TSG01: Technical Model for Access to Non-Public Registration Data

Technical Study Group on Access to Non-Public Registration Data

30 April 2019

Table of Contents

Abstract

The purpose of this document is to propose a technical solution for access to non-public domain name registration data in generic Top-Level Domains (gTLDs). This document contains a determination and scope of the nature of the problem, requirements necessary to address the issue, analysis of the solution space, and a high-level, technical proposal based on the Registration Data Access Protocol (RDAP) and OAuth 2.0 / OpenID Connect.

Executive Summary

Following the adoption of the European Union's General Data Protection Regulation (GDPR), ICANN organization (org) and the ICANN community have worked to balance the law's data protection requirements with the legitimate interests of third parties seeking access to non-public gTLD registration data. In designing a solution to balance these interests, it is important to reduce the potential liability faced by gTLD registries and registrars and ICANN when providing access to non-public gTLD domain name registration data.

In October 2018, ICANN President and CEO Göran Marby asked Ram Mohan, the Chief Operating Officer of Afilias and the former Security and Stability Advisory Committee liaison to the ICANN Board of Directors, to form a Technical Study Group on Access to Non-Public Registration Data (TSG) and asked the group to explore an implementation approach that would place ICANN as the coordinating party for third-party queries for non-public domain name registration data in the gTLD space.

The TSG's work is not an effort to replace the ICANN community's policy development process. Rather, the work of the Group is intended to help ICANN org determine whether such a model would diminish the legal liability for gTLD Contracted Parties (CPs), who would provide access to non-public gTLD domain name registration data.

Building on the technology available via the Registration Data Access Protocol (RDAP) and its extensions, the TSG recommends a technical model for authenticating, authorizing, and providing access to non-public gTLD domain name registration data to third parties with legitimate interests based on existing technologies. The technologies and various scenarios in which they could be used are explained in this paper.

The technical model would support a process that allows a Requestor to authenticate their identity and legitimate purpose for requesting data, come to a central service managed by ICANN, and receive approval or denial of the request. If approved, ICANN would ask the appropriate gTLD registry and/or registrar to provide all domain name registration data to ICANN, which in turn would filter it appropriately and return it to the requestor.

The TSG has not made decisions or recommendations on policy questions, e.g., which requestors get access, to which data fields, under what conditions should access be given, and what is a legal legitimate interest for requesting such data. These are not technical decisions or recommendations. The TSG has, with this document, delivered an outline of its working assumptions, requirements, and its proposed solution. The TSG is grateful for the support and robust feedback from the ICANN community, and now submits this document to the ICANN President and CEO for further consideration and appropriate action(s).

1. Background

The TSG explored technical solutions for authenticating, authorizing, and providing access to non-public gTLD domain name registration data for third parties with legitimate interests. The work focused on examining technical implementation solutions built on RDAP. In parallel with community efforts to develop an gTLD RDAP profile[1] prior to deployment of the protocol, the TSG focused its efforts on developing technical solutions for providing access to non-public gTLD registration data.

In a blog[2] published 24 September 2018, ICANN President and CEO Göran Marby wrote that ICANN is exploring possible technical solutions to be built on RDAP. This approach, upon which the TSG based its discussion, was further described during a data protection/privacy update webinar[3] held 8 October 2018. The implementation approach described during that webinar would place ICANN in the position of providing third-party access to non-public gTLD domain name registration data. If the query is approved, in accordance with relevant policy, ICANN would ask the appropriate gTLD registry or registrar to provide all domain name registration data to ICANN, which in turn would filter appropriately and provide it to the third party. This approach informed the TSG's subsequent deliberations.

The TSG, by design, has not made decisions or recommendations on purely policy questions, e.g., which users get access, to which data fields and under what conditions should access be given, and what is a legitimate interest for requesting such data. The TSG did consider the technical impact of policy choices, for example, policy recommendations that arose from Phase 1 of the Expedited Policy Development Process[4] (EPDP) on the Temporary Specification for gTLD Registration Data or other policy initiatives. Where there are multiple alternatives in either the proposed policy(ies) or the technical implementation(s), the TSG chose technologies that allow the technical model to be configured according to future policy choices.

With its adoption of the Temporary Specification for gTLD Registration Data[5] (Temporary Specification), the ICANN Board of Directors noted in the Annex as an important issue for further community action the development of "an accreditation and access model that complies with GDPR, while recognizing the need to obtain additional guidance from Article 29 Working Party/European Data Protection Board." The European Data Protection Board (EDPB), in a 5 July 2018 letter[6] to ICANN, also noted that "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, including the provision of clear information to data subjects."

[1] https://www.icann.org/gtld-rdap-profile

[2] https://www.icann.org/news/blog/icann-gdpr-and-data-protection-privacy-update

[3] https://www.icann.org/resources/pages/data-protection-meetings-2017-12-08-en

[4] https://community.icann.org/display/EOTSFGRD/EPDP+on+the+Temporary+Specification+for+gTLD+Registration+Data

[5] https://www.icann.org/resources/pages/gtld-registration-data-specs-en/

[6] https://www.icann.org/en/system/files/correspondence/jelinek-to-marby-05jul18-en.pdf

The EPDP, as part of its charter, plans to consider a standard access mechanism in Phase 2 of its work. To assist the community in its policy development work, ICANN org produced the Draft Framework for a Possible Unified Access Model[7] as a starting point for conversations with European data protection authorities, including the EDPB.

In addition, the Temporary Specification also directed the implementation of an RDAP service across the gTLD space. This service, which is set to launch on 26 August 2019, will be key to the development of a technical solution for providing third parties with a legitimate purpose with access to non-public gTLD domain name registration data through development of Contracted Party implementations that can be used as a foundation for addressing future requirements.

2. Conventions Used in this Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14[8] [RFC2119[9]] [RFC8174[10]] when, and only when, they appear in all capitals, as shown here.

2.1 Other Terms

Authentication Request - An OAuth 2.0 Authorization Request for Requestor authentication by an Identity Provider.

Authorization Endpoint - A service implemented by an Identity Provider to perform authentication of Requestors.

Access Token - An opaque data structure that is issued by an Identity Provider to allow authenticated access to service endpoints.

Authentication - The process or action of verifying the identity of a Requestor.

Authorization Grant - An OAuth 2.0 data structure that is used to exchange an Authorization Code for an Access Token.

Authorization - The process of specifying access rights/privileges to protected resources.

Authorization Code - An OAuth 2.0 data structure that is returned from an Authorization Endpoint to describe successful authentication of a requestor.

Browser User Agent - A web browser used by a Requestor to obtain an Access Token.

Consensus Policies - Policies established (1) pursuant to the procedure set forth in ICANN's Bylaws and due process, and (2) covering those topics listed in (i) Section 1.2 of the Consensus and Temporary Policies Specifications of the Registrar Accreditation Agreement and the Registry Agreements that are

[7] https://www.icann.org/en/system/files/files/framework-elements-unified-access-model-for-discussion-20aug18-en.pdf

[8] https://tools.ietf.org/html/bcp14

[9] https://tools.ietf.org/html/rfc2119

[10] https://tools.ietf.org/html/rfc8174

modeled after the Base Registry Agreement and (ii) comparable provisions in Registry Agreements that are not modeled after the Base Registry Agreement.[11]

Contracted Party (CP) - A registry operator of a generic Top Level Domain (gTLD) or an ICANN-accredited registrar that offer domain name registration services.

CP (Contracted Party) Servers - RDAP servers operated by gTLD domain registries and registrars.

HTTP Access Request - A HyperText Transfer Protocol (HTTP) operation using the GET or POST methods to obtain information from a resource.

ICANN Access Service - A browser-based web service used by the Requestors to obtain an Access Token from the OAuth/OpenID Connect process. In OAuth/OpenID Connect terms, this would be the Relying Party.

ICANN RDAP Gateway - A central RDAP proxy server through which all queries are directed and all responses are filtered. This server would receive RDAP queries and, depending on the outcome of an authorization check, forward that query to the appropriate Contracted Party Server. Any response from that server would then be returned to the user. It is **not** envisaged that the RDAP Gateway would store the response; such a "non-caching reverse proxy" has both security and privacy advantages.

ICANN Browser-based RDAP Client - A web-based interface RDAP client as an option to users who do not have their own RDAP clients

ICANN Browser-based Web Portal - A web-based interface for "exceptional" requests (requests not pre-authorized) which will be submitted – and reviewed by – a human.

ID Token - An OpenID Connect data structure that includes Requestor identity attributes, known as "claims".

Identification - The process of recognizing and naming someone or something.

Identity Providers - Organizations assigning credentials to and authenticating Requestors.

Public Domain Name Registration Data (Public Data) - Any registration data that is not required to be redacted by the Temporary Specification or any replacement Consensus Policy.

Reverse proxy: A type of proxy server that retrieves resources on behalf of a client from one or more servers. These resources are then returned to the client, appearing as if they originated from the proxy server itself. Unlike a forward proxy, which is an intermediary for its associated clients to contact any server, a reverse proxy is an intermediary for its associated servers to be contacted by any client.

Non-caching reverse proxy[12]: A reverse proxy which does not store copies of resources in a local cache, but fetches them anew from the origin server every time it receives a request from a client.

[11] See Registry Agreement, Specification 1, 1.1 and 2013 Registrar Accreditation Agreement, Consensus Policies and Temporary Policies Specification, 1.1.

[12] https://en.wikipedia.org/wiki/Reverse_proxy

Non-public gTLD domain name registration data - any registration data that is required to be redacted by the Temporary Specification or any replacement Consensus Policy.

RDAP User Agent - An RDAP client which uses an Access Token obtained by a Requestor to conduct RDAP queries (in some cases, this user agent may be an application in a web browser and indistinguishable from the Browser User Agent).

Requestors - The entities submitting queries, the results of which gain them access to non-public gTLD registration data.

Resource Server - As defined in RFC 6749 Section 1.1, a server hosting protected resources such as non-public gTLD domain name registration data records.

Third Party Authorizers - Organizations determining the data elements to be accessed by authenticated Requestors.

Token Endpoint - A service implemented by an Identity Provider to return ID and Access Tokens.

2.2 Document Naming Convention(s)

The TSG imagines a future where other Technical Study Groups may be formed, either to conduct further work on access to non-public registration data, or to work on other topics important to the ICANN community. The TSG proposes a document naming convention (TSG*nn*) so as to ensure continuity in the series, referenceability and naming uniqueness across documents. This document is named TSG01, and we suggest the next TSG output would be named TSG02, and so forth.

3. Architectural Assumptions

In formulating this project, ICANN org set forth specific architectural assumptions and requirements that it felt were important and central to its objectives. The TSG accepted these assumptions and requirements as given. The TSG developed additional assumptions that reflect the experience of the group. Many of these are based on best common practices as described in the list of references in Appendix 1.

The TSG acknowledges that some readers may question whether ICANN's architectural assumptions and requirements are necessary, optimal, or effective for the intended purpose. Such questions are appropriately addressed to ICANN org and are outside the scope of this document.

3.1 ICANN Organization's Architectural Assumptions

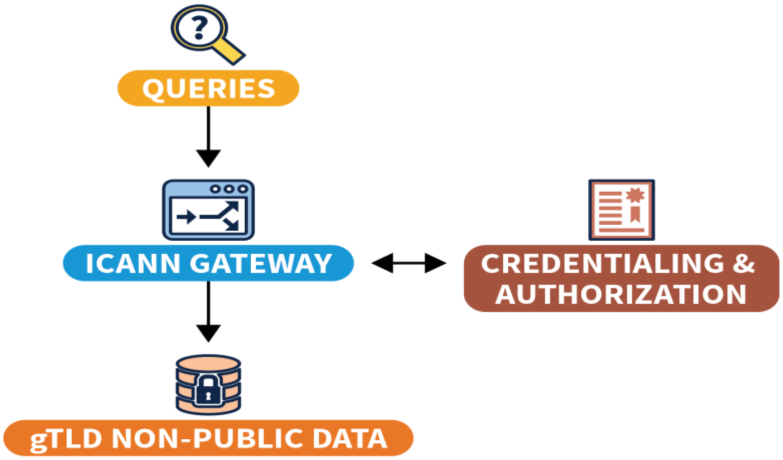

Figure 1.

1. RDAP will be used to access public and non-public gTLD domain name registration data; traditional “port 43” WHOIS (described in RFC 3912) will eventually be deprecated in gTLDs.
2. A standard model will apply equally to all parties requesting access to non-public gTLD domain name registration data.
3. ICANN org will be the sole party through which access to non-public domain name registration data is obtained in the gTLD space as part of a unified access model.
4. By acting as the sole party as described above, ICANN org thereby reduces CPs' legal liability arising from disclosure of non-public gTLD domain name registration data.

3.2 TSG's Additional Architectural Assumptions

5. The system MUST accommodate changes to data sets or data access.
6. All credentials will be protected in a reasonable way throughout their lifecycle.
7. For the purposes of access to non-public gTLD domain name registration data, scope of technical implementation for this solution is limited to RDAP, extension mechanisms to RDAP as defined by RFC 7480, 7481, 7482, and 7483, and other mechanisms an RDAP client implementer would find "natural" to implement.
8. It is expected that RDAP services provided by CPs will answer queries from unauthenticated sources, and when doing so will follow policies whereby data is redacted such as those listed in the Temporary Specification for gTLD Registration Data and the policy recommendations from the EPDP.
9. The solution will take into consideration the existing practices and currently deployed uses of RDAP.

10. The RDAP pilot working group will complete its RDAP profile, and such output will be ratified through the appropriate processes.
11. ICANN will ensure validity of credentials.
12. Policy choices may change the technical implementation of the proposed technical model.
13. The system will log all meta-data related to queries, and responses, but not registrant data. This data is required for both auditing of proper use of the system and for monitoring system operation and performance. Logs are treated as “first class” objects, that is, accessible via formal access methods. In anticipation that some queries will be deemed sensitive and hence must not be made available to those who usually have access to logs, the system MUST provide segregated, controlled storage and access to such data.
14. If adopted, the technical model will require more work to become an implementable specification.

4. Use Cases

In its early conversations, the TSG developed a number of use cases which it used to help frame its discussions about the functional requirements of the system. These use cases were then categorized as being Critical (must haves), Important (nice to haves), and Useful (but not necessary).

The following table lists these use cases, and the category into which they were placed.

Use Cases	Critical (Must have)	Important (Nice to have)	Useful (But not necessary)
Use Case #1: Authorized users require access to domain records, which might include single queries or multiple queries.	x		
Use Case #2: User receives authorization online and gets domain name registration data immediately. Authorization can be broad and ongoing, or specific and constrained.	x		
Use Case #3: Unauthorized, unauthenticated users request access to data elements associated with domain records. The system returns a well-formed (but negative) RDAP response.	x		
Use Case #4: Authenticated user requests data for which user is not authorized. The system returns a well-formed (but negative) RDAP response.	x		
Use Case #5: Data subject requests their own data via this system.			x

5. System Requirements

The following system requirements dictate function and features of the Technical Model. These requirements were used to evaluate technical solutions. Therefore, some of the requirements are mandates on the capabilities of the chosen technologies but are not mandates on their usage in operational practice; they are merely options to be utilized as policy dictates.

1. Overall
    a. The technologies used to implement Requestor identification, Authentication and Authorization MUST be based on current Internet standards.
    b. Requestors MUST be able to discover the base URL for the centralized access and authorization system. For example, CP RDAP Servers could include a URL for the ICANN Access Service in responses to unauthenticated queries, so that users may discover the means to use the service.
    c. The system's components will run on both IPv4 and IPv6 transport.
    d. The system MUST support a distributed data model, where registration data is stored by the CPs and non-public gTLD domain name registration data is only transferred through ICANN.
    e. All usage of RDAP and any other associated systems MUST use Transport Layer Security (TLS) for HTTP (HTTPS) following Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) (RFC 7525), DNS Security Introduction and Requirements (RFC 4033), and other applicable security protocols.
    f. The needs of Requestors SHOULD be assessed while implementing a system for access to non-public gTLD domain name registration data.
    g. The system MUST be able to determine whether a Requestor is authorized for access to non-public gTLD domain name registration data.
    h. The system MUST be able to associate attributes to the Requestor, and these attributes MUST be passed by the Requestor to the ICANN RDAP Gateway.
2. ICANN Browser-based Web Portal
    a. The system MUST provide a Web-based interface for "exceptional" requests (requests not pre-authorized) which will be submitted – and reviewed by – a human. Once authorized, data may be provided either via an ICANN browser-based RDAP client or via a user's own client and tooling.
    b. The system MUST allow triage of requests to identify high-priority requests that are to be handled first.
    c. The system MUST provide notifications of the progress of a request through the triage-review-fulfilment process, so Requestors are promptly notified of the result of their request.
    d. The system MUST assign each Requestor a unique identifier and assign each request a unique identifier.
3. ICANN Browser-based RDAP Client

    a. This client MUST be implemented entirely in-browser as a “single page application” and MUST NOT process credential information, queries, or results for any duration longer than is necessary to service a single RDAP query.
    b. Use of the client is optional (i.e. users may use their own RDAP clients and tooling if they desire).

4. Authentication and Authorization Determination
    a. Authentication and authorization determination MAY be delegated to agents that are qualified and appointed by the coordinating party.
5. ICANN RDAP Gateway
    a. The system MUST be able to process both unauthenticated and authenticated requests, and MUST redirect unauthenticated and unauthorized requests to the appropriate RDAP server able to answer for public domain name registration data.
    b. The system MUST be able to support multiple authenticated Requestor identities, each of which MAY be assigned a role.
    c. The system MUST be able to support multiple authorization policies based on the role assigned to the Requestor, and on the query.
    d. The system MUST be able to allow granular access to various data elements in RDAP based on authorization policies.
    e. The system MUST support passing Requestor *attributes* (see 1.h) to the authoritative CP RDAP servers. Whether the system passes attributes is dictated by policy.
    f. The system MUST support passing the Requestor and/or request *identifier* (see 2.d) to the authoritative CP RDAP servers. Whether the system passes the identifier is dictated by policy.
    g. The system MUST be able to receive and redirect queries from Requestors who are not authorized for access to non-public gTLD domain name registration data.
    h. The system MUST NOT prohibit the ability of non-interactive clients to issue requests (i.e. there MUST be no requirement requiring user interaction such as might be necessary with a browser-based RDAP client).
6. CP RDAP Servers
    a. If a CP RDAP Server can respond to a query with a result it is able to fulfill, it MUST receive and respond to queries from ICANN RDAP Gateway with all available domain name registration data.
    b. If a CP RDAP Server cannot respond to a query with a result that it would otherwise be capable of fulfilling, it MUST return an appropriate HTTP error code along with a descriptive RDAP error payload regarding the nature of the rejection.
    c. CP RDAP Servers MUST receive and respond to queries from unauthenticated requestors with all available public domain name registration data.
7. Logging / Auditing
    a. Logging and audit data held by all parties MUST be stored securely to prevent unauthorized disclosure of requests. Query logs contain sensitive information and if not treated carefully, could disclose private information inadvertently.

   b. There MUST be an ability to attribute each query to the Requestor issuing the query. This attribution MUST distinguish each query from every other query so that each user-to-query pairing will be unique and independently verifiable.
   c. The ICANN RDAP Gateway MUST log each query.
   d. Every Identity Provider MUST have the ability to download a query log containing only the queries of the users of said Identity Provider. Whether this feature is actually made available is dictated by policy.
   e. There MUST be a common format for the query log.
   f. The query logs should only be available under appropriate controls, which need to be considered in the policy development process.
   g. ICANN MUST publish aggregate statistics of queries for non-public domain name registration data.
   h. Data MUST be retained in accordance with requirements specified by policy.
   i. The system MUST provide the ability to reconcile queries between ICANN, Identity Providers, Third-Party Authorizers, CPs, and Requestors.

8. Performance / Service Level Agreements (SLAs)
   a. ICANN MUST develop and publish SLA commitments for all the service subsystems' (e.g., ICANN RDAP Gateway, CP RDAP servers, Identity Providers, Authorizers) availability, and request resolution times. Speed, reliability and responsiveness are important for usability of the system, and any SLA commitments should take these factors into consideration. Rate limiting should be applied to safeguard against abuse, denial of service attacks or to preserve stability, rather than to justify underprovisioning of systems.
9. Information Security Requirements
   a. The security controls for the system are expected to be determined and maintained based on risk assessments (for example, Article 32 of the GDPR[13]).
   b. ICANN, Identity Providers, and Third Party Authorizers are expected to undergo an annual security audit by a third-party auditor and provide the audit report as requested by the interested parties.
   c. All actors in the system MUST adopt best current practices for credential management lifecycle (e.g. multi-factor authentication, hardware tokens, quarterly account reviews and so on). See Appendix 1.
   d. There MUST be a mechanism for reporting breaches of data privacy and security (for example, to be in compliance with Article 33 of the GDPR[14]). See Appendix 1.
10. Information Security Guidelines
   a. The system MUST be governed by a business continuity management program and disaster recovery/incident response plans.
   b. The system MUST be developed and operated under an appropriate systems development life cycle.

[13] http://www.privacy-regulation.eu/en/article-32-security-of-processing-GDPR.htm

[14] https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1532348683434&uri=CELEX:02016R0679-20160504

c. Cryptographic techniques such as encryption and signing are expected to be adopted across the infrastructure to protect the confidentiality and integrity of data at rest and data in transit.

6. Functional Requirements

The System Requirements described in Section 5 include a mix of functional requirements, operational requirements, and management requirements. The set of functional requirements can be used to describe specific features that are mandatory for inclusion in the Technical Model. In this section, requirements that the TSG determined to be functional requirements are used to assess two client authentication technologies.

6.1 Functional Requirements Mapping and Analysis

The TSG considered two technologies as candidates for inclusion in its Technical Model proposal. The first candidate is authentication and authorization using OpenID Connect and OAuth 2.0. This technology is similar to the “Single Sign On” (SSO) technology that is commonly used to identify, authenticate, and authorize an end user for access to an online resource based on a credential (commonly a shared secret, but the credential may also be a digital certificate) issued to the end user by an Identity Provider. The second candidate is mutual Transport Layer Security (TLS) authentication using a digital certificate issued to the end user by a Certification Authority (CA). This technology is commonly used to identify and authenticate web servers, but TLS includes the ability for a web service to request a certificate from a client that can be authenticated prior to granting creation of a connection between a client application and the web service.

The TSG believes that OpenID Connect/OAuth authentication meets all of the identified functional requirements. We identified two requirements that were problematic for mutual TLS authentication:

> 1.h.: With digital certificates, attributes that can be used to identify and authorize an end-user are encoded when a certificate is created by the CA. They persist for the duration of the certificate validity period. We believe this requirement can be met, but there may be more of an administrative and operational burden due to the need to reissue and reinstall a client certificate if the attributes need to be adjusted on a per-query basis. We also felt that the overhead required to request, create, and install a client certificate may impose an operational burden for an end-user who needs to perform a one-time query. The relatively long-term validity period associated with a digital certificate would require periodic reviews of end-user eligibility to be associated with those attributes. For example, it would be necessary to periodically review the role assigned to an end-user to determine if the end-user remains eligible to assume that role.
>
> 4.a.: Digital certificates can be encoded with information that can be used to make end-user authorization decisions, but the CA that issues a certificate plays no role in the authentication and authorization transaction that takes place when a TLS connection is established beyond optionally determining if the certificate has been revoked. It is not currently possible to transmit the certificate to a third-party for authorization determination when a TLS connection is being established.

7. Out of Scope Requirements

The TSG identified several functional requirements that it decided were out of scope. These requirements are listed below, and the rationale for placing them out of scope is outlined.

7.1. Reverse Search

In the context of domain name registration data, “reverse search” refers to the ability to identify domains based on common attributes such as a nameserver IP address or registrant email address. This is in contrast to the traditional model of accessing registration data by means of a direct lookup of a resource by an already-known identifier (e.g., a domain name). A reverse search query may be sent to a specific CP’s RDAP server, but may also be sent to an RDAP gateway service, which “fans out” the same query to multiple CPs’ RDAP servers, and aggregates the responses into a single result.

While RDAP already supports limited search capabilities, and an Internet-Draft extending these capabilities exists, the TSG does not believe that consideration of a reverse search system - and especially a cross-TLD reverse search system - is within its remit. The technical feasibility of such a system (in which a single client query could result in thousands of search queries to CPs’ RDAP servers), and the policy developments required to allow such a system to exist, are too great for it to be considered at this point in time.

Should reverse search become feasible (from a technical and/or policy perspective) in the future, the TSG believes that the solution proposed in this document provides a platform in which it could be implemented.

7.2. Pseudonymity of Registrants

As an alternative to reverse search, some have proposed that CPs add pseudonymous, hash-based identifiers to RDDS records, which would allow third parties to correlate domains with common registrants, without disclosing the personal information from which the identifiers are derived.

The TSG believes that this proposal is out of scope, since it proposes a change to the basic data elements collected and processed by CPs: such a proposal is best addressed in the policy domain first, with a technical implementation to follow afterwards.

As with reverse search, the TSG believes that registrant pseudonymity could be supported by the Technical Model in the future.

7.3. Bulk Queries and Bulk Access

The TSG defines “bulk queries” and “bulk access” as follows:

- “Bulk queries” is defined as a single transaction, containing multiple discrete queries, which produces multiple responses.
- “Bulk access” is defined as requestors gaining access to entire datasets in an encapsulated format such as a CSV, XML or relational database file.

The TSG believes that neither feature is within scope since neither feature is currently possible within RDAP.

8. Actor Models

To implement access to non-public gTLD domain name registration data, several organizational entities, or actors, have been proposed, the combinations of which constitute several actor models.

The following is a list of these organization entity actors:

1. Requestors - The entities submitting queries, the results of which gain them access to non-public gTLD registration data.
2. ICANN RDAP Gateway - A central RDAP proxy server through which all queries are directed and all responses are filtered.
3. Identity Providers - Organizations assigning credentials to and authenticating requestors using functions provided by OpenID Connect and OAuth 2.0.
4. Third-Party Authorizers - Organizations determining the types of data to be accessed by authenticated requestors.
5. CP RDAP Servers - RDAP servers operated by gTLD domain registries and registrars.

Mapping these organizational entities to the actors (or participants) in a technical interaction using OAuth/OpenID Connect yields the following:

1. Requestors - The individuals submitting queries.
2. Browser User Agent - A web browser used by a requestor to obtain an Access Token.
3. RDAP User Agent - An RDAP client which uses an Access Token obtained by a requestor to conduct RDAP queries (in some cases, this user agent may be an application in a web browser and indistinguishable from the Browser User Agent).
4. ICANN Access Service - A browser-based, web service used by the requestors to obtain an Access Token from the OAuth/OpenID Connect process. In OAuth/OpenID Connect terms, this would be the Relying Party.
5. ICANN RDAP Gateway - An RDAP server proxy evaluating access based on an Access Token to which all queries are submitted and through which all responses are filtered. In OAuth/OpenID Connect terms, this would be the Resource Server.
6. Identity Providers - Organizations authenticating requestors.
7. Third Party Authorizers - Organizations determining the type of data to be accessed by authenticated requestors.
8. CP Servers - RDAP servers operated by domain registries and registrars.

The interactions between these actors are illustrated in Figure 2 below.

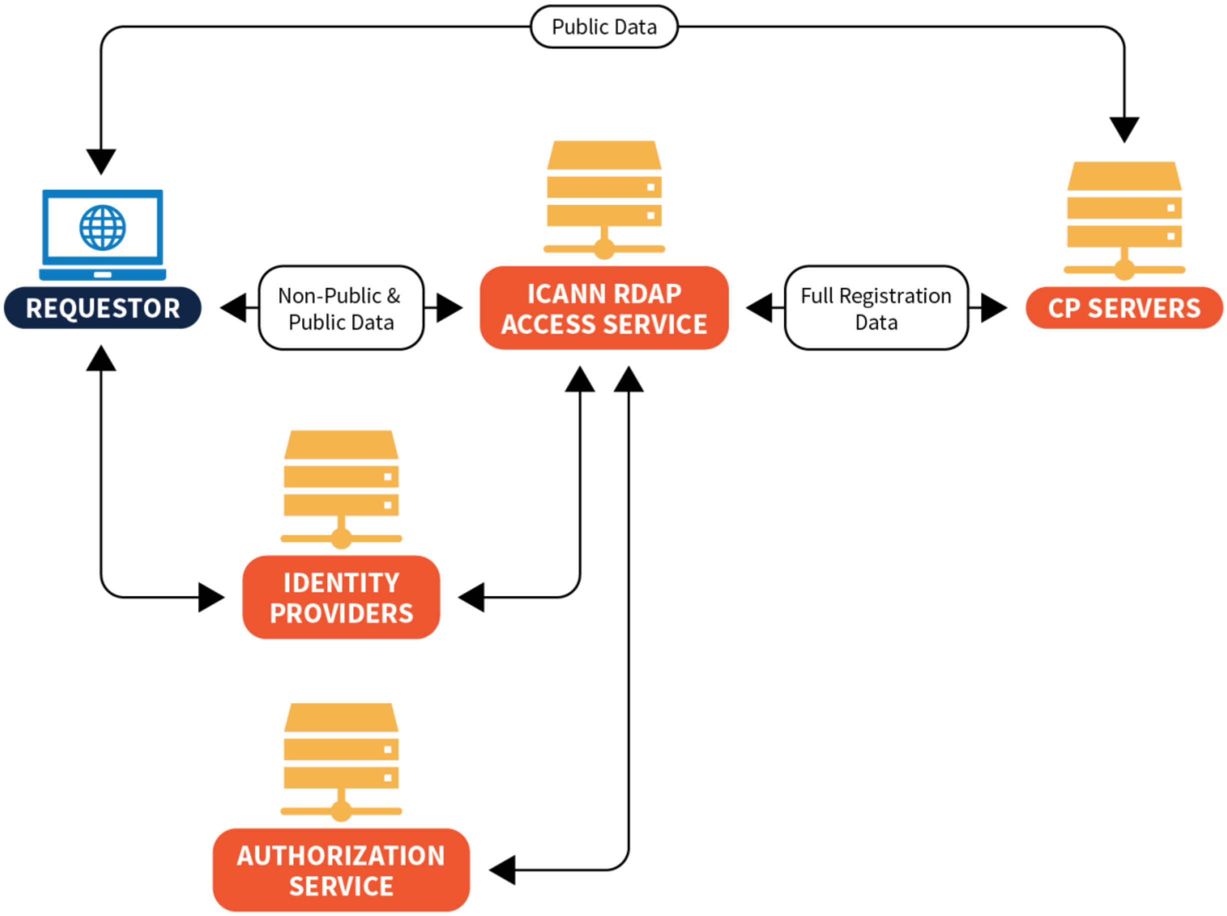

Figure 2.

Each of the following actor models are supported by the proposed solution in Section 8. Policy requirements will determine which actor model is best suited.

<u>Note</u>: The TSG takes no position on which of these models is ideal for selection to be deployed but is merely presenting possible combinations of actors and responsibilities that meet the stated requirements.

Actor Model 1: ICANN as Gateway and Sole Identity Provider and Authorizer

The simplest actor model is one in which the coordinating party takes on responsibility for identity management and authorization. From a technical point of view, this model offers the least number of interactions.

However this model requires the coordinating party to have knowledge that enables them to identify and authenticate each entity that is requesting access to the system. For example, a request to authenticate that an individual is a member of law enforcement would require knowledge of specific government bodies.

In regards to authorization, this model does not suffer from potential inconsistencies since only one party is responsible to implement the policy that dictates who gets access to what under what circumstances.

From a more practical perspective, this model would likely encumber ICANN with a burdensome, and likely politically unpalatable, need to vet and credential all requestors.

Actor Model 2: ICANN Gateway Using Multiple Identity Providers with ICANN as Sole Authorizer

In this model, ICANN delegates identity management to third-party Identity Providers, including for example national or regional law enforcement bodies and civil legal organizations, where vetting and credentialing of requestors may be already in-use, and natural.

By keeping ICANN as the sole Authorizer, this model lowers the number of interactions, at least in regards to the authorization steps. Similar to model one, it does not suffer from potential inconsistencies in implementing the authorization policy.

Actor Model 3: ICANN Gateway Using Multiple Identity Providers with Third-Party Authorizers

In this model, ICANN delegates identity management to third-party Identity Providers and authorization of data policy to third-party Authorizers.

While this model relieves ICANN of the burden of vetting requests and credentialing requestors similar to Actor Model 2, unlike the previous model, it delegates control of authorization decisions to third parties, which raises the possibility of inconsistencies in implementing the authorization policy.

Actor Model 4: ICANN as Gateway and Sole Identity Provider with Third-Party Authorizers

Like Model 1, ICANN takes on responsibility for identity management and authorization with the pros and cons described there.

By having multiple authorizers, this model raises the potential for inconsistencies in the implementation of the authorization policy as described in Model 3.

9. Implementation Considerations

While not hard requirements, there are several considerations on the implementation and ongoing operation of this system influencing the proposed solution.

Given the nature of a system of this type, there will be complexity. Therefore, burdensome complexity should be pushed, when possible, to the fewest and most capable actors.

The largest contingent of actors in this system will be the requestors, for example law enforcement agents. Any proposed solution should attempt to keep burdensome, complex and technical matters from impacting their primary duties.

Tooling, such as open-source RDAP user agents, may be expected to grow over time even as other parts of the system remain static. Any proposed solution should attempt to lower the implementation threshold necessary for the creation of tooling, which should also impact the complexity upon requestors.

A pure mutual TLS authentication system has many advantages with respect to simplicity. However, such a solution does not support a multiple authorizer model and places a significant burden upon Identity Providers (in the form of running a Certificate Authority) and Requestors (in that generation of cryptographic key pairs and installation of certificates may not be allowed by internal policy in many organizational information technology environments).

Likewise, the device flow of OAuth/OpenID Connect requires substantial revision and complication in RDAP user agents.

Therefore, the proposed solution is to use OAuth/OpenID Connect with a browser-based RDAP Access Service to obtain an Access Token to be used by RDAP user agents to access non-public gTLD domain name registration data using RDAP and well-known HTTP bearer token methods.

10. Proposed Solution

This proposed solution will accommodate any one of the four actor models in Section 8. In this section each component is described separately.

10.1 System Components

This document proposes two parallel systems for processing requests for non-public gTLD domain name registration data: a browser-based Web portal, which allows asynchronous, manual submission, review, authorization and (optionally) completion of requests for non-public registration data through an ICANN browser-based RDAP client, and an RDAP Gateway, which allows synchronous, automated machine-to-machine requests for non-public registration data. Either or both of these systems may be deployed, according to policy development outcomes.

10.2 Web Portal Description

The Web Portal is intended to be used by human beings to submit requests for non-public registration data which are then manually reviewed. As a result, the user experience of this system would be substantially similar to that of a web-based support or helpdesk system. It may be the case that such a system could be deployed without significant software development.

The Web Portal would still use OpenID Connect to authenticate Requestors who, once authenticated, would complete a form that gathers the required information needed to process their request.

An internal administrative interface would allow an individual to review requests, request further clarifying information from Requestors, and ultimately approve or reject requests. Approved requests could then be completed by back-end systems which would obtain the required data from the CP RDAP Servers and add the information to the case in the system.

Requestors would receive updates on the progress of their requests by email. The system SHOULD NOT transmit non-public gTLD domain name registration data or other sensitive data by email, but instead prompt the recipient to log in to the portal in order to view the contents of the update.

On approval, an Access Token and an Identity Token are attached to the request. These tokens will have a limited lifetime, and will be constructed to grant access to the Requestor for the specific data sought. The tokens can then be used with the Web RDAP client (see below) to retrieve the requested data.

10.3 Web RDAP Client Description

The Web RDAP Client is an ICANN-operated RDAP client providing a web-based user interface to access RDAP information. The client is a convenience for users who may not have or may not need RDAP tooling. This web-based RDAP client will be loosely coupled to the ICANN web portal used for applying for authorization to access non-public domain name registration data.

As an ICANN web system, it is important that this RDAP client not keep, cache, or persist any RDAP results beyond that which is necessary for servicing the specific query given to the RDAP client.

10.4 RDAP Gateway Description

The authentication mechanism used between the client and the ICANN RDAP Gateway will be based on OpenID Connect and OAuth 2.0 using shared secrets (e.g., usernames and passwords). Other authentication methods may be added once approved by ICANN. As other authentication methods become standardized, they may be considered for adoption. OpenID Connect and OAuth 2.0 are the recommended mechanism because it meets all of the identified functional requirements.

Mutual TLS authentication will be used to secure RDAP communications between ICANN and the CPs, and preferably between the ICANN Access Service, the Identity Provider(s) and the Third Party Authorizer(s). This method is recommended because ICANN will be fully authorized to access non-public domain name registration data, and only needs to authenticate itself without CPs having to make detailed authorization decisions on a per-query basis. The functional requirements not met by this method do not apply to interactions between ICANN and the CPs.

10.5 Prerequisites

Identity Providers may be appointed and approved to perform client identification and authentication functions. Third-party Authorizers may be appointed to perform authorization and information association functions. ICANN may serve as an Identity Provider, an Authorizer, or both. The functions (i.e. Identity Provider or Authorizer) can also be delegated to duly appointed, independent third party operators who are affiliated with Requestor communities of interest.

Identity Providers, Third-Party Authorizers, and ICANN exchange or publish configuration information to identify service endpoints. Service endpoints can be discovered dynamically or exchanged statically as a matter of implementation policy; this information is needed to facilitate web service interactions between these actors.

Requestors will register with and obtain credentials from an Identity Provider.

Identity Providers will assign attributes to Requestors. These attributes are associated with, for example, their functional role, the purpose of their RDAP queries, and any other information that is required by policy to make data access decisions when an RDAP query is processed. (As described in Section 5 of the OpenID Connect Core specification, these attributes are known as “claims” that are encoded in a data structure known as an “ID token”.) Policies must be developed to determine the set of attributes or claims that are needed to make data access decisions. Policies must also be developed to determine the attribute or claim sets that can be managed by specific Identity Providers. For example, claims associated with a law enforcement role should be limited to Identity Providers who are responsible for providing services to law enforcement agencies.

Some of the functionality described in this document, such as support for OpenID Connect and OAuth 2.0 and enhanced search capabilities, is documented in proposed extensions to the RDAP protocol. These extensions will need to be developed further by the Internet Engineering Task Force (IETF) before they can reliably be implemented.

10.6 Processing Steps

A requestor who wishes to submit an RDAP query first submits an Access Request (as described below). An Access Request is followed by processing to identify and authenticate the Requestor. A Requestor who has been identified and authenticated may then request tokens that can be used to submit an RDAP query. The RDAP query and tokens are submitted for processing, and an appropriate RDAP response is returned. The error-free flow of information associated with each of these steps is described in more detail below.

1. Access Request

The Requestor who wishes to perform an RDAP query uses an RDAP User Agent to send an HTTP Access Request to the Access Service. The Access Service will be operated by ICANN. The Access Service receives the request and returns an HTTP redirect to the client that prompts the client to send an Authentication Request to an Authorization Endpoint operated by an Identity Provider.

2. Identification and Authentication

The Identity Provider that operates the Authorization Endpoint prompts the client for the Requestor's credentials. The Requestor provides the credentials, and the Identity Provider attempts to authenticate the Requestor. The authenticated Requestor selects the attributes to be associated with the identity they are using to perform their RDAP query and submits their consent to share this information with the ICANN RDAP Gateway to the Identity Provider. The Identity Provider responds to their consent submission by returning an Authorization Code and an HTTP redirect (to the Access Service) to the client. The client validates the Authorization Code as described in Section 3.3.2.10 of the OpenID Connect specification and begins RDAP query processing.

3. Setup for RDAP Query

The client submits the received authorization code to the Access Service by following the redirect received from the Identity Provider. The Access Service receives the request from the client and submits a request for an ID token and an Access Token to a Token Endpoint operated by the Identity Provider. The Token Endpoint validates the Authorization code and returns an ID Token and an Access Token to the Access Service, which in turn returns them to the client application.

The client prepares an RDAP query. The RDAP query, an ID token, and an OPTIONAL Access Token are sent to the ICANN RDAP Gateway.

4. RDAP Query Processing

The ICANN RDAP Gateway receives the RDAP query, an ID token, and an OPTIONAL Access Token. The ICANN RDAP Gateway sends this information to a Third Party Authorizer (this service can be operated by ICANN using, for example, OAuth claims or the W3C Verifiable Credentials Data Model) for verification and validation. The tokens are validated as described in Sections 3.1.3.7 and 3.1.3.8 of the OpenID Connect specification, and the identity attributes (known as "claims" in OAuth 2.0) are retrieved from the ID token. The Third Party Authorizer maps the set of claims to a set of policies to determine if the requestor is authorized for access to any non-public gTLD domain name registration data elements. The Third Party Authorizer sends a response to the ICANN RDAP Gateway that indicates[15] the result of authorization processing. If the Requestor is authorized, the ICANN RDAP Gateway sends RDAP queries to the specific CP RDAP servers that are authoritative (e.g., have the closest relationship to the data subject) for the individual data elements within the requested data. These queries from the ICANN RDAP Gateway to the CP Servers may contain secure metadata as specified by the system requirements and relevant policy. The CP RDAP Servers each return RDAP responses containing the full set of data elements for which they are authoritative. The ICANN RDAP Gateway receives, processes, and filters the data to form a complete RDAP response that contains non-public data in accordance with the Requestor's level of access. The ICANN RDAP Gateway returns the RDAP response to the client.

10.7 Data Flow Diagram

The processing steps described above are illustrated in Figure 3 below.

[15] The TSG-RD notes that more work will be done in this area before finalizing this document.

Data Flow Diagram

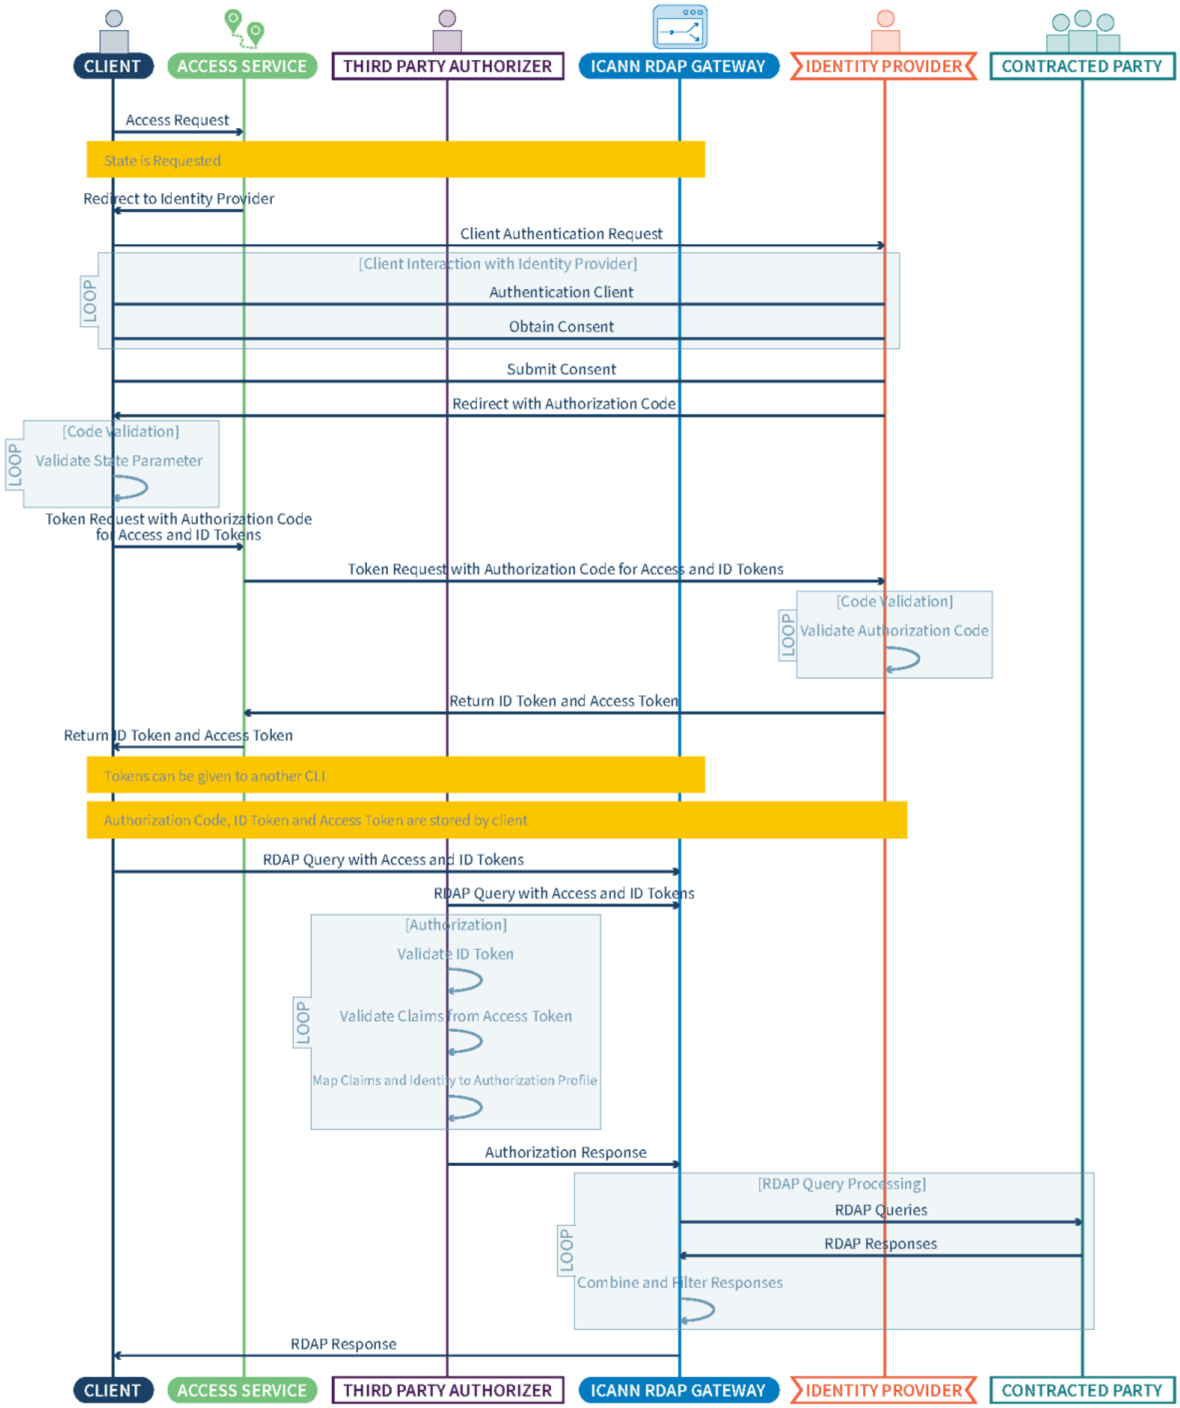
CLIENT
ACCESS SERVICE
THIRD PARTY AUTHORIZER
ICANN RDAP GATEWAY
IDENTITY PROVIDER
CONTRACTED PARTY
Access Request
State is Requested
Redirect to Identity Provider
Client Authentication Request
LOOP
[Client Interaction with Identity Provider]
Authentication Client
Obtain Consent
Submit Consent
Redirect with Authorization Code
LOOP
[Code Validation]
Validate State Parameter
Token Request with Authorization Code for Access and ID Tokens
Token Request with Authorization Code for Access and ID Tokens
LOOP
[Code Validation]
Validate Authorization Code
Return ID Token and Access Token
Return ID Token and Access Token
Tokens can be given to another CLI
Authorization Code, ID Token and Access Token are stored by client
RDAP Query with Access and ID Tokens
RDAP Query with Access and ID Tokens
LOOP
[Authorization]
Validate ID Token
Validate Claims from Access Token
Map Claims and Identity to Authorization Profile
Authorization Response
LOOP
[RDAP Query Processing]
RDAP Queries
RDAP Responses
Combine and Filter Responses
RDAP Response
CLIENT
ACCESS SERVICE
THIRD PARTY AUTHORIZER
ICANN RDAP GATEWAY
IDENTITY PROVIDER
CONTRACTED PARTY

Figure 3: Data Flow Diagram

11. Considerations for ICANN Community and Organization

During its deliberations, the TSG identified a number of issues that it believes need consideration by stakeholders and the community. These are outlined below.

11.1 Data Retention

The system requires various parties (such as ICANN org, Identity Providers, and Authorizers) to collect and store data of various kinds, including user account information, transaction logs, and metadata such as date-and-time of requests. As a matter of best practice and to comply with data protection law, the TSG believes that policies regarding retention and deletion of this data, should be established, communicated to the data processors, audited, and enforced. These policies are outside of the TSG's narrow technical scope.

In contrast to the retention of logging data, the system MUST NOT retain non-public gTLD domain name registration data in any persistent way or for any longer than needed to fulfill the query for which the data were acquired.

11.2 Service Level Agreements (SLAs)

In order to ensure a reliable system, Service Level Agreements (SLAs) MUST be established for each of the parties that operate the elements of the system. These SLAs would define the service performance levels expected of each party and the penalties for failing to meet them.

The TSG understands that the CPs will be subject to SLAs for operating their respective RDAP services. However, the TSG believes that the other actors in the system should also be subject to corresponding SLAs to carry out their respective obligations, specifically:

- ICANN org (as the operator of the RDAP Gateway, ICANN browser-based Web Portal and ICANN browser-based RDAP Client)
- Identity Providers (upon which both Requestors and ICANN org will rely)
- Third-Party Authorizers (also relied upon by requestors and ICANN org)

The TSG believes that defining the service level performance requirements for each party and the manner in which they are established, audited, and enforced, as well as whether SLA performance should be reported publicly is outside the scope of its remit and should be determined at the policy level.

It is recommended that ICANN org provide transparent reporting on the service level performance of each of the actors in the system, such as a "status page" or "dashboard" giving information on the status of component services. This provides its users with a clear view of any disruptions that might affect their use of the service.

11.3 Obligations on ICANN Org

The TSG recognizes that it proposes a solution that could potentially impose significant operational burdens on ICANN org, especially if the community determines that the operator of the RDAP Gateway must meet stringent service level requirements, and operate at significant scale.

It is recommended that ICANN org review the spectrum of potential operational outcomes for deployment and operation of the system proposed to determine the feasibility of such outcomes, operational and financial impact, and how challenges might be addressed.

It is recommended that ICANN org publish its review for Public Comment and that it solicit feedback from technical experts on its feasibility.

11.4 Risk to ICANN Org

The TSG notes that ICANN org will function as the coordinating party of the system in the gTLD space and, depending on the policy development outcome, may result in ICANN org shouldering the burden of vetting and credentialing requestors. This may expose ICANN org to significant operational and legal risks. It is recommended that ICANN org identify, assess, and where possible take steps to mitigate these risks.

11.5 Risks to Contracted Parties

The TSG was established to determine the feasibility of a system that would mitigate some or all of the legal risks to CPs from the disclosure of non-public registration data. The TSG cannot comment on the validity of this assumption, and expects that the CPs will come to their own determination based on their own legal advice.

11.6 Notification of Data Subjects by Contracted Parties

The TSG acknowledges that CPs may find it necessary to notify data subjects when their personal information is disclosed to third parties through this system, and recognizes the tension that such notification may occasionally be prohibited by law; for example, when a disclosure is made in pursuit of a criminal investigation. Nothing in the system described in this document prohibits such disclosure. The system provides a mechanism by which details of the Requestor as well as the nature and purpose of their request MAY (according to policy configuration) be communicated to CPs for this purpose.

11.7 Transparency

The Group believes that openness and transparency will be vital to ensuring the acceptance of the proposed model by the wider stakeholder community. Therefore, the TSG recommends that ICANN consider publishing a regular transparency report[16] similar to those published by other organizations providing statistics on requests for access to non-public gTLD domain name registration data.

11.8 Mechanism for Handling Complaints

The TSG believes that users of the system who are unsatisfied with the outcome of their requests (for example, because their request has been denied, or because they believe their request was not fully satisfied) should have a means to escalate these requests through a complaints process. Complaints relating to requests that have been triaged as high priority should also be treated as a high priority.

It is likely that ICANN org (and other actors within the system) may receive requests to delete personal data under Article 17 of the GDPR. It is recommended that ICANN org establish a process for handling such requests, which may involve directing the submitter to the appropriate CP.

[16] https://en.wikipedia.org/wiki/Transparency_report

Data subjects may also wish to submit complaints about any disclosures of their personal information which they believe to be inappropriate or contrary to the established policy. ICANN org should establish a system for receiving and reviewing such complaints.

11.9 Confidentiality of Requestors' Requests

The system proposed by the TSG allows - but does not require - metadata about a request for non-public gTLD domain name registration data to be disclosed to CPs when the request is transmitted to their CP servers. The TSG recommends that an Acceptable Use Policy be crafted that addresses the potential risks to both Requestors and data subjects associated with such disclosures.

11.10 Privacy-by-Design Considerations

The TSG reviewed the Technical Model with respect to the 7 Foundational Principles of Privacy by Design.[17] Presented below are its observations with regard to the conformance of the Technical Model to these principles and the TSG's comments for areas where we believe some exceptions may exist.

Principle	Areas of conformance	Comments
1. Proactive not Reactive; Preventative not Remedial	Only legitimate legal queries for non-public gTLD domain name registration data will be allowed to be fulfilled within the ICANN system.	The WHOIS system and its predecessors predate the work of the TSG. Thus, most of the TSG work is necessarily remedial.
2. Privacy as the Default	Non-public gTLD domain name registration data redacted in CP RDAP services, only accessed via ICANN system.	The implementation of the ICANN RDAP Gateway must "fail closed", e.g. should not disclose non-public data unless all input/authorization validation tests pass successfully.
3. Privacy Embedded into Design	Use of OAuth and OpenID Connect (see Section 10). Data minimization by limiting storage of non-public gTLD domain name registration data to CPs.	The WHOIS system and its predecessors pre-date the work of the TSG. Thus, most of the TSG work is modifying existing design.
4. Full Functionality – Positive-Sum, not Zero-Sum	Non-public gTLD domain name registration data still available for legitimate purposes, governed by privacy and security controls (Section 5, 7-10)	
5. End-to-End Security –	The model requires the use of	The model requires secure

[17] https://iapp.org/media/pdf/resource_center/Privacy%20by%20Design%20-%207%20Foundational%20Principles.pdf

Lifecycle Protection	TLS for communication between the subsystems, which provides both confidentiality and integrity. Non-public gTLD domain name registration data is not stored by the system. Logging data will not contain non-public gTLD domain name registration data nor identifying data of the Requestor and is expected to be encrypted at rest.	transmission of data between the components and actors within the system. Data handling processes outside the boundary of the system (such as within the CP's system and at the Requestor) cannot be audited or enforced by technical means, but through contractual provisions.
6. Visibility and Transparency	TSG recommends publication of a transparency report (see Section 11.7). CPs can notify data subjects if or when their data is disclosed.	
7. Respect for User (data subject) Privacy	The data subject's influence on how the system processes their data is limited. The TSG's focus has been on transparency as well as allowing CPs to respect the privacy of data subjects.	Legal or other considerations would require data providers to allow access to user data for authorized requestors. In some cases, it may not be feasible to provide appropriate notice.

12. Conclusion

The Temporary Specification for gTLD Registration Data[18] established a restricted environment in which there is no uniform way to obtain non-public gTLD domain name registration data. This was precipitated by the strong privacy requirements of GDPR and other such regulations, which exist in tension with the operational need to use registration data for legitimate purposes.

The TSG was charged with the task of developing a technical solution that strikes a balance between these opposing needs while still observing all of the requirements expected from a service that must be cautious about serving all of its constituents. With this document, the TSG has delivered an outline of its working assumptions, requirements, and its proposed solution. The TSG now submits this document to the ICANN President and CEO for further consideration and appropriate action.

To the maximum extent practical, the TSG has endeavored to avoid either influencing policy decisions or being influenced by them by way of open implementation decisions.

The TSG has proposed a technical solution that it believes:

- Accommodates all of the actor models described in Section 8.

[18] https://www.icann.org/resources/pages/gtld-registration-data-specs-en/

- Allows parties claiming legal legitimate purposes for accessing non-public gTLD registration data to get access to it in a uniform way.
- Provides sufficient logging as to enable auditing.
- Ensures integrity of the data delivered.
- Enables trust in the proposed system by way of regular transparency and performance reports.
- Requires adherence to established, deployed standards, allowing for a design that supports wide interoperability.
- Exhibits a relatively simple design that enables high availability, redundancy, and scalability.
- Further assures public trust by identifying procedures for handling deviations from policy or regulation.

12.1 Future Work

The TSG's Technical Model (TSG01) should be considered a proposal. The Technical Model (TSG01) should be considered the start of some work, rather than the end and is not sufficient to be directly applicable for implementation. Future technical work is necessary to complete the technical and specification process.

What remains after collecting and reviewing feedback is to embark upon the process of developing a detailed technical description of a working model for an RDAP system that meets these needs. The description must be as simple as possible to construct, deploy, operate, monitor, audit, and scale as demand on the system grows. It should include not only those components operated by ICANN directly, but by all participants in the service. This is not a task for the TSG. We believe that the work here might likely serve as a strong foundation to develop the detailed technical descriptions that could underpin the future of domain name registration data services.

The TSG believes that any future Technical Study Group on access to non-public domain name registration data should not be convened without:

1. Clarity on the legal and liability issues related to CPs participating in the system.
2. Policymakers responding to and providing answers to policy questions that abound.
3. Validating policy answers against technical requirements mapped to the TSG's Technical Model.
4. Protocol and technical work starting on important areas like, for example, log formats, authorization schema and OAuth claims, registration of JSON values for standardized values and referrals in IANA RDAP protocol registries, data minimization through RDAP partial responses, etc.
5. A plan for measuring user satisfaction, including a method to measure user behavior and use in order to inform future design or architectural changes.

Throughout this document, a number of policies were identified that will need development in order for future design or implementation work to be possible. These are summarized as policies which:

1. Result from the EPDP, or other policy initiatives, regarding access to non-public gTLD domain name registration data.
2. Identify and select Identity Providers (if that choice is made) that can grant credentials for use in the system.

3. Describe the general qualifications of a Requestor that is authorized to access non-public gTLD domain name registration data, such as which sorts of Requestors get access to which fields of non-public gTLD domain name registration data ("the authorization policy").
4. Detail whether a particular category of Requestors or Requestors in general, can download logs of their activity.
5. Describe data retention requirements imposed on each component of the system.
6. Describe service Level Requirements (SLRs) for each component of the system, including whether those SLRs and evaluations of component operators against them are made public, and for handling complaints about access.
7. Specify legitimate causes for denying a request.
8. Outline support for correlation via a pseudonymity query as described in Section 7.2.
9. Outline the selection of an actor model as described in Section 8 and the appropriate supported components and service discovery as described in Sections 10.1 through 10.5.
10. Describe the conditions, if any, under which requests would be disclosed to CPs.
11. Provide legal analysis regarding liability of the operators of various components of the system.
12. Outline a procedure for fielding complaints about inappropriate disclosures and, accordingly, an Acceptable Use Policy.

References

Federated Authentication for the Registration Data Access Protocol (RDAP) using OpenID Connect
https://datatracker.ietf.org/doc/draft-ietf-regext-rdap-openid/

Finding the Authoritative Registration Data (RDAP) Service
https://datatracker.ietf.org/doc/rfc7484/

HTTP Usage in the Registration Data Access Protocol (RDAP)
https://datatracker.ietf.org/doc/rfc7480/

JSON Responses for the Registration Data Access Protocol (RDAP)
https://datatracker.ietf.org/doc/rfc7483/

OAuth 2.0
https://oauth.net/2/

OAuth 2.0 Mutual TLS Client Authentication and Certificate-Bound Access Tokens
https://datatracker.ietf.org/doc/draft-ietf-oauth-mtls/

OpenID Connect
https://openid.net/connect/

OpenID Connect Core
https://openid.net/specs/openid-connect-core-1_0.html

Privacy by Design: The 7 Foundational Principles, Implementation and Mapping of Fair Information Practices
https://iapp.org/media/pdf/resource_center/Privacy%20by%20Design%20-%207%20Foundational%20Principles.pdf

Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)
https://tools.ietf.org/html/rfc7525

Registration Data Access Protocol (RDAP) Partial Response
https://tools.ietf.org/html/draft-ietf-regext-rdap-partial-response

Registration Data Access Protocol (RDAP) Query Format
https://datatracker.ietf.org/doc/rfc7482/

Registration Data Access Protocol (RDAP) Reverse search capabilities
https://tools.ietf.org/html/draft-ietf-regext-rdap-reverse-search

SAC101v2.0 SSAC Advisory Regarding Access to Domain Name Registration Data
https://www.icann.org/en/system/files/files/sac-101-v2-en.pdf

Security Services for the Registration Data Access Protocol (RDAP)
https://datatracker.ietf.org/doc/rfc7481/

Transparency Report
https://en.wikipedia.org/wiki/Transparency_report

W3C Verifiable Credentials Data Model
https://www.w3.org/TR/verifiable-claims-data-model/

WHOIS Protocol Specification
https://datatracker.ietf.org/doc/rfc3912/

Appendix 1. Frameworks and Guidelines for Secure Deployment of RDAP

Information Security

ISO/IEC 27001:2013 Information technology -- Security techniques -- Information security management systems -- Requirements
https://www.iso.org/standard/54534.html?browse=tc

ISO/IEC 27002:2013 Information technology -- Security techniques -- Code of practice for information security controls
https://www.iso.org/standard/54533.html?browse=tc

SP 800-171 Rev. 1 Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations
https://csrc.nist.gov/publications/detail/sp/800-171/rev-1/final

SP 800-53 Rev. 4 Security and Privacy Controls for Federal Information Systems and Organizations
https://csrc.nist.gov/publications/detail/sp/800-53/rev-4/final

Risk Management

ISO 31000:2018 Risk management -- Guidelines
https://www.iso.org/standard/65694.html

SP 800-30 Rev. 1 Guide for Conducting Risk Assessments
https://csrc.nist.gov/publications/detail/sp/800-30/rev-1/final

Business continuity

ISO 22301:2012 Societal security -- Business continuity management systems -- Requirements
https://www.iso.org/standard/50038.html

SP 800-34 Rev. 1 Contingency Planning Guide for Federal Information Systems
https://csrc.nist.gov/publications/detail/sp/800-34/rev-1/final

Incident Response

ISO/IEC 27035-1:2016 Information technology -- Security techniques -- Information security incident management -- Part 1: Principles of incident management
https://www.iso.org/standard/60803.html

ISO/IEC 27035-2:2016 Information technology -- Security techniques -- Information security incident management -- Part 2: Guidelines to plan and prepare for incident response
https://www.iso.org/standard/62071.html

ISO/IEC CD 27035-3 Information technology -- Security techniques -- Information security incident management -- Part 3: Guidelines for incident response operations
https://www.iso.org/standard/74033.html

SP 800-61 Rev. 2 Computer Security Incident Handling Guide
https://csrc.nist.gov/publications/detail/sp/800-61/rev-2/final

Credential Management

SP 800-63-3 Digital Identity Guidelines
https://csrc.nist.gov/publications/detail/sp/800-63/3/final

SP 800-63A Digital Identity Guidelines: Enrollment and Identity Proofing
https://csrc.nist.gov/publications/detail/sp/800-63a/final

SP 800-63B Digital Identity Guidelines: Authentication and Lifecycle Management
https://csrc.nist.gov/publications/detail/sp/800-63b/final

SP 800-63C Digital Identity Guidelines: Federation and Assertions
https://csrc.nist.gov/publications/detail/sp/800-63c/final

SAC 074 | SSAC Advisory on Registrant Protection: Best Practices for Preserving Security and Stability in the Credential Management Lifecycle
https://www.icann.org/resources/files/1194801-2015-11-03-en

ISO 21188:2018 Public key infrastructure for financial services -- Practices and policy framework
https://www.iso.org/standard/63134.html

Appendix 2. Team Composition

The TSG is composed of invited members with technical backgrounds, including expertise in RDAP and authentication/authorization technologies.

Role	Name	Affiliation/Employer
Sponsor	Göran Marby	ICANN Org
Coordinator	Ram Mohan	Afilias
Team Members	Benedict Addis Gavin Brown Jorge Cano Steve Crocker Scott Hollenbeck Jody Kolker Murray Kucherawy Andy Newton Tomofumi Okubo	Registrar of Last Resort CentralNic NIC Mexico Shinkuro Verisign GoDaddy Facebook ARIN DigiCert
ICANN Org Support Team	Eleeza Agopian Francisco Arias Amy Bivins John Crain Yvette Guigneaux Daniel Halloran Gustavo Lozano Diana Middleton Erika Randall	ICANN Org

Appendix 3. Frequently Asked Questions

Throughout the development of this Technical Model, questions have arisen by members of the ICANN community and others regarding many aspects of the technical relationships of the actors of this model, the envisaged operating practice, policy implications, and other areas of concern. The purpose of this appendix is to list many of the questions and provide answers.

A3.1 Does the Technical Model pseudonymize the identity of requestors? How do the Contracted Parties know who is querying them?

Pseudonymity cannot be directly addressed until further policy requirements are known as there are many solutions to this problem depending on need. The Technical Model has requirements to pass information from the requestor all the way through to the CP, and this requirement would facilitate pseudonymity if policy requires it.

A3.2 Can Requestors use Facebook or Google OAuth or Open ID credentials to obtain Non-public Data Registration Data?

The Technical Model envisages the role of Identity Providers. If ICANN policy allows the credentials of public systems (such as Facebook or Google) to be used, the TSG believes that the expectation of protecting the privacy rights of individuals may not be met.

A3.3 What are the technical considerations of using the ICANN RDAP Gateway vs queries being sent directly to the Contracted Parties?

Beyond the policy and legal considerations for using an ICANN RDAP Gateway, there are several technical benefits.

First, the authentication and authorization method used by the CPs can be simplified to be Mutual TLS. This means the CPs authenticate only ICANN and no other entity, and the CPs need only implement one authorization policy. Otherwise, technical mechanisms would be needed to synchronize all CPs with an up-to-date list of authentication credentials, and a computer language would need to be invented to describe policy along with a mechanism to synchronise all CPs with those policies.

Second, having CPs interact only with ICANN provides a convenient funnel through which compliance, auditing, reconciliation, and performance can be measured. This would enable higher transparency of the system.

A3.4 Shouldn’t Requestors always be required to use ICANN’s browser-based RDAP client?

The TSG does not require requestors to use ICANN’s browser-based RDAP client. Many requestors, (for example, law enforcement), may have special tooling to aid in their investigations and lawful requirements of prosecution. Some tooling may cryptographically sign query results to meet chain-of-custody for evidence requirements.

A3.5 Does the Technical Model support “reverse search”?

No. The TSG’s rationale for not including this feature in its model can be found in Section 7.1.

A3.6 Does this Technical Model support bulk access to Non-public Data?

No. The TSG’s rationale for not including this feature in its model can be found in Section 7.3.

A3.7 Will data subjects receive notice of their information having been queried?

Notifications to data subjects are a matter of policy and legal requirements as noted in Section 11.6.

A3.8 How many queries can a Requestor make and for how long?

The number of queries a Requestor can send, and the amount of time in which they can be sent must be determined by policy. For example, it is possible for the system to grant a Requestor using the ICANN Access Service website only one query, even if the Requestor is using their own RDAP client. It is also possible for the system to grant a Request from a law enforcement agency an unlimited number of queries, or set number of queries in a predefined time period. The TSG’s Technical Model accommodates these and other cases.