Guide to Name Collision Identification and Mitigation for IT Professionals

1 August 2014
Version 1.1

ICANN

Table of Contents

1. Introduction

Following the entry of a new top-level domain name into the global DNS root, organizations may see that queries to resolve some of the “internal” names specific to their network return different values, providing users and programs with different results. There are two basic issues: “internal” names that are leaking into the global Internet, and private name spaces that are defined in conflict with the global DNS name space.

The cause of such different results is that a DNS query a network administrator intended to have resolved locally, using an internal name space, is now being resolved using the new top-level domain data in the global DNS. Under these circumstances, queries that were never anticipated to leave the internal network are now getting results in the global DNS, and those results are different. At a minimum, leaked names that produce differing results may be bothersome to users (e.g., they may cause delayed access to web pages). They may also pose security problems (such as email being sent to the wrong recipients).

This document covers mitigation and prevention strategies for the most common types of private namespaces used by organizations. This document describes what organizations may encounter when internal names leak to the global DNS and specifies recommended mitigation practices. The description and advice offered here is aimed at IT professionals (network administrators, system administrators, and IT department staff) who understand in general how the DNS works and how their own internal name systems work. Readers who want more background are referred to the documents in Appendix A. Readers concerned with security are directed in particular to the reports from ICANN’s Security and Stability Advisory Committee (SSAC).

ICANN, the organization that administers the contents of the global DNS root, has prepared this document in consultation with name space subject matter experts to assist organizations whose private namespaces may be in conflict with the global DNS root. ICANN has published other documents describing how the global DNS is organized, how new names are added to the DNS root, and more. Appendix A of this document lists references on many topics for further reading. In addition, ICANN has recently begun helping organizations that are using private namespaces know when those namespaces will begin having collisions; this is described in Section 1.4 and Section 6.

Note that although this document addresses mitigation measures for name collisions, it only discusses problems that organizations may encounter when resolving names. It does not address other issues related to the operation of the global DNS itself. For example, the root name servers of the global DNS have always been inundated with queries that were never intended to be processed by the global DNS (see *SAC 045* in Appendix A), but the root name servers have also always been provisioned well enough to be able to field these excess queries. Related issues regarding root names servers are not addressed by this document. It addresses only the consequences of queries that are inadvertently leaked to the global DNS root name servers.

ICANN has developed a webpage that provides informational materials regarding name collisions available at http://www.icann.org/namecollision. The page also includes a process for reporting demonstrably severe harm as a consequence of name collisions caused by the new generic Top-level Domains (gTLDs).

1.1 Name Collisions

The global DNS is a hierarchical namespace, and the names in the DNS are composed of one or more labels that create a full name. At the top of the hierarchy is the DNS root zone that contains a set of names such as `com`, `ru`, `asia`, and so on; these are the global TLD*s* (top-level domains), commonly referred to simply as “the TLDs”. An example of a full domain name (often called a *fully-qualified domain name* or *FQDN*) might be `www.ourcompany.com`.

Almost all private namespaces are also hierarchical. There are three major types of private namespaces:

- **Namespaces that branch off of the global DNS** – Private namespaces that branch off the global DNS are rooted under a name that is resolvable in the global DNS, but the entire directory structure under that name is managed locally with names that IT administrators never intended to be seen in the global DNS. For example, consider a private namespace rooted under `winserve.ourcompany.com`: the names in that private namespace (`winserve`) are managed by the private nameserver and are not visible in the global DNS.

- **Namespaces that use their own roots with private TLDs** – The root of the private namespace is a single label that is not a global TLD. The entire directory structure, including that of the private TLD, is managed by private nameservers that are not visible in the global DNS. For example, if the private namespace is rooted in `ourcompany`, then private nameservers are also responsible for `www.ourcompany`, `region1.ourcompany`, `www.region1.ourcompany`, and so on. There are many different types of namespaces that use their own roots with private TLDs. Examples include Microsoft’s Active Directory (in some configurations), multicast DNS (RFC 6762), and older LAN directory services that are still used in some corners of the Internet.

- **Namespaces that are created though the use of search lists** – A search list is a feature of a local name resolver (either for a private namespace or a recursive resolver for the global DNS). A search list allows a user to enter shorter names for convenience; during resolution, the nameserver appends configured names to the right of the name in a query. (These configured names are also called *suffixes*.)

Namespaces that branch off of the global DNS only cause name collisions when combined with search lists. Any query that involves an FQDN that comes from the global DNS will by definition never have a name collision with a different name in the global DNS. Such a query could only cause name collisions when inadvertently created through the use of search lists.

The concept of “private namespaces” confuses many people who are largely accustomed to typical Internet usage, i.e., people who are only familiar with global DNS naming and who may be surprised to learn that some requests for name resolution do not or should not result in a query to the global DNS. They may be even more surprised to learn that some queries for names are purposely meant to start in the private namespace, but end up in the global DNS. One reason name collisions may occur is that queries intended for a nameserver of a private namespace incorrectly started in the global DNS instead.

1.2 Name Collisions Due to Private TLDs

Name collisions occur as the result of two events. First, a query for a fully-qualified domain name that is rooted in a private TLD leaks from the private network to the global DNS. Second, the query locates in the global DNS the exact same name that exists on the private network under the private TLD.

A common cause of such name collisions is the use of a name in a system like Microsoft’s Active Directory that is not a TLD in the global DNS at the time that the system is configured, but is later added to the global DNS. This kind of name collision has already happened many times before and is expected to continue with the introduction of new TLDs in the global DNS (see *Introduction to the New gTLD Program* in Appendix A).

1.3 Name Collisions Due to Search Lists

Another cause of name collisions is search list processing. If a query is not an FQDN, it is a *short unqualified name*. A search list contains one or more suffixes. These are iteratively appended to the right side of a query. When a resolver is unable to resolve a short unqualified name, it appends suffixes from the list as it attempts to resolve the name until a matching name is found. A search list is a useful feature; however, search list processing accommodates the use of short unqualified names that are not FQDNs and thus inadvertently creates namespaces that are not rooted in the global DNS. In this case, the name collision occurs when a string the user intends to use as a short unqualified name is instead completed by the search list and resolved as an FQDN.

For instance, assume a name resolver has a search list that consists of the suffixes `ourcompany.com` and `marketing.ourcompany.com`. Further assume that a user enters `www` into a program that uses that resolver. The resolver might then first look up `www`, and if that didn't return a result, it might then look up `www.ourcompany.com` and the `www.marketing.ourcompany.com`.

Note the use of the word "might" in this example's description. Rules for how search lists are to be applied when doing name resolution vary across operating systems or applications. Some systems will always try to resolve a name either in the private namespace or the global DNS before applying the search list. However, other systems will use the search list first if the string being searched doesn't contain a "`.`" character. Still others will use the search list if the string being searched ends with a "`.`" character. Some operating systems and applications (such as web browsers) have changed their search list rules multiple times. It is thus impractical to predict when search lists will or will not be used, what is or is not a short unqualified name, and thus whether or not short unqualified names are likely to leak out to the global DNS. See *New gTLD Concerns: Dotless Names and Name Collisions* in Appendix A for more detail about diversity of search list processing.

This description of search lists may come as a surprise to some readers because they are so common in places that at first glance do not appear to create "private namespaces". Every suffix in a search list defines another namespace that may be consulted during name resolution. This creates a private namespace that operates reliably only when the client queries the particular resolvers for that namespace. Depending on the search list implementation, some name resolvers might even try the short unqualified name entered by the user or configured in software before appending any of the names in the search list. For example, typing `www.hr` in one location on the Internet might produce one result from the DNS resolver, but typing it in a different location might produce a different result. When this occurs, one of those namespaces is "private" relative to the other.

Using search lists instead of resolving FQDNs via the global DNS contributes to name resolution uncertainty. Name collisions produced by search lists are hard to predict because search lists are so common. They are part of the name resolver software in many operating systems, network equipment, servers, and more. Resolver software acts differently from system to system, between different versions of the same operating system, and even as a function of the operating system's or application's view of where on the network the request is coming from. Deploying a name resolution service that resolves names using only the global DNS is the best assurance against such uncertainty and unpredictable results.

1.4 Assisting the Detection of Name Collisions in New gTLDs

From August 18, 2014 onward, when a gTLD is delegated from the DNS root zone, the gTLD is required to conduct a *controlled interruption* service for 90 days. During the controlled interruption period, easily identifiable answers are sent from the authoritative name servers for the new gTLD for a variety of DNS queries. The purpose of these answers is to warn organizations that will experience name collisions that they need to take immediate action to prevent possible damage due to the leaked queries.

Further, from the same date, some gTLDs that are already in the root zone are required to conduct a controlled interruption service for 90 days before they delegate certain second-level names into the global DNS. The purpose here is the same as above: to warn organizations that are leaking private queries that they need to mitigate possible damage as soon as possible.

Note that these rules only apply to gTLDs, not to TLDs that are for country codes (usually called “ccTLDs”). When a ccTLD is added to the root zone, its operator can choose to have a controlled interruption, but is not obliged to do so.

2. Problems Caused by Name Collisions

Name collisions based on queries that leak into the global DNS from private networks can have many unintended consequences. When a query gets a positive response, but with an answer that is from the global DNS instead of the expected private namespace, the application making the query will attempt to connect to a system that is not part of the private network, and may succeed. Such a connection could be a nuisance (by introducing delay during name resolution). It could also prove to be a security issue, i.e., it may create a vulnerability that might be exploited for malicious purposes, depending on what the application does after connecting.

2.1 Direction to Unexpected Web Sites

Assume that a user enters `https://finance.ourcompany` in his or her web browser while on a private network, and that the network has a namespace whose private TLD is `ourcompany`. If the browser's query for the name `finance.ourcompany` resolves as expected, the browser gets an IP address for the internal web server of the finance department. Imagine, though, that the TLD `ourcompany` is also part of the global DNS, and that that TLD has a second-level domain (SLD) name `finance`. If the query leaks, it will resolve to a different IP address than it did when the query was resolved in the private namespace. Now imagine that this different IP address might host a web server. The browser would attempt to connect to a web server on the public Internet, not on the private network.

As shown earlier, the same problem can happen even in networks that do not have private TLDs, but do use search lists. Consider a browser that is normally used on a network where users have a search list that has the name `ourcompany.com`, and the user enters the name `www.finance` in order to get to the host `www.finance.ourcompany.com`. Now imagine that the browser is being used by an employee from a mobile device at a coffee shop. If that query leaks out to the Internet, and there is a TLD called `finance`, the query could resolve to a different IP address, e.g., a completely different host whose name in the global DNS is `www.finance`. That query would cause the browser to attempt to connect to a web server at a completely different part of the public Internet than it would have if the query had gone to the resolver on the private network.

A common user response to this scenario is that the user would recognize that this was the wrong web site and would just leave immediately. However, a browser can expose a great deal of information to a web server if the browser "trusts" the web server because it has the same domain name as one that the browser has visited earlier. The browser might automatically enter login or other sensitive data, thereby exposing that information to capture or analysis outside the organization. In other circumstances (e.g., a carefully formulated attack against the organization), the browser might connect to a site hosting malicious code that installs dangerous programs on the computer.

Note that the use of TLS and digital certificates might not help prevent the damage due to name collisions; in fact, it might make it worse by giving users a false sense of security. Many of the certificate authorities (CAs) that issue certificates for names in the global DNS also issue certificates for short unqualified names in private address spaces, so it is possible that a user who is misdirected to a site would still see a valid certificate. See *SAC 057* in Appendix A for more detail on certificates with names from private namespaces.

2.2 Direction of Email to the Wrong Recipients

The possible consequences arising from name collisions is not limited to web browsers. Email intended for one recipient can be sent to a different recipient if the host names in the recipient's addresses are the same; for example, email to `chris@support.ourcompany` could be delivered to a completely different user account if `ourcompany` becomes a TLD in the global DNS. Even if the message is not delivered to a particular email user, there might be an attempt to send it, and such attempts might expose the email contents to capture or analysis outside an organization..

Many network devices such as firewalls, routers, and even printers may be configured to send notifications or log data by email. If the recipient name that was entered for email notifications later is subject to name collision in the global DNS, the notification might be delivered to a completely unintended recipient. Event or log data in the message body that may reveal network configuration and host behavior may leak to an unintended recipient. Routine network performance or traffic analysis by IT staff may be interrupted if the intended recipient of such data never receives the log data, or events that trigger notifications may not be investigated or mitigated.

2.3 Security Reductions

Name collision occurrences that are left unmitigated may expose systems in private networks to unintended behavior or harm. Systems that rely on name resolution for correct operation and that also perform security functions *can* perform reliably when they use FQDNs and resolve them from the global DNS.

For example, in firewalls, security rules are often based on the source or destination of a packet flow. The source and destination of packets are IPv4 or IPv6 addresses, but many firewalls allow them to be entered as domain names as well. If short unqualified names are used and name resolution is not performed as expected, the rules may fail to block or allow traffic as the administrator intended. Similarly, firewall logs often use domain names, and using short unqualified names that resolve in unpredictable ways can interfere with event monitoring, analysis, or response. IT staff reviewing the logs might, for example, misunderstand the severity of an event because a short unqualified name in the log might identify different hosts depending on where the log was created (that is, in the log, the same short unqualified name might appear to be associated with two or more different IP addresses). This problem may be compounded by the fact that most firewalls can act as their own DNS resolvers or allow administrators to use or configure search lists.

2.4 Systems Affected by Name Collisions

All network-attached systems should be checked for use of host names that are rooted in a private TLD or host names that are based on search lists. All of these "use" instances will need to be updated to use a FQDN from the global DNS. A non-exhaustive list of systems or applications to check would include:

- **Browsers** – Web browsers allow users to specify the location of HTTP proxies, and these are very often on the private network. Check whether a user or IT staff has made custom home pages, bookmarks or search engines: these can have links to servers on the private network. Some browsers also have configuration options for where to get revocation information on SSL/TLS certificates that might point to host names on the private network.

- **Web Servers** – Web servers offer HTML content that contains links and metadata that have embedded host names. Check whether web servers on a private network have content with short unqualified names. Check whether the configuration files for the web server have short unqualified names of other hosts on the private network.

- **Email User Agents** – Email clients such as Outlook and Thunderbird all have configuration options for where to receive email using the POP or IMAP protocols, and where to send email over the SUBMIT protocol; all of these might use host names on the private network. Check whether these applications are configured to obtain revocation information on SSL/TLS certificates from hosts assigned short unqualified names.

- **Email Servers** – Check whether email servers have configurations that list the short unqualified names of other local hosts, such as backup email gateways, offline storage servers, and so on.

- **Certificates** – Check whether applications that employ X.509 certificates, such as telephony and instant messaging programs have configuration data that use short unqualified names to identify where to get revocation information on SSL/TLS certificates.

- **Other Applications** – Custom applications can have many configuration parameters where host names might be stored. The most obvious space would be in configuration files, but host names could appear in many kinds of application data, links on social media or wiki sites, or even hard-coded in source code. Check these configuration data for short unqualified names.

- **Network Devices** – Check network infrastructure devices – firewalls, security information and event management (SIEM) systems, routers, switches, network monitoring devices, intrusion detection or prevention systems, VPN servers, DNS servers, DHCP servers, log servers – to determine whether these are configured with short unqualified names of other devices on the private network.

- **Client Administration** – Check whether centralized client administration tools such as those that configure an organization’s workstations and network devices have short unqualified names in the configurations (particularly search lists) that are controlled and reset by the systems.

- **Mobile Devices** – Consumer devices such as phones and tablets may have similar configuration options as some of the applications listed above, and thus possibly have configuration choices that might contain short unqualified names from the local network.

All of these systems should be checked for configuration data that store short unqualified names to ensure that such names can be changed when the root of the private namespace changes or when search lists are no longer used.

3. When to Mitigate Against Name Collisions

Names are sometimes added to the global DNS root zone, such as when a country's name changes, or when ICANN delegates new TLDs. Both types of top-level domains have been added nearly every year for well over two decades. New TLDs have been added in 2013 and 2014, and it is a certainty that more will be added in coming years.

History shows that some name collisions have occurred when TLDs are added to the DNS. History also shows that names from private namespaces have leaked for many years, in some cases with very high frequency; see *SAC 045* in Appendix A for more detail. History illustrates that name spaces and name resolution intended for private networks are never as thoroughly segregated as administrators think they are, and that name queries that administrators intend to be resolved by internal nameservers are instead sometimes sent to resolvers in the global DNS.

Network administrators sometimes make choices of names based on assumptions that the list of names in the root of the global DNS is immutable, but that list in fact has and will change over time. For example, when the `cs` TLD was added almost 25 years ago for the country of Czechoslovakia; many universities were using search lists that allowed a user to enter a name ending with `cs` for the Computer Science department that would be fully qualified with the university's domain name, and these decisions resulted in name resolution uncertainty when the new TLD was added to the root zone because the names ending in `cs` were now FQDNs in the global DNS. Even when current global DNS root names often do not overlap with those in a private namespace (either a private TLD or search list), network administrators often forget to stay up-to-date on which names are in the global DNS root.

It is recommended that an IT department begin mitigation efforts as soon as feasible. Taking a stance of "we'll just do our firewalling better" may reduce some collisions, but will never eradicate all of them. Similarly, saying "We'll get our users to be sure to use our nameservers" or "We'll make remote workers use VPNs" are likely to reduce some collisions, but these may also make the remaining collisions harder to diagnose.

Name collisions can happen regardless of the characters in the name; however, the use of non-ASCII characters such as ä and 中 and й in private TLDs complicates analysis of collisions. Resolvers may send out queries for these in ways that are hard to predict, and may not match the Internet's standards, so determining when name collisions will happen becomes much more difficult.

Although the global DNS root will end up larger than it has been in past years, the adding of names to the root is really not all that unusual. For each new TLD being added, there is a chance that there will be name collisions with private namespaces that have been leaking to the Internet mostly unnoticed. Organizations have been using names and assuming the risk of collisions for years.

Note that the addition of new names to the DNS root is not, and will never be, a problem to organizations already using FQDNs from the global DNS in their network. These organizations will see no difference to their own use of DNS names, because there are no name collisions. The problems only appear for organizations that are using private TLDs, or organizations that use search lists that allow entry of short unqualified names where the shortened name itself might be a valid name in the global DNS.

3.1 Determining the Potential for Collisions

In order for you to determine whether or not there will be name collisions with your organization's private namespace, you need to identify and catalog all of the private namespaces and DNS search lists your organization uses, and then compile a list of the top-level names in these sources. For most organizations, there is typically only one namespace with just one top-level name, but some organizations, particularly those that have combined with other organizations that were also using private namespaces (for example, as a result of a business merger or acquisition), have multiple private top-level names.

Next, you need to determine both the current and expected contents of the global DNS zone. The names in the current root zone for the global DNS can be found at http://data.iana.org/TLD/tlds-alpha-by-domain.txt. To determine whether a name from a private namespace is being considered for allocation through the current new gTLD program:

1. Go to https://gtldresult.icann.org/application-result/applicationstatus
2. Click the arrow in the "String" column
3. Scroll through the pages until find the range that contains your private namespace's name

If there is any overlap between the list of private TLDs you have just made and the list of names in the DNS zone, there is a chance that there will be name collisions, and therefore mitigation is needed now.

Note that after the current round of new TLDs are entered into the root zone, more may be proposed; in particular, the list of new TLDs may change and name collisions between private namespaces and future new TLDs may occur. Also, organizations with private TLDs consisting of two letters (such as `ab`) should be aware that two letter top-level domain names are reserved for use as country codes, and these are added to the root zone through a completely different procedure.

3.2 Global DNS gTLDs Whose Delegation Is Deferred Indefinitely

ICANN has stated that it will indefinitely defer delegating three TLDs: .corp, .home, and .mail. These gTLDs are still in common use in private namespaces, and thus pose a significantly higher risk for collisions than other TLDs. The deferral is not guaranteed to be forever, so any organization using one of those names as a private namespace should still follow the directions in Section 4 or Section 5 for migrating from the private namespace. However, such organizations have more time to perform the migration than an organization that has used a different name that might appear in the global DNS root in the foreseeable future.

4. Steps to Mitigate the Problems Associated with a Private TLD

Using private TLDs has not been recommended as a best practice for decades. In fact, the instructions that come with Microsoft's Active Directory and Server products have explicitly discouraged the use of private TLDs for many years. The most effective mitigation for name collisions due to names that end in a private TLD leaking to the global DNS is to change from using a private TLD to a one that is rooted in the global DNS.

The steps in this section apply to any network that has for its own reasons chosen to use a private TLD as its root and to use search lists to resolve short unqualified names instead of rooting its namespace in the global DNS and querying the global DNS for resolution of FQDNs. This section applies to any organization that uses a private TLD, not just the ones that are already leaking name queries to the global Internet. If your organization is using what you perceive to be a "safe" private TLD, that is, a name that is not yet applied for or approved to be delegated into the global DNS root, you should still seriously consider changing to a name rooted in the global DNS. If you work in a large organization with more than one private TLD (such as a company that has merged with another company and has not merged its two namespaces), the steps in this section must be performed for each private TLD.

Chances are that when the organization chose to use a private TLD, it did so with a particular naming convention in mind. The steps here may be in conflict with that original model. In order to mitigate reliably the problems associated with name collisions due to private TLDs, users and systems both need to change the way that they use domain names, and local nameservers need to be reconfigured in a way that some users may find inconvenient. Use the explanations of the unintended or undesirable consequences that can affect your organization in order to raise awareness and foster acceptance among your user community.

Important note: At the same time as you are performing the steps in this section, you will probably also need to mitigate for name collisions caused by search lists, which is covered in Section 5. Many of the steps in that section are the same as these, and can be performed at the same time.

4.1. Monitor the requests coming into the authoritative nameservers

In order to mitigate problems with a private TLD, list all of the computers, network equipment, and any other system that uses the current private TLD in any requests. When you change the names being used, all devices that use the old private names in an automated fashion will need to be updated.

There are three common ways to perform this monitoring and enumeration of systems:

- The authoritative nameserver (such as Active Directory) may have a logging feature. Turn on the logging feature to gather details of all queries for the private names.

- Many modern firewalls can also be configured to detect and log queries for private names. This may not be as effective as logging from the naming system itself, depending on the topology of your network. For example, if a query does not go through a firewall, the firewall cannot see the query, and it will thus be missed.

- If neither of the above can be used, monitor and collect traffic delivered to and emitted by the authoritative nameserver using a packet capture program such as Wireshark. However, this method requires that the captured data be processed with a program in order to find the queries for just the private names.

Some organizations will (and should) choose to do more than one of the above to increase the chances of finding all the requests. Note that this step can produce confusing results. Devices such as computers and phones have applications into which users type names; those devices will show up in the survey even though there might not be any stored versions of the old private names. For this step, it is only necessary to know all the places in your network where the old private name is being stored and used for applications.

4.2. Create an inventory of each system using the private TLD in an automated fashion

You need a summary of the log data obtained from the previous step. That summary should be a list of all the devices and all of the names being queried rather than every instance of the device making a query. The reason you need all the names that are being queried is that some devices will have multiple applications that will each need to be fixed. Thus, the summary must include both all the systems and all the applications on each system that use the private TLD. This summary becomes the manifest for devices that need to be changed.

4.3. Determine where your global DNS names are administered

It is likely that you already have a global DNS name for your organization and that the domain name can be used for the root of your private namespace. You need to determine who is in charge of your DNS names and what processes they use to create and update names in the DNS. This may be done within your IT department, or it may be done through a service provider (often the same company as the one from which you get your Internet connectivity).

4.4. Change the root of your private namespace to use a name from the global DNS

A common strategy for using a global DNS name as the root of your private namespace is to have a publicly accessible name delegated from the global DNS but then use your existing authoritative nameserver to administer all the names below that. For example, if your company has the global domain name `ourcompany.com`, you might choose `ad1.ourcompany.com` as the root name.

If your organization has more than one domain name in the global DNS, you should root your names under one that can be most easily controlled by the IT staff in your organization. In some cases, additional names are controlled by other entities, such as a marketing department. If possible, it is best to root your name under a name over which the IT organization already has control.

The steps for making this change depend on which private nameserver software you have, the specific version of that software, the topology of the nameservers on your private network, and the existing configuration of the nameserver. These details are beyond the scope of this document, but should be covered in the instructions from your vendor for your current system. Also, in many organizations, this change will require authorization from some levels of management, particularly if the management of the global DNS names is different from the management of the private namespace.

As part of this step, if you have certificates for any hosts using names in the private namespace, you need to create certificates for those hosts using the new (fully qualified) names. The steps for getting these certificates depend on your CA and are thus also beyond the scope of this document.

4.5. Allocate new IP addresses for hosts, if needed

If you have TLS certificates that are based in your old private TLD name, you will need to get new certificates for the new names. If your web server does not support the Server Name Indication (SNI) extension to TLS which allows more than one domain name to be served under TLS on the same IP address, you will need to add IP addresses to the hosts so that the host supports the old private name on the original IP address and the new name on a new IP address. Alternatively, you can update your web server software to a version that handles SNI extensions correctly.

4.6. Create a system for monitoring equivalence between the new and old private names

When you change all the private names to use the new root, you will continue to serve addresses and to log queries for your old private names in order to check for systems not in your inventory that were not updated to use the DNS-rooted names. Because of this, you need to be sure that the new and old private names have the same values for the IP addresses.

Some private namespace software allows you to keep the two trees in parallel, but if you have older software or multiple authoritative nameservers, it is likely you will have to do monitoring for equivalence using custom tools. These custom tools need to query all the names in both the old and name namespaces often, and alert you if there is a mismatch so that you can determine which system changed without a parallel change in the other system.

If you needed to add IP addresses in the previous step due to having SSL/TLS certificates, the mismatch needs to be allowed by the equivalence monitoring software.

4.7. Train users and system administrators to use the new name

In addition to changing the systems where names are entered in configurations, you need to change the ways users think in order to get them to change from the old private names to the new ones. This training should be done before implementing the following steps so that users have a chance to get accustomed to the new names, but the training should make it clear that the change is coming and that they should start thinking in terms of the new names soon. This is also a good time to train users about using FQDNs. Use the explanations of the unintended or undesirable consequences that can affect your organization in order to raise awareness and foster acceptance.

4.8. Change every affected system over to the new names

This is the point where the migration from the old private names to the new ones becomes real for all the systems (PCs, network devices, printers, and so on) on the network. The private names are replaced by the new DNS names on a system-by-system basis. Every instance of the old private name is found in all the software on the system and replaced with the new DNS name. At the same time, you should deprecate the use of short unqualified names in search lists.

The monitoring that was started above is exceptionally important in this step. You are unlikely to be able to determine all of the applications in all the systems that have the old private names embedded in them. Instead, the monitoring system needs to be consulted after each system is changed to see whether that system is still making requests for the old private names.

Many systems run some initialization applications when they are first turned on. These applications may have system names embedded in them, and finding all of these can be difficult. After changing all of the names in a system from the old private names to the new DNS names, reboot the system and use the monitoring software to watch for name lookups. If the system is looking for any of the old private names, you need to determine which software is causing that request and change it to use the new names. This process might take a few reboots in order to fully configure a system correctly.

4.9. Begin monitoring for use of old private names at the nameserver

You should configure your authoritative nameserver to start monitoring all requests for names that have the old root. Because your users should not be using these names any more, the log created by this monitoring step may not be very large; if it is, you will have to repeat some of the steps above for particular systems on your network.

4.10. Set up long-term monitoring at perimeters to watch for old private names

The previous steps should have found the vast majority of uses of the old private names, but a few (possibly key) systems may still be using the old private names, but perhaps only rarely. One way to detect these name queries is to add rules to all firewalls at the edge of your network to look for any requests that are leaking. These rules should have a high priority associated with them and should be configured to generate event notifications so that IT staff is promptly alerted. You could instead find these events in the firewall logs, but doing so has a higher chance of being missed. Alerts that are triggered when requests occur will allow staff to detect these now hopefully rare events. Some firewalls only support this type of rule by adding additional features at additional cost; if that is true for your firewall, you need to assess whether or not the benefit of finding stray requests is worth the additional cost.

4.11. Change all names from the old root to point to a non-functioning address

After the users have been trained, the most effective way to be sure that they stop using the old private names before removing them is to have all the old private names point to a server that you have configured to not respond to service requests of any kind. This also helps flush out any systems that are still using the old namespace but were not detected in the earlier steps.

The address pointed to should be a server that is guaranteed not to be running any services. By doing this, there is no chance that any system using an old private name gets erroneous information and that applications will report errors that should be easily detectable or understood by users; as part of the awareness training, you can recommend that users report all errors of this kind to IT staff. As this step is implemented, the monitoring system that is checking for equivalence between the old and new names (described above) needs to be kept up to date with the changes.

The names should be changed one at a time, probably with at least a few hours between each change or batch of changes. This step is likely to cause calls to the IT department, so staging the changes will help balance the load of calls as names that were still in use begin to stop working.

4.12. If certificates were issued for any hosts under the old private names, revoke them

If your organization had SSL/TLS certificates issued for any servers in your network using the old private names, those certificates should be revoked. This is fairly easy to do if your organization acts as its own CA. If you used a commercial CA to issue certificates for the private namespace, you need to determine that CA’s process for requesting revocation; different CAs might have different requirements for such requests.

4.13. Long Term Operations with the New Name

Note that the old private name and domains beneath it are still being served, and will continue to be served for as long as you run the nameserver. There is no reason to remove them, and in many systems such as Active Directory, it can be difficult to remove the first name that was configured in the system.

There is actually a good reason to leave the name there: this allows you to see if there are any residual traces of the old private name in systems on your network. As long as all the addresses associated with all the names under that private TLD point to a host with no services running, you can use both the logs from the nameserver (and, for extra benefit, a system logging all the traffic to that server) to determine how thorough you were in removing the old private name.

5. Steps to Mitigate Name Collisions Associated with Search Lists

In order to mitigate reliably the problems associated with name collisions due to search lists, users and systems need to change the way that they use domain names. It can be helpful to prepare users in advance by way of change notifications, awareness programs, and training.

Note that if you are already doing centralized administration, these actions are probably less difficult than you might think. Many people who normally use search lists know that they can also type full names if needed (such as if they are accessing a server from outside the organization's private network), and they will need less training than those who only understand the short unqualified names.

5.1. Monitor the requests coming into the nameserver

In order to mitigate problems caused by search lists, you need to know all of the computers, network equipment, and any other system that uses search lists in any request. All devices that use search lists in an automated fashion will need to be updated.

There are three common ways to perform this monitoring and enumeration of systems:

- The recursive nameserver (such as Active Directory) may have a logging feature, and you can turn on the logging feature to obtain details of all queries that have short unqualified names.

- Many modern firewalls can also be configured to detect and log queries of all names. This may not be as effective as logging from the naming system itself, depending on the topology of your network. For example, if a query does not go through a firewall, the firewall cannot see the query and it will thus be missed.

- If neither of the above can be used, the nameserver can be monitored using a packet capture program such as Wireshark. However, this method requires that the captured data be processed with a program in order to find the queries for just the short unqualified names.

Note that this step can produce confusing results. Devices such as computers and phones may have applications into which users type names; those devices will show up in the survey even though there might not be any stored versions of the short unqualified names. For this step, it is only necessary to know all the places in your network where a short unqualified name is being stored or used for applications.

5.2. Create an inventory of each system using short unqualified names in an automated fashion

You need a summary of the logs from the previous step. That summary should be a list of all of the devices and all of the short unqualified names being queried rather than every instance of the device making a query. The reason you need all the names that are being queried is that some devices will have multiple applications that will need to be fixed. This summary becomes the manifest for devices that need to be changed.

5.3. Train users and system administrators in using FQDNs

In addition to changing the systems where short unqualified names are entered in any configuration (either a system-wide configuration or the configuration for an individual application), you need to change the ways users think to get them to change from using shortened names to full ones Use explanations of the unintended or undesirable consequences that can affect your organization in order to raise awareness and foster acceptance.

5.4. Change every affected system over to FQDN use

Replace the short unqualified names with their equivalent FQDN on a system-by-system basis. Every instance of a short unqualified name that is found in all the software on the system needs to be replaced with the full domain name.

The monitoring that was started above is exceptionally important in this step. You are unlikely to be able to determine all of the applications in all the systems being changed that have short unqualified names embedded in them. Instead, the monitoring system needs to be consulted after each system is changed to see whether that system is still making requests for the short unqualified names.

Many systems run some initialization applications when they are first turned on. These applications might have system names that rely on search lists embedded in them, and finding all of these can be difficult. After changing all of the names in a system to use FQDNs, reboot the system and use the monitoring software to watch for name lookups. If the system is looking for any short unqualified names, you need to determine which software is causing that request and change it to use FQDNs. This process might take a few reboots in order to fully configure a system correctly.

5.5. Turn off search lists at shared name resolvers

This is the point where the migration away from short unqualified names becomes real for all the systems (PCs, network devices, printers, and so on) on the network. Search lists can exist in any system that does name resolution or that serves configuration to other systems, such as a DHCP server. These systems are often stand-alone nameservers, but they can also be firewalls or other network devices. Regardless of the type of system, search lists need to be turned off on each of them in order to prevent users from trying to use short unqualified names within a given namespace.

5.6. Begin monitoring for use of short unqualified names at the nameservers

You should configure your nameserver to start monitoring all requests for names that need to use search lists. If you provide advance notice and training, your users should not be using these names any more, so the log created by this monitoring step may not be very large; if it is, you may need to repeat some of the steps above for particular systems on your network.

5.7. Set up long-term monitoring at perimeters to watch for short unqualified names

The previous steps should have found the vast majority of uses of the old private names, but a few (possibly key) systems may still be using short unqualified names, although perhaps only rarely. The best way to detect these name queries is to add rules to all firewalls at the edge of your network to look for any requests that are leaking. These rules should have a high priority associated with them and should be configured to generate event notifications so that so that the IT staff is promptly alerted. You could instead find these events in the firewall logs, but doing so has a higher chance of being missed. Alerts that are triggered when requests occur will allow staff to detect these now hopefully rare events. Some firewalls only support this type of rule by adding additional features at additional cost; if that is true for your firewall, you need to assess whether or not the benefit of finding stray requests is worth the additional cost.

6. Detecting Name Collisions in New gTLDs

Beginning August 18, 2014, ICANN is requiring that gTLDs that are newly delegated in the root zone assist organizations to detect when they are leaking queries to the global DNS for names falling under the new TLD. This assistance will last for 90 days, most likely, the first days that the new gTLD is in the root zone; after that, the new gTLD will act like any other TLD in the root zone. The assistance is provided through a "controlled interruption" service that is described in this section.

Clearly, an organization that needs to mitigate name collisions between its private namespace and the global DNS should do so before the corresponding new TLD enters the root zone: it should not wait for this 90-day period. (This is particularly true for organizations that chose a two-letter TLD for their name, because these names are not required to perform a controlled interruption.) Controlled interruptions are meant as a last warning to an organization that it needs to quickly perform mitigation before the TLD starts giving "real" answers to queries.

This section describes how a controlled interruption is implemented on an authoritative name server, and how it appears in the answers to queries. It also gives advice to organizations that have private namespaces to determine whether operational changes that they are observing are due to controlled interruption and, if so, what to do about those changes.

6.1 Description of Controlled Interruptions

The controlled interruption service that is being required by ICANN for new gTLDs added to the root zone after August 18, 2014 is designed to cause an interruption to devices whose requests for domain names in private namespaces that leak into the global DNS. Currently, when such a DNS request leaks into the global DNS, the root name servers send back a response with a code that indicates that the domain does not exist. (Technically, this is the RCODE field of the response's header being set to a value of 3, mnemonically defined to be an "NXDOMAIN" response.)

During the controlled domain service period, instead of an NXDOMAIN error in the response, the response contains no error indication of an error, but instead contains data that has the highest chance of being noticed by the system that sent the request. It is impossible to design a response that will always be noticed because there are so many different types of software that make DNS requests; however, the ICANN-mandated controlled interruption will be observable on systems with adequate logging of errors, and on networks where DNS traffic can be observed by network administrators.

gTLDs operating in the controlled interruption fashion will respond to a wide variety of DNS queries in a predictable fashion. Section 6.2 explains how to observe the behaviors of systems that get controlled interruption responses to these DNS queries.

- By far the most common DNS query is for A records, that is, for the IPv4 address(es) associated with a domain name. Those queries will always come back with the IPv4 address of 127.0.53.53. This address is a loopback address for the host that sent the query, so that if the application uses that address to initiate any kind of contact, it will send the message to itself. This, of course, is likely to fail, since almost all programs that are making DNS lookups intend to use the address in the response to contact another server.
- Another common DNS query is for records that contain text, commonly known as "TXT records". In the controlled interruption service, the TXT record response will always be the exact string "Your DNS configuration needs immediate attention see https://icann.org/namecollision". A system that displays such text records gives the viewer information on name collisions.

- For DNS queries that are for mail servers (technically, for mail exchanger or MX records), the controlled interruption service will respond with the domain name your-dns-needs-immediate-attention.<TLD>, where "<TLD>" is the TLD in the DNS request. This domain name may be visible in error responses from the mail client or mail server. Looking up the address of the domain name your-dns-needs-immediate-attention.<TLD> will return 127.0.53.53.

- The controlled interruption service will respond to queries for service (SRV) records with the domain name your-dns-needs-immediate-attention.<TLD>. Queries for SRV records are not as common as those for IPv4 addresses, text records, and mail server names, but are becoming more common for newer applications such as instant messaging and voice transmission.

A gTLD added to the root zone before August 18, 2014 may also have a controlled interruption service for a subset of the possible second-level domains in the TLD. The records returned in the controlled interruption for these names are identical to the records described above. ICANN required that some SLDs be blocked from the TLD, and those names might become active soon after a 90-day controlled interruption for the SLDs

6.2 Observing Controlled Interruptions

It is important to note that there is no guarantee that an application that receives a controlled interruption response will act visibly different than it did before the controlled interruption. However, it is highly likely that the application will behave differently, and the difference will most likely be a failure; hopefully that failure will have an error messages associated with it and the application user will report this to a system administrator who is tasked to handle this. If the error message contains the IPv4 address 127.0.53.53, that is a very strong indication that the error is due to the program using a name from a private namespace that leaked to the public Internet.

Errors due to the controlled interruption service appear when a program that was earlier getting NXDOMAIN responses to queries starts getting actual responses. Of course, these errors would appear later when the new gTLD was responding with real data, and the controlled interruption service is likely to last only the 90 days mandated by ICANN. During that time, the errors will be more obvious because error messages contain the IPv4 address of 127.0.53.53, the text "Your DNS configuration needs immediate attention see https://icann.org/namecollision", or a domain name containing "your-dns-needs-immediate-attention".

Controlled interruptions can also be observed on an organization's network if the network administrator is actively searching for DNS messages that contain those responses. Such searching can be done through a network tap at appropriate ingress points, or can be done on a firewall. This type of observation does not rely on seeing error messages in the affected computer; instead, the network administrator can determine which computer whose requests for names in private namespaces are leaking from the organization's network.

Regardless of how the controlled interruption is discovered, the result should be that the computer getting the controlled interruption response should be reconfigured to only make DNS queries to the organization's nameserver, not to the global DNS. There is no standard way to specify such a setting, although the setting is normally part of the operating system. If the computer gets its network settings from a server in the organization's network, commonly called a "DHCP server", that server needs its settings changed to have DNS queries go to the organization's nameserver, not to the global DNS.

Any observation of a computer getting a controlled interruption response is a sign that other computers on that organization's network might be getting them as well. A system administrator should immediately check the DNS settings for all computers on the same network, even if those computers are not showing observable signs of getting controlled interruption answers. Remember that the controlled interruption only last 90 days, so there is limited time to find computers that have incorrect DNS settings.

Of course, making such changes is only a temporary mitigation for the underlying problem of name collisions. Sections 4 and 5 of this document give instructions on how to make permanent mitigations.

7. Summary

Name collisions have the potential to create unanticipated results for organizations that use private namespaces. This document lists some of those potential results and specifies best practices for changing the way that private namespaces are used within organizations. The document also describes controlled interruption as a means to identify where the effect of name collisions might become apparent.

For namespaces that used a private TLD that is becoming (or is already) a TLD in the global DNS, mitigation best comes in the form of migrating the namespace to a namespace that is rooted in the global DNS. For namespaces that use name shortening with search lists, mitigation can come only by eliminating the use of search lists. Steps to achieve these mitigations also include long-term monitoring in the private network to be sure that all instances of names that might cause collisions are no longer being used. There will be means for organizations to tell when they are going to experience name collisions as some new TLDs are delegated in the root zone.

The comprehensive mitigation for the problems of name collisions is to use FQDNs in all places where a domain name is used. In a network that is already using the global DNS, this means not using search lists. In a network that uses a private namespace, this means that the private namespace should be rooted in the global DNS, and should not be using search lists.

Appendix A: Further Reading

The following documents were produced by various organizations within ICANN. Other organizations provide documents that might be useful as well. Most significantly, the vendor of your nameserver software and/or hardware may have valuable information on their tech support web site.

A.1. Introduction to the New gTLD Program

This page describes the history, implementation, and progression of the program to add hundreds of new gTLDs to the global DNS.
http://newgtlds.icann.org/en/about/program

A.2. Name Collision in the DNS

ICANN commissioned Interisle Consulting Group, LLC, to create this in-depth report about potential name collisions. It gives an overview of name collisions, presents data on currently non-existent TLDs that are currently queried at the root servers, and gives a great deal of background on the problems that name collisions might present.
http://www.icann.org/en/about/staff/security/ssr/new-gtld-collision-mitigation-05aug13-en.pdf

A.3. New gTLD Collision Occurrence Management Plan

This is the plan adopted by ICANN on how to manage name collision occurrences between new gTLDs and private namespaces. It also includes many pointers to comments received by ICANN to earlier proposals that relate to name collisions in the root zone.
http://www.icann.org/en/groups/board/documents/resolutions-new-gtld-annex-1-07oct13-en.pdf

A.4. Name Collision Occurrence Management Framework

This document component of the New gTLD Collision Occurrence Management Plan. It defines the specifics of the controlled interruption service for gTLDs that are delegated in the DNS root zone beginning August 18, 2014.
http://www.icann.org/en/system/files/files/name-collision-framework-30jul14-en.pdf

A.5. New gTLD Concerns: Dotless Names and Name Collisions

Search lists on different systems can give very different results depending on what is in the short unqualified name that is being queried. This article focuses on search lists for dotless domains (TLDs that have address records at their apex), but the description of search list processing is valuable in many other contexts as well.
https://labs.ripe.net/Members/gih/dotless-names

A.6. SAC 045: Invalid Top-level Domain Queries at the Root Level of the Domain Name System

This ICANN SSAC report describes the types of queries for TLDs that were seen in by root servers at the time of writing.
http://www.icann.org/en/groups/ssac/documents/sac-045-en.pdf

A.7. SAC 057: SSAC Advisory on Internal Name Certificates

This ICANN SSAC report describes the security and stability implications for certificates that contain private (internal) names. It identifies a practice by CAs that can be exploited by attackers and could pose a significant risk to the privacy and integrity of secure Internet communications.
http://www.icann.org/en/groups/ssac/documents/sac-057-en.pdf