

The Public Suffix List: A Guide for TLD Administrators

ICANN Office of the Chief Technology Officer

Paul Hoffman
OCTO-011
18 May 2020



TABLE OF CONTENTS

1	INTRODUCTION	3
2	WHAT IS A PUBLIC SUFFIX?	3
3	WHAT IS THE PSL?	3
4	HOW TO MAKE CHANGES TO THE PSL	4
5	HOW THE PSL IS MANAGED	5

This document is part of the OCTO document series. Please see the [OCTO publication page](#) for a list of documents in the series. If you have questions or suggestions on any of these documents, please send them to octo@icann.org.

This document was originally prepared by Jothan Frakes under contract to ICANN. ICANN also thanks the PSL volunteers for contributions to this document.

1 Introduction

The Public Suffix List (PSL) is a machine-readable list of domain names that support subdomains operated by unaffiliated organizations. It can be useful for software, particularly web browsers, to know which domains allow registration beneath them and which do not. All top-level domains (TLDs) in the Domain Name System (DNS) are part of this list, and the PSL is an important component of how many applications function, so every TLD administrator should be aware of it and ensure the list's data for their TLD is up-to-date and accurate.

This guide is intended for TLD administrators. The document first defines what a public suffix is, and then describes what the PSL is and how it is commonly used. It also details how TLD administrators can submit changes to the PSL. Finally it explains how the PSL is managed and sets some expectations for how changes are incorporated across the Internet.

2 What is a Public Suffix?

The volunteers who run the PSL define a public suffix as: a domain under which multiple parties that are unaffiliated with the operator of the domain may register subdomains. ICANN's Security and Stability Advisory Committee (SSAC) defines a public suffix slightly differently: a domain that is controlled by a public registry, under which subdomains can be registered by third parties.¹

Common examples of public suffixes are TLDs such as *.com*, and some second-level domains such as *.co.uk*, where domain names can be acquired through domain name registrars. Public suffixes can include third-level domains like those used by Tumblr, where a subdomain is created under *.tumblr.com* for every user of the service.

3 What is the PSL?

The PSL is a data file. The file is a list of all known DNS public suffixes and rules pertaining to each. The data file and full documentation can be found on the PSL's website.² The PSL started as an initiative from Mozilla, but it is now maintained by a group of external volunteer maintainers. As such, it is outside the direct management or control of ICANN, the IANA functions, or the IETF.

A significant goal of the PSL is to help developers write code to handle public suffixes. The PSL was originally primarily intended for browser vendors to implement functionality like protecting against "super cookies" being set on the wrong domain name boundary. But use of the PSL has expanded to many other applications. Common uses now also include:

- ⦿ **Web security:** Many security features of browsers, including cookies, process isolation, and scripting access, are restricted to being part of the same boundary of a domain

¹ See SAC070, "SSAC Advisory on the Use of Static TLD / Suffix Lists", <https://www.icann.org/en/system/files/files/sac-070-en.pdf>

² See <https://publicsuffix.org>

name. When a service such as Tumblr is shared between multiple users through subdomains, browsers need to know how to appropriately isolate these different users.

- ⦿ **Domain resolution:** In applications it is sometimes unclear whether a string is intended as a domain name or something else. For example, some browser location bars can accept a URL, or a domain name that is used as a URL shortcut, or a search term. One common heuristic used in such programs is to assume that a domain name that matches a name in the PSL is a search term, while other domain names are shortcut URLs.
- ⦿ **Organization domains:** The Domain-based Message Authentication, Reporting & Conformance (DMARC) anti-spoofing protocol³ lets an organization publish a blanket sending policy for all of its subdomains. DMARC code uses the PSL to find the organization domain for any subdomain.
- ⦿ **Program logic:** Browsers and other applications can use the PSL to highlight the most relevant part of a domain name (for instance, *example.com* rather than *www.example.com*). Also, some applications attempt to identify links within message content, and convert them to hyperlinks using the PSL to narrow what is interpreted as a domain name.

Because the PSL is used in some popular applications and is incorporated in many widely-used software libraries, it is very important that TLD administrators curate their data in the PSL to ensure accuracy and to make updates when a TLD's configuration changes. A correct entry for a TLD could mean the difference between seeing search engine results instead of a website, or malicious spam being treated as legitimate business email.

The content of the PSL adheres to ICP-3⁴, which defines a single, authoritative root, and recognizes ICANN, the IANA, and the IETF as the authorities on this matter. Alternative root namespace entries submitted to the PSL are not accepted.

There is a mechanism for exceptions to ICP-3 in the PSL. These exceptions are limited to cases where there are suffixes clearly defined by an IETF RFC, or where there is a TLD that ICANN has included in their list of those TLDs with a contractual agreement and is scheduled to be added to the IANA root zone. At the time this document is being published, there is only one such exception in the PSL: *.onion* from RFC 7686.⁵

4 How to Make Changes to the PSL

Updating the PSL is not difficult, but it requires some understanding of the change process. All updates are done through the PSL's repository on GitHub.⁶ The process is secure and stable, with a number of checkpoints to reduce the potential for unauthorized changes. Before requesting changes to the PSL, you need an account on GitHub, and a reasonable knowledge of how to use GitHub, particularly how to use GitHub's *pull requests*.⁷

The PSL has two sections, ICANN and private. The ICANN section covers TLDs that are assigned by IANA; the private section covers non-TLD domains, and is mostly used to show

³ See <https://dmarc.org/>

⁴ See <https://www.icann.org/resources/pages/unique-authoritative-root-2012-02-25-en>

⁵ See <https://datatracker.ietf.org/doc/rfc7686/>

⁶ See <https://github.com/publicsuffix/list>

⁷ See <https://help.github.com/en/github/collaborating-with-issues-and-pull-requests/about-pull-requests>

their domain security policy (which can be useful for cookie policies in browsers). Some PSL applications focus on the ICANN section, while others also incorporate the private section. As a TLD administrator, you are likely submitting changes only in the ICANN section.

The PSL is kept in a single text file, called `public_suffix_list.dat`. Each line contains a public suffix. Lines can also contain an optional wildcard character to indicate that every name that matches the wildcard is also a public suffix. TLD administrators should be sure they understand the file format before proposing changes to it.

To submit a proposed change to the PSL:

1. Carefully read the PSL guidelines.⁸
2. Clone the PSL git repository.
3. Make the desired edits to the `public_suffix_list.dat` file in the cloned repository.
4. Submit the changes to that file by creating a pull request on Github.
5. Put a TXT record in the zone to link the pull request with the zone administrator.

The pull request for changes to the PSL should include as much information as possible on the context of the change and on your administrative relationship to the TLD(s) being modified. You can look at existing pull requests as reference. Functional and complete contact information is needed in the pull request in case the PSL maintainers have questions or need to clarify an aspect of the submitted change.

After the pull request is submitted, the PSL maintainers will be notified and will review the request in due time. They may respond to the pull request with questions, so you should be following it through its tracking email and also by checking the GitHub repository. It is considered bad form to submit multiple requests for the same change just because you think that the first request is not being acted on quickly enough. In fact, submitting multiple requests will likely delay action on the initial request because it adds additional processing that the maintainers must perform.

The validation process requires a DNS TXT record with the request number from the GitHub system be added to the zone corresponding to the domain being updated by the request. For changes to names where TXT record validation is not possible manual validation is necessary. This takes longer, and usually works through the contact email for the TLD in the public IANA TLD database.

5 How the PSL is Managed

Because the PSL is run by volunteer maintainers, not paid staff, timing in the process is not always predictable and there can be delays in reviewing changes. Although anyone can propose changes to the PSL, all proposed changes are validated by the PSL maintainers before being accepted to ensure that changes are coming from a valid administrative source.

⁸ See <https://github.com/publicsuffix/list/wiki/Guidelines>

If a TLD registry publishes its registration policies on its website, the change submission is likely to be expedited. Submissions that can be successfully validated against the registry operator's own public documentation and policies may not require any further validation or authorization. Submissions can come from parties other than the direct administrator of a TLD or domain name, and the PSL maintainers have some protections in place to ensure that the requests are validated. For generic TLDs (gTLDs), the PSL maintainers automatically pull the authoritative list of gTLDs from ICANN. In doing so, the output of the tool could conflict with manual edits made by gTLD operators.

After a change is made in the PSL, it's important to understand that some software and services on the Internet that rely on the PSL get updated slowly or not at all. Although some developers have their software regularly fetch the current PSL, others build a static list into their software or update with lower frequency. Some of the default mobile and tablet browsers only update when the underlying operating system updates; this can clearly cause problems for users.

Some developers may choose to not include all PSL entries, or to introduce additional items. The PSL does not compel or prescribe how the list gets used or if it is used in its entirety. A developer is free to make their own determination of inclusion or exclusion of portions of the list at their own discretion or modify it for their own purpose. Therefore it is important to realise that while keeping entries in the PSL as up to date as possible is important, discrepancies may still be seen due to factors beyond your control.

The PSL has a very low-volume mailing list for general discussion topics.⁹

⁹ See <https://groups.google.com/forum/#!forum/publicsuffix-discuss>