The Environment Surrounding DNS

**DNS is used in many applications, serving as an important Internet service.
Here we discuss name collision issues that have arisen with recent TLD additions,
and examine the latest trends in the environment surrounding DNS.**

3.1 The Latest DNS Trends

DNS is a service that returns records in response to queries, and is mainly used for name resolution when looking up the IP addresses that correspond to domain names. Most applications on the Internet use name resolution via DNS, so it is a very important service. DNS involves authoritative servers for each zone that store records corresponding to domain names, and clients that make queries. In most cases, clients ask for DNS cache servers maintained by ISPs, etc., to perform laborious DNS recursive lookups, and only receive the results. DNS cache servers only know the IP address of the authoritative server that provides top-level zone information, which is known as the root. Based on the information gained from there, they track down authoritative servers likely to have more detailed information to locate the required record. Also, because server load and latency become an issue when making recursive queries each time, they cache the records obtained for a while, and retrieve records from the cache if the same query is received again. Recently, functions related to DNS have also begun to be implemented in devices on the communication route, such as broadband routers or firewalls. These may be involved in relaying DNS queries or applying control policies.

A governing organization known as a registry is specified and managed for each top-level domain (TLD), to prevent duplicate domain name spaces being registered. The root zone information is managed by ICANN, and from here authority is passed on to each TLD registry for the processing of domain name registrations from registrants. When registering a new domain name, an application is made to the TLD registry for the domain name you want to register through an intermediary called a registrar. However, each top-level domain and its associated subdomains have their own registration policies, and the details of who can register domain names for what purposes may vary. For example, .jp domains are governed by Japan Registry Services Co., Ltd. (JPRS), and any individual or company with an address in Japan that can be contacted is able to register the second-level domain names known as general-use jp. However, only companies and organizations registered in Japan can register domain names with a .co.jp suffix. Some TLDs are also run without any restrictions set, under a policy that allows anyone to register domain names. Because the management of registered domain names is delegated to the registrant, each registrant must have their own operation policy to manage and operate their domain appropriately.

3.2 Name Collision Issues

Systems introduced for the sake of convenience with the idea that some form of action is better than nothing can create problems down the line. Name collision issues are an example of this. For instance, in environments with clearly defined management that are only used by certain people, such as companies or homes, independent internal domain name spaces with a non-existent TLD (private TLD) may be used. For small scale operations, host names may be registered directly on clients using a hosts file, etc. When larger numbers of clients are involved, the DNS may be configured to respond to private TLD queries from an internal cache server or firewall, making access using an internal domain name possible without requiring any changes to client settings in particular. A system may appear to be working according to plan when configured, but the Internet is continually evolving, and cracks can start to appear when standard technology is used under non-standard configurations. Over 300 new TLD have already been added to the root zone in recent years, and that number continues to grow. If a private TLD configured internally for convenience conflicts with an added TLD, it results in issues such as no longer being able to use legitimately registered domain names, or unintentional connection to sites. This is called name collision issues (Figure 1). To avoid these issues it is important to preserve uniqueness for domain names, even if they are only for internal use. When you already have a domain name registered, you can be secure in the knowledge that the uniqueness of the domain name you use can be maintained in the future by configuring a subdomain for internal use under it, or preferably registering a new domain name for internal use.

Name collision issues also affect digital certificates for servers. Public certificate authorities that issue digital certificates have issued them even for private TLD domain names for internal use, so that digital certificates can also be used on the internal servers of organizations. Digital certificates for standard servers enable the ownership of domain names to be confirmed through the registration of specific character strings to email or websites. However, this kind of confirmation cannot be performed with private TLDs, so it was possible to obtain digital certificates without any confirmation in particular. As a result, when a domain name was registered to a newly added TLD, someone could have obtained a digital certificate for the corresponding domain name in the past. In response to these name collision issues, the CA/Browser Forum, a private organization involved with digital certificates, developed operational standards to gradually limit internal domain name digital certificates in the future. The digital certificates already issued for internal domain names are set to expire no later than November 1, 2015. When a new TLD is added, associated digital certificates are revoked within 120 days. There are also plans to revoke all digital certificates for internal domain names using a private TLD or domain names that cannot be confirmed to exist from the Internet, including those previously issued, in October 2016.

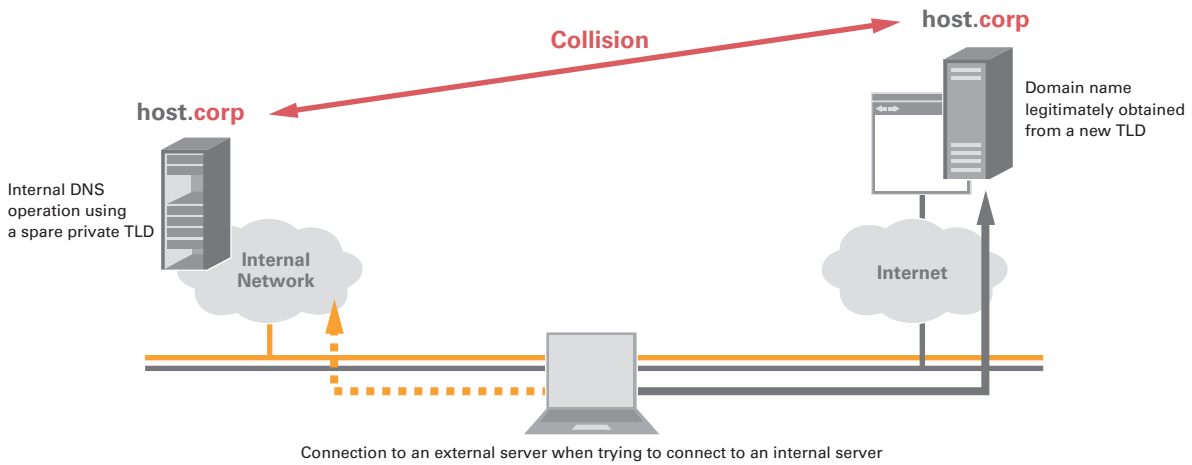

Figure 1: Name Collision Examples

Clients sometimes implement functions such as search lists and DNS suffixes for supplementing the domain name during DNS name resolution. These functions enable users to specify the intended domain name by entering the first part, rather than the fully qualified domain name (FQDN). Within organizations the domain name part is usually shared, so this function is used to enable connection to multiple servers or devices without having to repeatedly enter the same domain name. However, this is another area in which name collision issues crop up. Although largely dependent on the client implementation, when a domain name containing at least one "." is specified as the destination, it seems that in many cases the domain name is first queried via DNS, and then queried again with the addition of the search list domain name if no record exists. Consequently, when name resolution using the intended domain name was previously successful after the domain name was supplemented because a domain name corresponding to the initial query did not exist, there is a possibility of unintentionally connecting to the wrong site when a response to the initial query is obtained due to subsequent changes in the environment. Newly registered TLDs include regional names such as tokyo and nyc, so it is necessary to take particular care when using these regional names for subdomains. It goes without saying that thought must be given to the overall impact, considering all the TLD added, the operation of internal subdomains, and the DNS search lists distributed via DHCP, etc. To avoid trouble in the future, it would be best to instruct users to connect using fully qualified domain names whenever possible.

ICANN recognized the issues with name collision from an early stage, and they have taken a range of measures to lessen the impact. The CA/Browser Forum operational standards mentioned previously are the result of discussions with ICANN, and they have evaluated the risks associated with implementing new TLD based on actual DNS query status. Their investigation used data from major authoritative servers collected mainly during DNS-OARC's A Day in the Life of the Internet (DITL) project*1. Internal domain names are designed to be used within organizations, but DNS queries leak when there are configuration errors, or when mobile devices attempt to connect to servers from an external network. This can also be detected on root servers, allowing private TLD usage to be estimated. In this survey it was discovered that queries using "home" and "corp" private TLDs were remarkably frequent. Because there would be too many problems if these were recognized as new TLDs, it was decided that they should remain undelegated indefinitely*2. Other new TLDs can be delegated, with the caveat that registering certain second-level domain names that are highly likely to cause name collision is prohibited, based on their frequency of appearance in data such as DITL.

In Japan, JPNIC established a team of experts to evaluate risk and recommend strategies for the launch of a large number of gTLDs. They looked into name collision issues, and summarized their recommendations in writing*3. Because the impact of name collision issues can show up in unexpected places, I recommend a thorough review to check whether there are actually any dangers present. This should also cover areas that have functioned without issue before, including internal settings and URLs listed in documents.

*1 https://www.dns-oarc.net/oarc/data/ditl
*2 https://www.icann.org/en/system/files/files/resolutions-new-gtld-annex-1-07oct13-en.pdf
*3 https://www.nic.ad.jp/ja/dom/new-gtld/name-collision/

3.3 DNS and Communication Control

DNS has the capability to control and monitor the communications of users, such as redirection to websites and control of mail delivery. In other words, it is possible to control client communications by manipulating DNS responses (Figure 2). For example, providers that distribute content on a worldwide scale may have an implementation that responds to DNS queries from clients with the IP address of a delivery server close to the client, to reduce latency and optimize delivery. Because many users actually use the DNS cache servers provided by ISPs, etc., it seems the authoritative servers managed by content providers group users by DNS cache server, and respond with the IP address thought to be most suitable. Meanwhile, there have also been cases in which this has been exploited by attackers. Attackers often cause users to reference DNS cache servers they manage, resulting in the download of malicious content. In DNS Changer cases, it was reported that the DNS lookup address for a device was overwritten, and the DNS lookup address of broadband routers was changed to one controlled by the attacker. To avoid being redirected to malicious data like this, the utmost care must be taken with DNS settings, but the environment surrounding DNS is becoming more complicated.

Currently, most devices set the DNS cache server to look up based on DHCP information. In some cases the DHCP function is consciously activated by an administrator within an organization, while for consumer usage it may be provided as standard via a broadband router. Many broadband routers implement functions for simply relaying DNS queries to the DNS cache servers of ISPs, etc., and configure the router itself as the DNS lookup address for devices at home. However, with regard to DNS specifications, it seems these implementations sometimes only feature very limited functionality. They may lack support for queries over TCP, or not be fully compatible with EDNS0. In light of this situation, guidelines for implementing DNS relay functions in devices such as broadband routers were published as RFC5625/BCP152*4. These guidelines recommend that DNS relay functions be implemented with a focus on transparency, so they can continue to be used without issue in the future.

Some broadband router models forcibly send all DNS queries that pass through them on to the DNS cache server set on the router, regardless of the device's DNS lookup address settings. In this case, it is not possible to determine which DNS cache server is referenced simply by looking at the DNS lookup address settings on the device, no matter what IP address

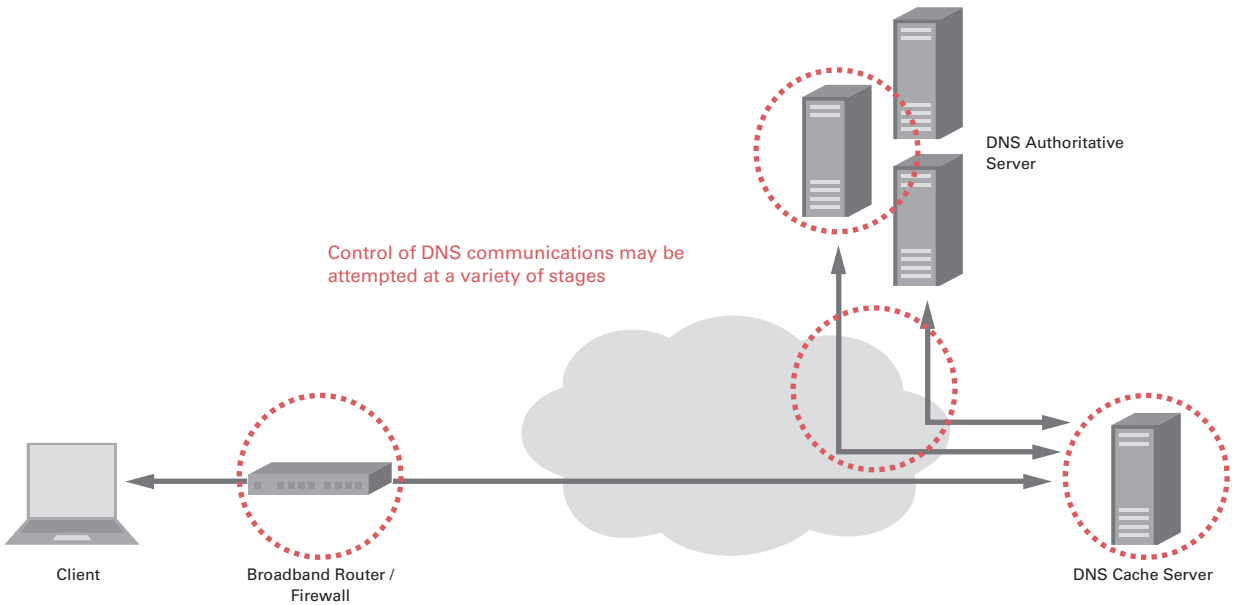

Figure 2: DNS and Communication Control

*4 https://tools.ietf.org/html/rfc5625

is registered to it. That is because as soon as a DNS query from the device passes over the broadband router, it is rewritten as a query to the specified DNS cache server. This may be a convenient function for implementing operation policies that enforce specific DNS cache servers, but it is harder for users to notice when a problem occurs, making isolation very difficult.

A variety of controls are also being implemented on ISP DNS cache servers. When devices connected to a walled garden IPv6 network communicate with servers on the Internet side via IPv4 alone, the IPv6 connection may fail, causing delays or connection faults due to IPv6/IPv4 fallback. Communication is also possible via IPv6 as long as an Internet connection exists, but because this requires an application and router update, there are concerns over the time it takes to implement. To alleviate problems in the user environment such as these, DNS cache servers may implement a function called an AAAA filter that gives no response for the AAAA records used with IPv6. Additionally, to prevent the circulation of child pornography, child pornography blocking*5 is sometimes implemented on DNS cache servers. This blocks responses for specific domain names.

In some countries and regions, DNS control has been introduced for the government-led blocking of access to undesirable content. Blocking is apparently carried out via the DNS cache servers provided by each ISP. Queries to specific domain names are also sometimes blocked, or the responses altered to contain falsified information, by implementing functions that monitor DNS queries on a network. Because users cannot determine whether issues are caused by communication faults or intentional blocking in most cases, and with blocking policies seldom being disclosed, it is tricky to figure out whether there is an issue and isolate the problem. However, as there are patterns in the content blocked depending on the country or region, it is possible to estimate whether intentional blocking is likely to be involved. This demonstrates that DNS control can be implemented in a variety of areas, and these must be considered when isolating an issue to pinpoint the obstruction.

3.4 DNS and Attacks

Because DNS is implemented on many devices, and many applications are practically reliant on it, it is sometimes exploited or targeted by attacks. DNS amplification attacks, which use DNS cache servers called open resolvers that respond to queries from anyone as stepping stones, have come to be particularly widely exploited due to their amplification efficiency and distributed nature. Many DNS cache servers were left to respond to queries from everyone without any countermeasures taken, so they were used as stepping stones. The configuration of DNS cache servers for ISPs in Japan has changed gradually in recent years, and they now usually only accept DNS queries from their users. However, some DNS relay functions implemented on broadband routers respond to DNS queries from the Internet without restriction by default, allowing them to be exploited as stepping stones in attacks. These must be dealt with individually by users, so continued reminders for users are required.

Attacks and the exploitation of DNS authoritative servers are also occurring. As with regular DDoS attacks, these aim to disrupt service by flooding authoritative servers with large volumes of traffic. There have also been cases in which authoritative servers have been exploited as stepping stones in DNS amplification attacks by directing large volumes of queries that are legitimate under DNS protocol at them. Straightforward traffic floods can be prevented through use of an appropriate packet filter. However when a server used as a stepping stone in an amplification attack, it is not easy to distinguish the queries intended as attacks, so the response must be considered. A number of authoritative DNS, including JP DNS, are making an effort to reduce the impact by implementing a function known as Response Rate Limiting (RRL), which limits a consecutive series of identical responses*6. However, this countermeasure is not foolproof, so it is necessary to continue scrutinizing attack techniques and searching for an appropriate response.

*5 http://www.netsafety.or.jp/blocking/
*6 http://www.redbarn.org/dns/ratelimits

Since the beginning of 2014, we have intermittently observed large volumes of DNS queries involving a number of domain names on ISP DNS cache servers. The purpose of these is not known, but our guess is that they are likely to be distributed attacks on the authoritative servers for the corresponding domain. However, because a large volume of communications with the corresponding authoritative servers is generated along with these attacks, even the DNS cache servers are overloaded, and at times faults such as delays involving name resolution for DNS cache server users occur. Regardless of whether or not this was the attacker's intention, when users are impacted some form of action must be taken. That said, because the communications appear identical to regular DNS queries from users, with broadband routers acting as open resolvers used as stepping stones in some cases, and DNS queries from bots that have infected user PCs used in others, it is not easy to apply generic countermeasures. We must pay close attention to the format of queries that appear with abnormal frequency, and take measures on a case-by-case basis.

UDP is usually used as the protocol for DNS queries. UDP communications are easier to spoof than TCP, and attackers may be able to inject fraudulent responses. For the injection of fraudulent responses to succeed, the query and IP address, port number, DNS ID, and QNAME information must match. To defend against this, it is necessary to ensure that information corresponding to queries does not match fraudulent responses. Assuming that the DNS ID is already generated from sufficiently random numbers, the only remaining option is to use a well-randomized number for the outgoing port number. When using old DNS implementations with a fixed outgoing port number, it will be necessary to upgrade to the latest DNS implementation to be able to send out queries that are hard to calculate. Also, because some firewalls and NAPT devices overwrite information such as the DNS ID and outgoing port number even when you've gone to the trouble of randomizing them, caution is required.

Spoofing of the source IP address is often used in attacks related to DNS. Once BCP38*7 is implemented on each network, creating environments in which the source IP address cannot be spoofed, it will be possible to eradicate most of the current attacks that exploit DNS. Routers are also equipped with a uRPF check function to make BCP38 easier to implement, so we recommend seriously considering the introduction of source IP address validation to prevent attacks in which the source IP address is spoofed, particularly in networks connected to devices.

3.5 Conclusion

DNS is a crucial Internet service that many applications depend on. To keep it available for use in a healthy condition, the authoritative servers, clients and DNS cache servers that carry out name resolution, and other equipment that interacts with DNS must be managed and operated with suitable coordination. DNS name spaces are changing significantly after the recent addition of new TLDs. When using domain names internally via private TLDs, or name resolution dependent on search lists, name collision issues may occur. Also, because DNS is widely used as a control system, and control attempts may be made in a variety of areas, it is becoming more complex. This complexity itself can be a source of problems, and could interfere with name resolution, so care must be taken. Although not limited to DNS, attack techniques are changing due to an abundance of bandwidth and CPU resources. We recommend a close eye be kept on attack techniques as they develop, as it will be necessary to gather and share information on a daily basis to keep up with the times. IIJ would like to contribute to the development of a healthy Internet by operating our facilities appropriately, and taking part in information sharing and discussion as needed.

Author:

Yoshinobu Matsuzaki
Mr. Matsuzaki is a Senior Engineer in the Network Engineering Section of IIJ Network Service Department.

*7 http://tools.ietf.org/html/bcp84