Vision of a Distributed Collection and Access DNS Registration Model

Steve Crocker, Edgemoor Research Institute
1 July 2024

This memo sets forth a vision for a unified but distributed DNS registration data collection and access control system. The key elements are flexibility to conform to a wide range of local, regional and national policies, accommodation for multiple levels of privacy, the ability to define fine grained collection and access rules, and the ability achieve this vision incrementally and with compatibility with existing registration systems.

Objectives

These are numbered for ease of reference and do not indicate priority.

1. Provide access to registration data to authorized requesters for legitimate purposes.
2. Protect registrants against privacy violations, harassment, spam and related ills.
3. Fit within existing laws, including the GDPR.
4. Operate at reasonable cost.
5. Operate with manageable operational risks.
6. Be acceptable to the multiple competing constituencies.

The Registration Process

The main actors are the registry, the registrar, the account holder, and the registrant.

Registration of a domain takes place when a registrar's account holder[1] requests a domain belonging to a specific registry, is told it is available, reserves it, and completes the registration process. The registration process includes payment and designation of the registrant.

For ease of reading, here's the same description using fictitious actors. Albright Operations is a registry for the .bbq domain. BeGood Services is a registrar licensed to serve Albright Operations.[2] Swifty Fingers has an account with BeGood. Jim Cole wishes to obtain the domain toastypeanuts.bbq. On his behalf, Swifty logs into his account at BeGood, finds toastypeanuts.bbq is available, and registers it. In the process, he makes the required payment and names Jim Cole as the registrant.[3]

[1] Account holders are also called customers.
[2] This arrangement is generally not exclusive. There may be other registrars that serve Albright, and BeGood may serve other registries.
[3] The registrant is also called the registered name holder (RNH)

Swifty provides BeGood with several pieces of data. Additional data is generated automatically during the registration process. All told, a hundred or so distinct pieces of data – "data elements" – are associated with the registration of toastypeanuts.bbq. These data elements fall into broad groups.

DNS records	Nameserver (NS) and other records that will be published in Albright's public DNS servers.
Contact data	Name, organization, address, email, telephone number for the registrant and possibly for other contacts such as the Admin, Tech and Billing contacts. Which of these data elements is required, optional or not collected may vary according to the registrar's specific policy. Collection policies are discussed below.
Payment data	BeGood keeps a record of the payments and additional data such as the IP address Swifty used when he logged into this session.
Registration period	The registration date and expiration date.
Locks	There are several locks that can be set to prohibit future changes. Some of these are set within the registrar. Others are passed along to the registry.

This list is approximate and not necessarily complete.

All of these are collected or generated within the registrar. Some are always forwarded to the registry. Some are never forwarded to the registry. Others are forwarded to the registry if the registry requires them. The details are controlled by the registry's policy and any higher authorities with jurisdiction over the operation of the registry and registrar.[4]

The usual reason for registering a domain name is to provide translation of the domain name into an IP address to reach the registrant's systems on the Internet. Thus, if Joe Cole is running the business Toasty Peanuts and want to make sales over the Internet, he will have a server connected to an Internet Service Provider (ISP). The ISP will assign an IP address, e.g. 10.9.8.7,[5] for Joe Cole to use. Swifty enters this address into BeGood's system. BeGood will communicate it to Albright Operations, and Albright will put an entry into its public DNS servers that translates toastypeanuts.bbq to 10.9.8.7.

---

[4] ICANN is an example of a higher authority. It has jurisdiction over generic top-level domains (gTLDs). Country code top level domains, (ccTLDs) are subject to the jurisdiction of the country's or territory's government. For our purposes, we use the term Policy Authority to refer to these higher authorities.

[5] 10.9.8.7 is a local IP address per RFC 1918, https://datatracker.ietf.org/doc/html/rfc1918, and cannot be accessed over the public Internet. It is used here as an illustrative example.

The Request Process

Although the express purpose of registering a domain name is provide translation of the domain name into an IP address to reach registrant's systems on the Internet, many others are interested in the registration data. Historically, registration data was accessed via the Whois protocol. Prior to the restrictions described in the next paragraph, it was estimated that five billion Whois queries were made each month across the approximately 400 million domain name registrations.

In the early days of the Internet, most of the registration data was publicly available. While access to this data served many socially desirable purposes, open access also permitted rampant use of the registration data for purposes harmful to the registrants. Examples include spam, propagation of malware, harassment, and other forms of abuse. In response to these widespread abuses, restrictions on access to registration data have become commonplace. The best-known restrictions are derived from the European Union's General Data Protection Requirement (GDPR). However, similar restrictions have also been imposed by other authorities. Further, in addition to the restrictions imposed by Policy Authorities, various registries and registrars have adopted their own restrictive rules.

Implementation of these restrictions

In response to the GDPR, ICANN and other organizations adopted a very cautious and conservative approach. The result has been to almost totally shut off access to non-public DNS registration data. That's improved protection, but it has had a deleterious effect on the use of registration data for legitimate purposes.

Our Approach

We envision a system design with the following characteristics.

Collection Rules

1. Each registrar has rules governing its collection of registration data. These rules specify which data elements are required, which are optional and which are not collected.

2. There is a separate specification regarding the accuracy level for each data element.

3. A sensitivity level is assigned to each data element. The sensitivity level embodies the basic idea of public vs private data but has four levels instead of just two.

4. The registrar may have different sets of rules for different subsets of potential registrants. The most common distinction is between natural and legal persons, but other distinctions such as nexus, whether the registrant is at risk, and whether the

registration data contains PII,[6] may be supported.

5. Registries, policy authorities and governments may have rules that set the boundaries on the registrar's rules.

Disclosure Rules

Requests for registration data come from a variety of parties. In principle, each request must be evaluated to see whether it meets the criteria satisfactory to the registrar and consistent with legal requirements. The details of these requirements may vary according to the registrar and the jurisdiction. In general, they will cover the following concepts.

Purpose	Does the stated purpose of the request fit within the registrar's, policy authority's and government's requirements?
Protection	The requestor is obliged to limit their use of the data to the approved purpose. The requestor is further obliged to protect the data from misuse by others, both within and outside of the requestor's organization.
Requestor	Is the requestor trustworthy and accountable? This determination will usually include some degree of identification of the requestor but will almost always include additional details related to enforcement of the use and protection requirements.
Data Elements	What data elements is the requestor authorized to receive? This provides a degree of control more detailed than assuming all non-public data is to be disclosed. For example, for some purposes it is sufficient to provide only the registrant's country code but not the city and street address.
Sensitivity Level	What is the maximum sensitivity level the requestor is authorized to receive. Together with the Data Elements list, these provide a two-dimensional sieve that can be tailored to a nuanced set of policies.

Distributed Design

A system consistent with this approach can be built without a central control. The policy decisions can be localized. The only aspects that need to be common across the system are the names of the data elements and the protocols for making requests, communicating responses, etc.

[6] In some of our field work, we have heard privacy advocates insist that some legal person registrations contain PII and should be protected similarly to natural persons. We take no position on this except to provide the hooks to implement this distinction for those registrars who wish to provide it.

A distributed design makes it possible for multiple organizations to provide the necessary software. It also makes it possible for this design to be adopted incrementally.

Optimization

Taken literally, the request and disclosure processes treat each request as if it is entirely unrelated to any previous request. In practice, a very large fraction of the requests will be nearly the same as previous requests. Usually, these similarities will be related to the patterns of usage within a community of requestors. Example communities of requestors are law enforcement agencies, intellectual property attorneys, security practitioners, and security researchers. (This list is illustrative, not complete.) Members of these communities are likely to make requests over time that share the same characteristics.

One form of optimization is for each requestor and each registrar to keep track of past requests. Requestors will tend to make subsequent requests that are like past successful requests. On the registrar side, if a request is like a past successful request, the evaluation process can be expedited.

A stronger form of optimization is when requestors and registrars form agreements in advance of the request and response process. Such agreements will usually include the details set forth above under Disclosure Rules. If such agreements are in place, the request and response processes can be reduced to checking that a request is consistent with an existing agreement, and the cycle can be automated.

Automation has the threefold advantage of certainty, speed, and greatly reduced cost. However, automation also requires trust and experience. We automation will appear incrementally across the system. And, of course, requests that do not fit into an existing agreement will need to be evaluated individually.

One way to achieve and administer agreements is for groups of requestors – a Requestor Group(RqG) -- and groups of registrars – a Registrar Group(RrG) -- to form consortia that work out the details on behalf of their members.

The slide deck, **A Holistic, Engineering Approach to Collection and Disclosure of Registration Data for Internet Identifiers,** covers the concepts in this note and includes a partial picture of the interactions.