Draft Technical Model for Access to Non-Public Registration Data

Technical Study Group on Access to Non-Public Registration Data

6 March 2019

Abstract

The purpose of this document is to propose a technical solution for access to non-public domain name registration data in 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 RDAP and OAuth 2.0 / OpenID Connect.

Executive Summary

Following the adoption of the European Union's General Data Protection Regulation (GDPR), ICANN org and the ICANN community have worked to balance the law's data protection requirements with the legal 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 registration data.

In October 2018, ICANN President and CEO Göran Marby asked Ram Mohan 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 funnel for third-party queries for non-public registration data in the gTLD space. This group's work does not interfere with community efforts to develop an RDAP profile prior to deployment of the protocol, and is not an effort to replace the 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, who would provide access to non-public registration data.

Building on the technology available via the Registration Data Access Protocol (RDAP), this approach would position ICANN as the sole access point to non-public registration data. In this paper, the group recommends a technical model for authenticating, authorizing, and providing access to non-public 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 would allow users to verify 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 registry and/or registrar to provide the requested data to ICANN, which in turn would provide it to the third party.

The TSG has not made decisions or recommendations on policy questions, e.g., who gets access, to which data fields and 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. After the group receives community feedback on this model, it will make any appropriate adjustments, and will submit this paper to Göran Marby.

1. Background

This group explored technical solutions for authenticating, authorizing, and providing access to non-public gTLD 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 RDAP profile prior to deployment of the protocol, this Technical Study Group focused its efforts on developing technical solutions for providing access to non-public gTLD registration data.

In a blog 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 Technical Study Group based their discussion, was further described during a data protection/privacy update webinar held 8 October 2018. The implementation approach described during that webinar would place ICANN in the position of providing a third party access to non-public gTLD registration data. If the query is approved, in accordance with relevant policy, ICANN would ask the appropriate registry or registrar to provide the requested data to ICANN, which in turn would provide it to the third party.

By design, the group has not made decisions or recommendations on policy questions, e.g., who gets access, to which data fields and under what conditions should access be given, and what is a legitimate interest for requesting such data. The group did consider the technical impact of policy choices, for example, policy recommendations that arise from the Expedited Policy Development Process (EPDP) or other policy initiatives. Where there are multiple alternatives in either the recommended policy(ies) or the technical implementation(s), the group 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, 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 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." 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 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 now 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 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", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

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.

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

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

ICANN RDAP 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.

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.

Non-public Data *

Public Data *

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 Owner *

Resource Server *

Third Party Authorizers - Organizations determining the types of data to be accessed by authenticated requestors.

Token Endpoint - A service implemented by an Identity provider to return ID and Access tokens.

[* Terms will be defined upon the document's completion.]

3. Assumptions

The group based its work on the following assumptions:

1. RDAP will be used to access 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 registration data.
3. ICANN will be the sole party through which access to non-public registration data is obtained in the gTLD space as part of a unified access model.
4. The system must accommodate changes to data sets or data access.
5. All credentials will be protected in a reasonable way throughout their lifecycle.
6. For the purposes of access to non-public data, scope 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.
7. It is expected that RDAP services 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.
8. The solution will take into consideration the existing practices and currently deployed uses of RDAP.
9. The RDAP pilot working group will complete its RDAP profile, and such output will be ratified through the appropriate processes.
10. Data holders assume that ICANN will ensure validity of credentials.
11. Policy choices may change the technical implementation of the proposed draft technical model.
12. If adopted, the draft technical model will require more work to become an implementable specification.

4. Use Cases

Use Cases	Critical (Must have)	Important (Nice to have)	Useful (But not necessary)
Use Case #1: Authorized users (e.g., security researchers, law enforcement, registrars, registries, etc.) require access to domain records, which might include single queries or multiple queries. This does not preclude a later mechanism supporting bulk queries and replies. Reverse search capabilities are contemplated, but the TSG recognizes that this is an advanced search capability that is not fully supported at this point in time.	x		
Use Case #2: User receives authorisation online and gets 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	x		
Use Case #4: Authenticated user requests data for which user is not authorized.	x		
Use Case #5: Data subject requests their own data via this system.			x

4.1 User Journey

- User should be able to discover the base URL for the centralized access and authorization system
- Correlate based on different aspects without seeking access to the underlying data. (eg., i can't tell who registrant is, but i can tell they are the same. Was name just registered? Registered to someone I've seen before? An abusive registrant?)
- Authorization is centralized within ICANN. Access of GDPR-protected data is centralized within ICANN.
- Users who have no authentication vs. wrongly authenticated. A lightweight mechanism to redirect such users properly.

5. System Requirements

1. Overall
    a. The technologies used to implement requestor identification, authentication and authorization MUST be based on current Internet standards.
    b. The system’s components will run on both IPv4 and IPv6 transport.

   c. The system MUST support a distributed data model, where data is stored by the authoritative contracted parties and non-public data is only transferred through ICANN.
   d. All usage of RDAP and any other associated systems MUST use TLS for HTTP (HTTPS) following recommendations in RFC 7525, DNSSEC, and other appropriately secure protocols.
2. ICANN Browser-based Web Portal
   a. The system MUST be able to determine whether a requestor is authorized for access to non-public data.
   b. 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.
   c. The system MUST provide a Web-based interface for “exceptional” requests (requests not pre-authorised) which must be submitted by, and reviewed by, a human. Once authorised, data is provided via this interface rather than via RDAP.
   d. The system MUST allow triage of requests to identify high-priority requests that must be handled first.
   e. The system MUST provide notifications of the progress of a request through the triage-review-fulfilment process, so requestors are notified promptly of the result of their request.
   f. The system MUST assign each requestor with a unique identifier.
3. Authentication and Authorization Determination
   a. Authentication and authorization determination MAY be delegated to agents that are qualified and appointed by the coordinating party (e.g., ICANN for gTLDs, RIRs for IP addresses).
4. ICANN RDAP Gateway
   a. The system MUST be able to process both unauthenticated and authenticated requestors.
   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 2.b) to the authoritative contracted party RDAP servers. Whether the system passes attributes is dictated by policy.
   f. The system MUST support passing the requestor *identifier* (see 2.f) to the authoritative contracted party RDAP servers. Whether the system passes the identifier is dictated by policy.
   g. The system MAY be able to receive and redirect queries from requestors who are not authorized for access to non-public data.
   h. The system MUST enable automation of client requests.
5. Contracted Party RDAP Servers

   a. The system MUST receive and respond to queries from ICANN RDAP Gateway with all available registration data.

6. Logging / Auditing
   a. Logging and audit data held by all parties MUST be stored securely to prevent unauthorised disclosure of requests.
   b. There MUST be an ability to attribute each query with the user 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. 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 available is dictated by policy. There MUST be a common format for the query log. The query logs SHOULD NOT be publicly available. ICANN MUST publish aggregate statistics of queries for non-public data.
   d. Data MUST be retained in accordance with requirements specified by policy.
   e. The system MUST provide the ability to reconcile queries between ICANN, identity providers, third party authorizers, contracted parties, and requestors.
7. Performance / SLA
   a. There MUST be SLA commitments for all the service subsystems' (e.g., ICANN RDAP Gateway, contracted parties RDAP servers, identity providers, authorizers) availability, and request resolution times.
8. Information Security Requirements
   a. The security controls for the system SHOULD be determined and maintained based on risk assessments (for instance, Article 32 of the GDPR).
   b. ICANN, Identity Providers, and Third Party Authorizers MUST 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).
   d. There SHOULD be a mechanism for reporting breaches of data privacy and security (for instance, to be in compliance with Article 33 of the GDPR).
      *Refer to Appendix 1 for supplemental information regarding best current practices and frameworks that could be leveraged to fulfill the information security requirements.
9. 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.
   c. Cryptographic techniques such as encryption and signing SHOULD 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 4 above include a mix of functional requirements, operational requirements, and management requirements. The set of functional requirements can be used to describe specific features that must absolutely be included in a technical model. In Section 5.1, requirements that the TSG determined to be functional requirements are used to assess two client authentication technologies. The requirements are described as “NA” if they were not identified as functional requirements, “YES” if the requirement can be met by the technology, and “NO” if the requirement cannot be met by the technology.

6.1 Functional Requirements Mapping and Analysis

The TSG considered two technologies as candidates for inclusion in our 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.

Requirement	OpenID Connect/OAuth Authentication	Mutual TLS Authentication
1.a.	YES	YES
1.b.	YES	YES
1.c.	YES	YES
1.d.	YES	YES
2.a.	YES	YES
2.b.	YES	YES*
2.c.	YES	YES
2.d.	NA	NA
2.e.	NA	NA
2.f.	YES	YES
3.a.	YES	NO*
4.a.	NA	NA
4.b.	YES	YES
4.c.	YES	YES
4.d.	NA	NA
4.e.	NA	NA
4.f.	NA	NA
4.g.	NA	NA
4.h.	YES	YES
5.a.	NA	NA
6.a.	NA	NA
6.b.	NA	NA
6.c.	NA	NA

6.d.	NA	NA
6.e.	NA	NA
6.f.	NA	NA
7.a.	NA	NA
7.b.	NA	NA
8.a.	NA	NA
8.b.	NA	NA
8.c.	YES	YES
8.d.	NA	NA
9.a.	NA	NA
9.b.	NA	NA
9.c.	YES	YES

As noted in the table above, the TSG believes that OpenID Connect/OAuth authentication meets all of the identified functional requirements. We identified two requirements, noted with an asterisk, that were problematic for mutual TLS authentication:

2.b.: 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.

3.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 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. Actor Models

To implement access to non-public gTLD 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.

4. Third Party Authorizers - organizations determining the types of data to be accessed by authenticated requestors.
5. Contracted Party 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 RDAP 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. Contracted Party Servers - RDAP servers operated by domain registries and registrars

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.

Actor Model 1: ICANN as Proxy 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 Proxy 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 Proxy 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 Proxy and Sole Identity Provider with Third Party Authorizers

In this model, 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.

8. 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 conduct access to non-public data using RDAP and well-known HTTP bearer token methods.

9. Proposed Solution

The authentication mechanism used between the client and the ICANN RDAP proxy will be based on OpenID Connect and OAuth 2.0 using either shared secrets (e.g., usernames and passwords) or digital certificates at the mutual choice of the Identity Provider and the client. An Identity Provider must support one authentication method and may also support the other. Other authentication methods may be added once approved by ICANN. As other authentication methods become standardardized, 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 Contracted Parties, and also between subsystems. This method is recommended because ICANN is fully authorized for access to non-public data, and the Contracted Parties only need to authenticate ICANN without 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 Contracted Parties.

9.1 Prerequisites

Identity Providers must be appointed and approved to perform client identification and authentication functions. Third party authorizers must be appointed to perform authorization and information association functions. ICANN may serve as an Identity Provider and/or an authorizer. The function 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 must exchange configuration information to identify service endpoints.

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. (In OAuth terms, these attributes are known as “claims” that are encoded in a data structure known as a “ID token”.) Policies must be developed to determine the set of attributes/claims that are needed to make data access decisions. Policies must also be developed to determine the attribute/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.

9.2 Processing Steps

A requestor who wishes to submit an RDAP query must submit an Access Request (as described below). An Access Request must be 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 also be operated by ICANN) 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 data elements. The Third Party Authorizer sends a response to the ICANN RDAP Gateway that indicates[1] the result of authorization processing. If the requestor is authorized, the ICANN RDAP Gateway sends RDAP queries to the specific contracted party RDAP servers that are authoritative (i.e., have the closest relationship to

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

the data subject) for the individual data elements within the requested data. These queries from the ICANN RDAP Gateway to the contracted party servers may contain secure metadata as specified by the system requirements and relevant policy. The contracted party RDAP servers each return RDAP responses containing the full set of data elements for which they are authoritative, which are received, processed, and filtered by the ICANN RDAP Gateway 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.

9.3 Data Flow Diagram

The processing steps described above are illustrated in Figure 1 below.

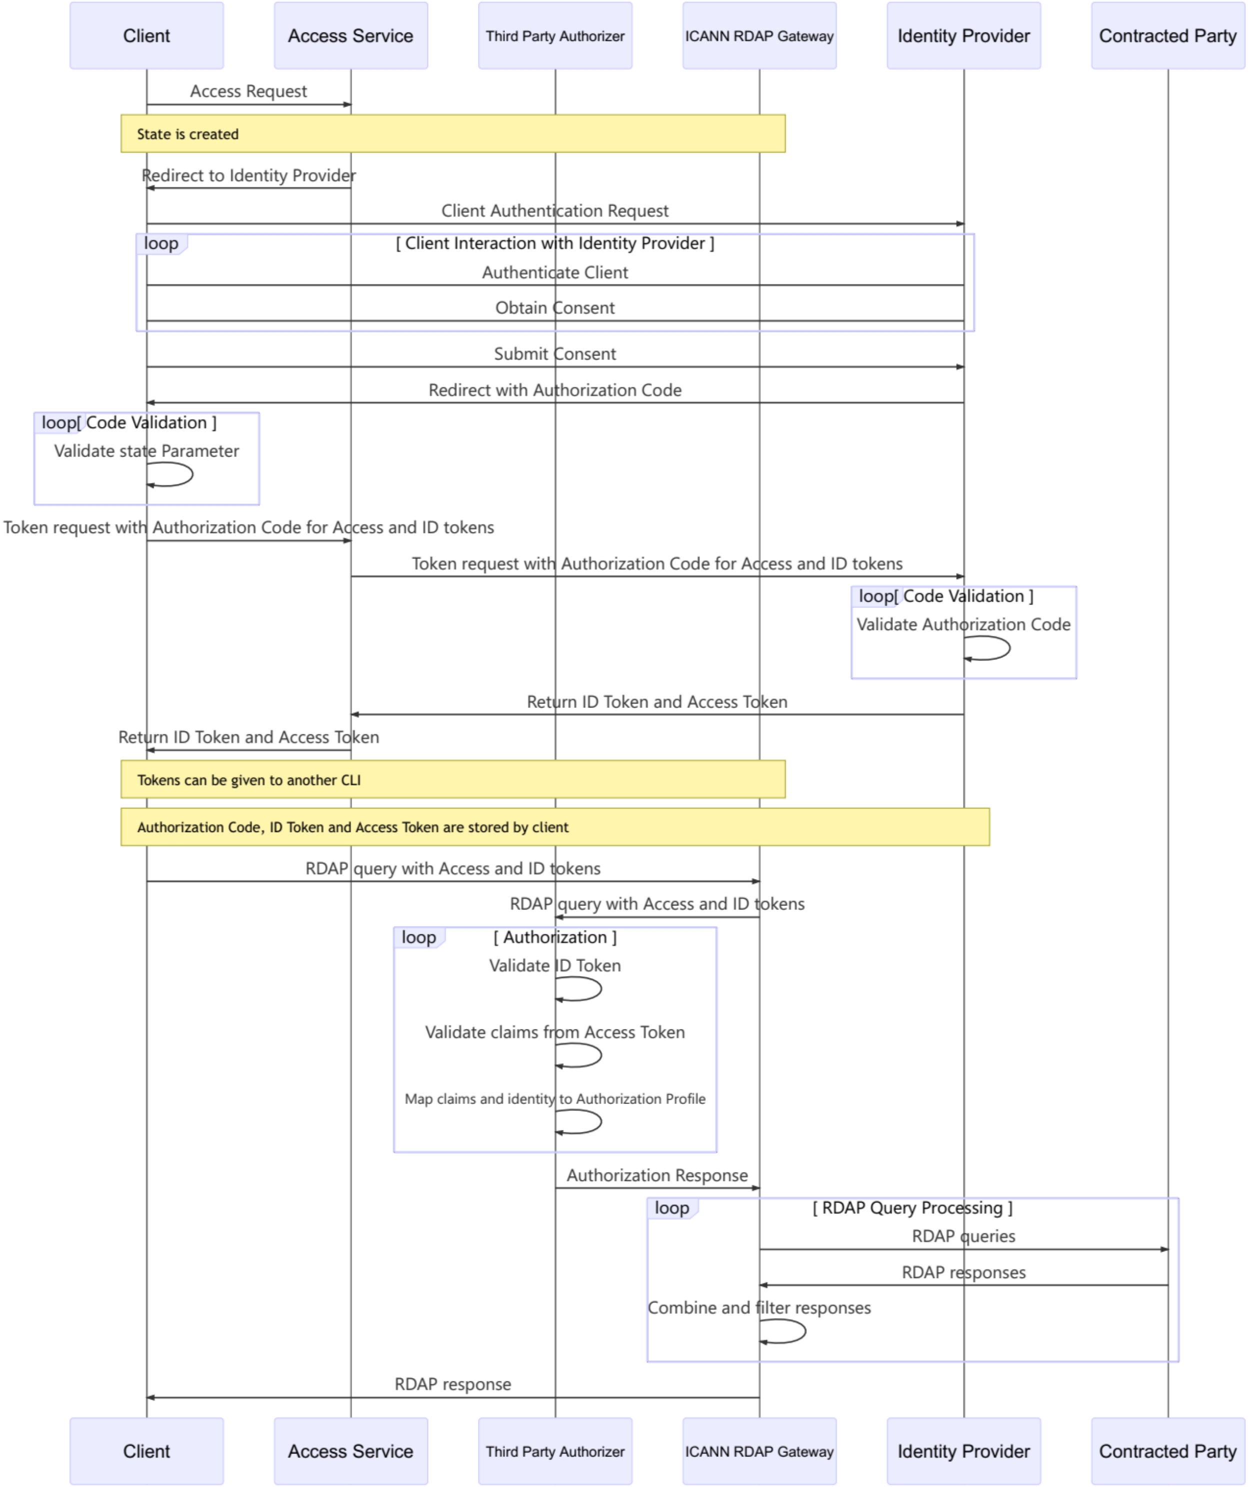
Client
Access Service
Third Party Authorizer
ICANN RDAP Gateway
Identity Provider
Contracted Party
Access Request
State is created
Redirect to Identity Provider
Client Authentication Request
loop
[Client Interaction with Identity Provider]
Authenticate Client
Obtain Consent
Submit Consent
Redirect with 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 1: Data Flow Diagram

10. 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.

10.1 Data Retention

The system requires various parties (such as ICANN org, identity providers and authorizers) to collect and store data of various kinds, such as user account information and transaction logs. As a matter of best practice, and also to comply with data protection law, the TSG believes that policies regarding retention and deletion of these data, which are outside of the TSG's narrow technical scope, SHOULD be established, communicated to the data processors, audited and enforced.

10.2 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 contracted parties will be subject to SLAs for operating their respective RDAP services. However, the Group believes that the other actors in the system should also be subject to SLAs which they would enter into with those parties who will rely on those actors to carry out their respective obligations, specifically:

- ICANN org (as the operator of the RDAP Gateway)
- Identity Providers (who both requestors and ICANN org will rely on)
- Third Party Authorizers (also relied on by requestors and ICANN org)

The TSG believes that defining the service level performance requirements for each party, the manner in which they are established, audited and enforced, and 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), to provide its users with a clear view of any disruptions which might affect their use of the service.

10.3 Obligations on ICANN Org

The TSG recognises that it is proposing a solution that could potentially impose significant operational burdens on the ICANN organization, especially if the community determines that the operator of the RDAP proxy must meet a stringent Service Level Agreement, and operate at significant scale.

It is RECOMMENDED that the ICANN organization review the spectrum of potential operational outcomes for deployment and operation of the system proposed, to determine the feasibility of such outcomes, their 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.

10.4 ICANN’s Role as Coordinating Party

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

10.5 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 contracted parties from disclosure of non-public registration data. The Group cannot comment on the validity of this assumption, and expects that the contracted parties will come to their own determination, based on their own legal advice.

10.6 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 Group recommends that ICANN consider publishing a regular Transparency Report[2] which provides statistics on requests for access to non-public gTLD registration data, similar to those published by other organizations.

10.7 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 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 contracted party.

11. Conclusion

The Temporary Specification for gTLD Registration Data 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 of registration data for legitimate purposes.

The TSG is charged with the task of developing a technical solution that enables ICANN to strike a balance between these opposing needs while still observing all of the requirements one would expect from a service that must be cautious about serving all of its constituents. The TSG has, with this document, delivered an outline of its working assumptions and requirements, and now solicits the feedback of the community.

To the maximum extent practical, the TSG has endeavoured to avoid either influencing policy decisions or being influenced by them by way of open implementation decisions, and the TSG will continue to do so.

The TSG has proposed a technical solution that it believes:

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

- accommodates all of the actor models described in Section 7
- 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

What remains after collecting and digesting community 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 while still being as simple as possible to construct, deploy, operate, monitor, audit, and scale as demand on the system grows, and to 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 registration data services.

References

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

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

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

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

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

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

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

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

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

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

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

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 group will be composed of invited members with technical backgrounds, including expertise in RDAP and authentication/authorization technologies.

Role	Name	Affiliation/Employer
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 John Crain Daniel Halloran Gustavo Lozano Diana Middleton Erika Randall	ICANN