, with timely deadlines. The creation and execution of this plan should be a top priority of the ICANN Board, ICANN Organization, and ICANN community.

A. **ICANN policy-making should result in a domain registration data policy, including statements of purposes for the collection and publication of the data.** This should clarify what uses are legitimate, and what data set will be published publicly versus by gated access. SSAC assumes that this work will be accomplished via the GNSO Policy Development Process (PDP/EPDP). The Board should work with the GNSO Council to ensure that the policy-making process avoids past process failures so that outcomes can be reached.
B. **The ICANN Board and the ICANN Organization should require contracted parties to migrate from using the WHOIS protocol to using the RDAP protocol.** Deploying RDAP will allow for the authentication of users and gated access. Existing contractual language allows ICANN to require such implementation.
C. **The remaining thin gTLD registries should be required to move to thick status, per the Thick WHOIS Consensus Policy and Board Resolution 2014.02.07.08.** Implementation will provide more uniform, predictable, and stable RDDS service.

[66] See "Guidance for Registrars on Use of the Whois Service", https://registrar-console.centralnic.com/pub/whois_guidance

[67] See "OpenSRS: WHOIS rate limiting", https://help.opensrs.com/hc/en-us/articles/204075306-WHOIS-rate-limiting

D. **The ICANN Board should support the creation of an accredited RDDS access program, with the ICANN Organization ensuring the creation, support of, and oversight of the supporting technical access mechanism.** This program will identify qualified users, enable their access under appropriate data protection measures, and will allow RDDS server operators to manage those users' access accordingly. The technical access mechanism should include a credential management system so that users do not need to negotiate and set up access with RDDS operators individually.

Recommendation 2: The ICANN Board should direct the ICANN Organization to work with the ICANN Community to: A) develop policy with clearly defined uniform purposes for RDDS rate-limiting and corresponding service level agreement requirements, and B) clarify current expectations for the use of rate limiting under existing policy and agreements.

Recommendation 3: The ICANN Board and PDP policy-makers should ensure that security practitioners and law enforcement authorities have access to domain name contact data, via RDDS, to the full extent allowed by applicable law.

Recommendation 4: The initiation of charges for RDDS access, or for any significant future changes in fees for RDDS access, must include a formal assessment of user impacts and the security and stability impacts, and must be conducted as part of a formal Policy Development Process (PDP). RDDS has vital uses in the public interest, and RDDS is a core service. Allowing contracted parties to charge for or significantly change access fees for RDDS service would be a highly consequential decision affecting diverse users, and presents a reasonable risk of a meaningful adverse effect on stability and security.

Recommendation 5: The SSAC reiterates Recommendation 2 from SAC061: "The ICANN Board should ensure that a formal security risk assessment of the registration data policy be conducted as an input into the Policy Development Process. A separate security risk assessment should also be conducted regarding the implementation of the policy." These assessments should be incorporated into PDP plans at the GNSO. Among other aspects, the risk assessment should assess how the policy will affect access to and use of domain registration data by law enforcement bodies and security practitioners.

Recommendation 6: The ICANN Board should direct the ICANN Organization to work to ensure that all methods of access to RDDS data provide an equivalent response to the same query. In particular, in the interest of security and stability, if a data field is published to a given user, ICANN should ensure that the registry or registrar publishes it via all contractually required RDDS access methods.

Recommendation 7: The ICANN Board should direct the ICANN Organization to work to ensure that RDDS access is provided in a measurable and enforceable framework, which can be understood by all parties. ICANN Compliance testing should also measure RDDS uptime and responsiveness as it is experienced from the public Internet perspective, and ICANN Compliance must have the means and tools necessary to enforce the policy.

8 Acknowledgments, Statements of Interests, and Objections and Withdrawals

In the interest of transparency, these sections provide the reader with information about four aspects of the SSAC process. The Acknowledgments section lists the SSAC members, outside experts, and ICANN staff who contributed directly to this particular document. The Disclosures of Interest section points to the biographies of all SSAC members, which disclose any interests that might represent a conflict—-real, apparent, or potential—-with a member's participation in the preparation of this Report. The Dissents section provides a place for individual members to describe any disagreement that they may have with the content of this document or the process for preparing it. The Withdrawals section identifies individual members who have recused themselves from discussion of the topic with which this Report is concerned. Except for members listed in the Dissents and Withdrawals sections, this document has the consensus approval of all of the members of SSAC.

8.1 Acknowledgments

The committee wishes to thank the following SSAC members for their time, contributions, and review in producing this report.

SSAC members
Greg Aaron
Joe Abley
Benedict Addis
Jaap Akkerhuis
Jeff Bedser
Don Blumenthal
Ben Butler
kc claffy
Jay Daley
James Galvin
Robert Guerra
Julie Hammer
Geoff Huston
Merike Kaeo
Warren Kumari
John Levine
Carlos Martinez
Danny McPherson
Ram Mohan
Russ Mundy
Dave Piscitello
Rod Rasmussen
Doron Shikmoni

ICANN staff
David Conrad

Andrew McConachie (editor)
Kathy Schnitt
Steve Sheng

8.2 Statements of Interest

SSAC member biographical information and Statements of Interest are available at: https://www.icann.org/resources/pages/ssac-biographies-2018-03-02-en

8.3 Objections and Withdrawals

There were no objections or withdrawals.

9 Revision History

9.1 Version 1

Version 1 of SAC101 was published June 14th, 2018.

9.2 Version 2

Version 2 of SAC101 was published to reflect evolving circumstances related to ICANN's Temporary Specification for gTLD Registration Data, and the ongoing Expedited Policy Development Process (EPDP) on the Temporary Specification for gTLD Registration Data. Version 1 of SAC101 has been retired and version 2 is authoritative.

The revisions included the addition of Section 5.2, references to subsequent ICANN Board actions, correspondence from the GAC and European Data Protection Board, and updates to the Recommendations section.

Appendix A: Examples of Stated Rate-Limits

- CentralNIC (multiple gTLDs): "Queries that come from 'untrusted'; sources are rate-limited at a maximum query rate of 240 queries per hour. Any queries in excess of this query rate will be refused until the query rate falls below this limit. If the query rate from an 'untrusted' source exceeds 300 queries per hour, the source is permanently blocked. Queries that come from "trusted" sources are rate-limited at a maximum query rate of 3,600 queries per hour. Any queries in excess of this query rate will be refused until the query rate falls below this limit."[68]
- Tucows: "The default limit for each IP or IP range is one (1) lookup per second. A daily rate-limit for each IP or IP range has also been set."[69]
  Public Interest Registry (.ORG): 30 queries per minute for the public (per correspondence). "Public Interest Registry accredited registrars who submit queries through the web-based WHOIS search mechanism are limited to 50 queries per minute."[70]
- Nominet (.UK): "5 queries per second with a maximum of 1,000 queries per rolling 24 hours."[71]
- Donuts (multiple gTLDs): ten queries per second.[72]
- SIDN (.NL): Registrars are limited in the number of domains they can query if the domains being queried are sponsored by other registrars: "extra-portfolio queries: 0.75% of portfolio (minimum 25, maximum 1,000 per day); A registrar with 10k domain names will therefore be allowed 75 extra-portfolio queries per day"[73]

[68] See "CentralNIC, Guidance for Registrars on Use of the Whois Service", https://registrar-console.centralnic.com/pub/whois_guidance

[69] See "OpenSRS, WHOIS rate limiting", https://help.opensrs.com/hc/en-us/articles/204075306-WHOIS-rate-limiting

[70] See "Public Interest Registry, Frequently Asked Questions", http://pir.org/resources/faq/

[71] See "Nominet, Acceptable Use Policy", https://registrars.nominet.uk/registration-and-domain-management/acceptable-use-policy

[72] See "Donuts, WHOIS limit", https://donuts.zendesk.com/hc/en-us/articles/201944666-WHOIS-limit?mobile_site=true

[73] See "SIDN, Impact of the GDPR on domain registration", https://www.sidn.nl/downloads/presentations/Synopsis%20of%20the%20webinar%20Impact%20of%20the%20GDPR%20on%20domain%20registration.pdf

Appendix B: SSAC Rate-Limiting Measurements

Test Cases

Case Number	Delay between queries (seconds)	Queries	Speed (queries/hour)	Duration (hours)
0	3600	10	1	10
1	1800	20	2	10
2	900	40	4	10
3	450	80	8	10
4	240	150	15	10
5	120	300	30	10
6	60	600	60	10
7	30	1200	120	10
8	15	2400	240	10

Data from SSAC Staff rate-limiting measurements taken February and March 2018 is available at the following locations in both Portable Document Format (PDF) and Microsoft Excel format.

https://www.icann.org/en/system/files/files/sac101-whois-rate-limit-testing-results-11jun18-en.pdf
https://www.icann.org/en/system/files/files/sac101-whois-rate-limit-testing-results-spreadsheet-11jun18-en.xlsx