Report on Root Zone Glue Handling

Andreas Gustafsson

Araneus Information Systems Oy

November 2009

1. Introduction

The Internet Assigned Numbers Authority (IANA) is responsible for maintaining the contents of the root zone of the Domain Name System (DNS). The root zone delegates the various top-level domains (TLDs) to their name servers by means of name server (NS) records and associated address records commonly referred to as *glue* records.

Maintaining the glue records presents particular challenges because they may be shared by more than one TLD. This is reflected in the current IANA procedures for processing changes to glue, which can involve obtaining authorization from multiple parties, and as a result have suffered from lengthy delays.

A number of alternative procedures aiming to reduce those delays have been proposed, including the use of automated mechanisms. This report examines existing and proposed procedures from a technical perspective, introduces new proposals, and makes recommendations for improvements.

1.1. Terminology

This report uses terminology established in the Internet Draft *draft-koch-dns-glue-clarifications-03* (Koch 2007). Additionally, the following terms are used:

Server name

A domain name referenced in an NS record. For example, given the delegation EXAMPLE. NS NS1.EXAMPLE., the domain name NS1.EXAMPLE is a server name. Although it is common to refer to the right-hand side of an NS record as a *name server*, this report uses the term *server name* to emphasize that it is referring to the domain name rather than the physical server machine (or machines) it represents. There is no simple relationship between server names and physical name server machines: a given server name can map to multiple IP addresses, and multiple server names can map to the same IP address. Also, multiple IP addresses can refer to the same physical machine, and through the use of anycast routing and load balancing, a single IP address can refer to multiple physical machines.

Resolver

An entity executing the DNS resolution algorithm of RFC 1034 section 5.3.3 (Mockapetris 1987), typically a component of a caching name server, a.k.a. a recursive name server. Not to be confused with a *stub resolver.*

1.2. DNS Stability Concerns

A primary goal of the root zone glue handling procedures should be to maintain the overall stability of the DNS. Stability concerns can be divided into three classes, namely concerns of integrity, availability, and efficiency:

Integrity - ensuring that no unauthorized party can gain control over DNS data

Availability - ensuring that the data within each TLD remains resolvable at all times

Efficiency - ensuring that resolution proceeds with a minimum of delay and network traffic

As a rule, concerns of integrity and availability should take precedence over those of efficiency.

1.3. Glue Policies

When parent domains other than the root delegate child domains, they generally cannot provide glue for every child server name, but only for those that fall within the parent domain (including its delegated subdomains). This inability stems from two technical restrictions: firstly, many authoritative servers restrict the data for a zone to consist of records below the zone apex, and secondly, even if authoritative servers sent glue records outside the delegating domain, most resolvers would consider them "out of bailiwick" and discard them for security reasons.

The root domain is a special case in that it is not affected by the these restrictions—because every DNS name is a subdomain of the root, arbitrary glue can be present in the root zone, and will be accepted by resolver bailiwick checks.

In addition to the technical restriction that glue must fall within the parent domain, many domain registries apply the additional restriction of only allowing registrants to register glue records for servers falling within the child domain. This is referred to as a *narrow* glue policy, whereas providing glue for any server name below the parent is referred to as a *wide* glue policy. (Koch 2007)

The root registry currently employs a wide glue policy, allowing glue for any TLD server name. In fact, the current policy goes even further than that, not only allowing wide glue but requiring it: at present, every TLD server name has at least one glue record in the root zone.

2. The Current Root Zone Glue Update Process

The procedure currently used by IANA to process updates to root glue records shares most of its steps with the procedures for other types of updates to the root zone and its associated registration records, such as the creation of a new TLD, or a change in the administrative control of a TLD.

The following is a somewhat simplified description of the steps involved in the case of a glue update.

1. A root zone update request is sent to IANA. Requests can be sent by anyone, but are typically sent by the TLD whose server is being renumbered, or in the case of shared server names, by one or more of the TLDs sharing the server name.
2. IANA checks that the request is well-formed and clear.
3. IANA checks the requested changes for conformance to applicable technical requirements (IANA 2009).
4. IANA obtains approval for the change from the administrative and technical contact of each TLD using the server name.
5. IANA obtains U.S. Government authorization for the change.
6. The change is implemented in the DNS root zone by VeriSign.

2.1. Delays in Updates to Shared Glue

The existing root glue update process has proven problematic particularly in the case of updates to glue records for server names shared by multiple TLDs. Because such updates currently require approval by both the technical and administrative contact of each TLD using the shared name, a large number of approvals have to be obtained, and this has led to lengthy delays in the processing of such update requests. In the last two years, there have been several cases where the fulfillment of a root glue update request took more than 30 days, and one case where it took 330 days.

2.2. The Prevalence of Shared Server Names

An analysis of the current root zone data[1] shows that out of the 1127 server names used by the TLDs, 992 are used by a single TLD, and the remaining 135 are shared among two or more TLDs. Most of these are shared by only two or three TLDs. The server name shared by the greatest number of TLDs is SUNIC.SUNET.SE, shared by 26 TLDs; if it were renumbered, updating the corresponding root zone glue would require the approval by a total of 52 contacts under the current procedures.

Of the 280 TLDs, 92 have no shared server names, and 46 have only shared server names; the remaining TLDs have a mix of both.

2.3. The Server Renumbering Process

The process of renumbering a name server must be carefully staged to avoid interruption of service. This report assumes the following order of events:

1. Authoritative DNS service is set up at the new IP address, for the same set of zones as at the old IP address, while continuing to provide service at the old IP address.
2. The authoritative address records at the server name, or the server names if the server is known by multiple names, are updated to the new IP address.
3. The corresponding glue records are updated.

[1] This and other measurements presented in this report are based on the state of the root zone as of October 29, 2009 (SOA serial number 2009102900).

4. A wait period is observed to ensure that the old IP address has expired from resolver caches. Since resolvers may have obtained the address either from glue or authoritative data, this should be at least the larger of the authoritative or glue TTL, plus the time needed to propagate the change to all authoritative servers. The SOA EXPIRY field can be used as a worst-case estimate of the latter.
5. Service at the old IP address is decommissioned.

Updating the glue before the authoritative address records would also work, but would be counter to the principle that glue should reflect the authoritative data and not vice versa. When a server is known by multiple names, this ordering applies to each name independently; the glue for one name can be updated as soon as its authoritative data has been updated, whether or not the authoritative address records for the other names have been updated.

2.4. Risks of Delayed Glue Updates

When glue updates are delayed, the length of time for which service needs to be provided at both the old and new IP address increases, and there is an increasing risk that at some point, the server operator will no longer be able to provide service at the old address. If that happens, it will impact resolution efficiency, as resolvers relying on the glue address will receive no response and will suffer a retransmission timeout before they retry using a different server. In the extreme, it could impact integrity, as the old IP address could eventually be reassigned to a third party, who could then run a server providing forged responses for any domain whose glue still points at the IP address in case.

Both of these risks are associated with delays in deleting an out-of-date glue record. Delays in adding a new glue record do not carry the same risks; a policy exploiting this asymmetry will be discussed in section 5.2.

2.5. The Need for Notification

It is not clear that the existing policy of requiring approval by the contacts of every TLD affected by a glue change is actually necessary. The following rationale for the current policy was given in (ICANN 2006):

> *Duty-of-care to TLD operators*
>
> *The current active confirmation seeks to ensure the TLD operator is fully aware of the renumbering of one of their authoritative name servers. This gives them the opportunity to identify possible configurations changes they need to make to ensure continuity of service. Changes to this role may need to consider the impact on TLD operators, and what IANAs duty of care is to inform TLDs of impending changes that impact them.*

I find this rationale flawed. The only party that is in a position to notify the TLDs of a renumbering in a timely and reliable manner is the server operator, and therefore the duty to do so should fall on the server operator, not on IANA.

Notification is not even necessary in many cases, because the renumbering of a shared server does not necessarily require configuration changes at the TLDs it serves. In the cases where configuration changes are needed, they typically need to take place early in the process, at a point where IANA is not yet involved. For example, if a TLD authorizes zone transfers using source IP addresses rather than cryptographic means, and the renumbering causes zone transfer requests to originate from a new IP address, the TLD may need to modify zone transfer access lists, but this needs to take place before the server operator makes the new IP address operational and publishes it as authoritative data, not at the time when the root glue is updated.

This is not to say that IANA cannot extend TLDs the courtesy of notifying them when a renumbering is known to have occurred, but since IANA may only discover the renumbering after the fact, this cannot be the primary mechanism by which TLDs are notified of such changes. Any such courtesy notification by IANA should be done fully independently of IANA's primary duty of maintaining the root zone, and should not be a reason to delay updates.

2.6. The Need for Approval

The current process not only requires notification of TLDs affected by changes to root glue, but requires their approval of the change.

It is important to distinguish between a TLD's possible disapproval of the renumbering of a server, and their possible disapproval of the corresponding update to the root glue. A TLD may have reasons to disapprove of a renumbering, but it is hard to imagine a valid reason why a TLD would disapprove of a root glue update while approving of the underlying renumbering. The growth of referral responses past 512 bytes might be one such case, but such changes are already blocked by technical checks.

If a TLD disapproves of the renumbering of one of its servers, its primary recourse should be to negotiate with the server operator to have the renumbering reversed or postponed. If no agreement can be reached, the TLD has the option of switching to a different operator. Either solution is preferable to delaying the root glue update, not only from the perspective of IANA's duty to keep the root zone up-to-date, but also from the TLD's own perspective, because delaying the root glue update will not be fully effective at stopping resolvers from using the new address given that some of them rely on the authoritative address record rather than the root glue.

3. The 2006 Public Comment Process

In December 2006, ICANN published an announcement seeking public comments on the DNS root zone glue policy (ICANN 2006). The announcement presented four proposals for consideration, including the baseline proposal of keeping the current policy, and solicited additional proposals.

The following sections will analyze the proposals given for consideration by ICANN as well as some of the proposals received in response to the announcement.

3.1. Reduce Acceptance Threshold

This proposal was described as follows in (ICANN 2006):

> *Instead of requiring every administrative and technical contact to approve, the threshold could be at a minimum the requesting administrative and technical contact, or be a lower figure than 100%. For illustration, a hypothetical criteria may be "The changes must be approved by the administrative and technical contacts of either two affected TLD operators, or 20% of affected TLD operators whichever is higher.*

I find the description of the hypothetical criteria somewhat unclear; it could be read as either *"IANA shall seek the approval of at least 20% of the affected TLDs, and those TLDs must approve the change unanimously"* or as *"IANA shall seek the approval of every affected TLD, and the change will be made if at least 20% approve, even if the remaining 80% all object"*. I assume the former was intended.

Clarity aside, the general principle of verifying the change with a sample of the affected TLDs rather than all of them seems like a reasonable approach.

3.2. Allow Changes with a Mandatory Advisement and Wait Period

> *Once the acceptance criteria is met, if it has not been accepted by 100% of affected parties, IANA can notify all administrative and technical contacts of the nature of the change and give them a fixed time (e.g. 30 days) to make necessary alterations before changes are made to the root zone.*

As discussed in section 2.5, IANA may certainly advise TLDs of changes to the root glue, but that should not be IANA's primary duty. Since configuration changes at TLDs need to be made in advance of the server renumbering (if at all), any wait period to allow for such changes should occur between the time the server operator notifies its served TLDs of an upcoming renumbering and the time of the renumbering, and is a matter between the TLD and the server operator. The root glue update should require no additional configuration changes, and therefore should not require an additional wait period.

3.3. Introduce Name Server Operators as Participants in the Process

> *Changes are requested and coordinated with the administrative and technical contacts of a domain. These are not necessarily the same parties that operate the authoritative name servers for a domain. These operators may authorize changes instead, although such an approach would be a radical departure from current operations and would require a new procedure that dealt with the roles and responsibilities of these operators, as well as IANA procedures to authenticate these parties.*

A further distinction should be made between the party operating the physical name server and the party maintaining the zone containing the server name, as these may or may not be the same. In the case of a shared TLD server name belonging to an independent name service provider, for example NS-EXT.ISC.ORG, they probably are the same, but in the case of an in-domain TLD server name pointing at a

server located in another country for geographic diversity, for example F.EXT.NIC.FR, they are probably not the same.

In the cases where they are not the same, the maintainer of the zone containing the server name seems like the more appropriate party for IANA to interact with for purposes of verifying changes to root glue, since they, not the operator of the physical server, maintain the authoritative data that the glue is supposed to be a copy of.

The number of new relationships that would have to be established between IANA and these zone maintainers can be very roughly estimated as the number of distinct zones that contain one or more TLD server names but are not themselves TLD zones; that number currently stands at 385. The actual number of new relationships will be smaller because many of these zones share maintainers or are maintained by TLD operators.

3.4. Move to a Narrow Glue Policy

Several of the responses to (ICANN 2006) supported the discussion, investigation, or long-term adoption of a narrow glue policy for the root zone, though none went as far as proposing its immediate adoption.

Moving to a narrow glue policy in the root zone would eliminate most, but not all, of the cases where root glue updates affect multiple TLDs and therefore require multiple approvals under the current procedures. The remaining cases are those where a server name is acting as an in-domain server for one TLD (and therefore would still be eligible to have root zone glue under a narrow policy), while simultaneously acting as an out-of-domain server for one or more other TLDs. In the current root zone, there are 44 such server names, and the largest number of TLDs sharing such a server name is 7.

Moving to a narrow glue policy could potentially have a negative impact on the efficiency and resiliency of the DNS, as resolvers would have to perform separate lookups for addresses that were previously included in referrals from the root servers, thereby causing additional traffic, additional round-trip delays, and additional points of failure.

Somewhat paradoxically, providing fewer glue records in referrals does not make resolvers any less dependent on glue for resolution, and may actually increase the number of glue records they need to rely on. The reason for this is that when a resolver executes the resolution algorithm for the first time, without the benefit of cached information from previous resolutions, each time it follows a referral from a parent domain to a child domain it must make use of glue, either directly or indirectly. Glue is used directly when it is included in a referral from the parent; it is used indirectly when the referral does not include glue. Such indirect use of glue is unavoidable, because when the referral lacks glue, the resolver must invoke itself recursively to do a separate name server address lookup, and that recursive invocation will itself either make use of glue provided in a referral, or cause yet another recursive

invocation, and so on, forming a chain of recursive invocations that can only successfully terminate through the use of glue.[2]

To determine the extent to which separate address lookups would be required for resolution under a narrow root glue policy, a simulation was run resolving each TLD using only the glue that would remain under a narrow policy. According to this simulation, of the 280 TLDs, 222 would still have at least one glue record in referrals from the root, thereby allowing an initial resolution to succeed without the need for a separate address lookup. Another 57 TLDs would get glueless referrals from the root, and therefore would require a separate address lookup, but that address lookup would itself be in a TLD whose referrals from the root zone contain glue. Finally, in the case of one TLD, not only would the TLD itself lack glue, but the address lookup could involve a second TLD whose referrals from the root zone lack glue, such that a second, nested address lookup must be performed. There are currently no cases where more than two levels of nested address lookups are needed, and no circular dependencies among TLDs that would make resolution fail entirely.

A significant part of the glue that would be included in referrals from the root servers under the narrow glue policy would be out-of-domain for the target TLD of the referral, but nevertheless present in the root zone because the server name meets the narrow glue policy with respect to some *other* TLD and is shared by the target TLD. Without this sharing, one address lookup would be required to resolve 50 TLDs, and two address lookups would be required for 22 TLDs.

The above analysis makes a number of assumptions regarding the algorithms used by the resolver. For example, it assumes that when the resolver receives a referral response having glue for some but not all of the servers, it will make use of the available glue rather than attempt to look up the missing addresses before proceeding to contact the servers for which glue is available. Also, it only considers address lookups that are strictly necessary for resolution; even if the initial resolution does not require a separate address lookup, practical resolvers will sooner or later look up additional addresses in order to spread the load evenly over all the authoritative servers, causing a significant amount of additional traffic. Finally, it does not consider addresses included in referrals as a result of the root servers being authoritative for zones other than the root itself. If switching to a narrow glue policy is seriously contemplated, additional tests using real-world resolver implementations should be performed.

3.5. Use Non-shared Server Names

Since the long delays in processing root glue updates under the current policy have been in cases where server names are shared by many domains, one approach to reducing the delays would be to reduce the degree to which TLDs share server names.

To switch from a shared server name to a non-shared one, a TLD needs to change an NS record in the TLD zone to point to a different, non-shared name, arrange for an A record containing the IP address of the server to be added at that name, and submit a root zone change request to change the NS record

[2] Strictly speaking, a successful resolution requires that address records are provided in referrals, but these address records need not necessarily be glue; for an example of a configuration where they are not glue, see section 5.3.

and add the new A record as glue. No operational changes are needed: the physical server and its IP address can remain shared, only now the server is being shared by IP address rather than shared by name.

The new, non-shared server name can be created in any domain, but typically it will be placed either within a domain belonging to the server operator or within the TLD itself. The option of placing it within the TLD itself will be examined more closely in section 3.6.

When a server is shared by IP address, the root zone will contain multiple glue records with the same IP address. Such configurations are already supported by the root registry, as can be seen from the presence of 37 such shared IP addresses as glue in the current root zone.

Although switching to non-shared server names would solve the current problem that a single objecting or unresponsive TLD can effectively veto a root glue update for a long time, it does not change the overall number of approvals needed when a shared physical server is renumbered: if the server is shared by name by N TLDs, there is a single root glue change request requiring approval by all N TLDs, but when it is shared by IP address, there will instead be N separate root glue change requests, each requiring approval by one TLD. However, because these change requests would be initiated by each TLD separately and would not require approval by other TLDs, they could take advantage of future systems for self-service automation as discussed in section 4.5.

3.6. Use In-domain Server Names

When switching from shared to non-shared server names, the new server name is commonly created within the TLD itself, making it an in-domain server name. Currently, most TLDs, 167 out of 280, are using a mix of in-domain and out-of-domain server names. 72 TLDs have no in-domain server names, and 41 TLDs have only in-domain server names.

If all TLDs currently using out-of-domain server names were to switch to in-domain server names, it would have the effect of eliminating the sharing of server names. It would also render the question of whether the root should use a narrow or wide glue policy moot, because there would no longer be a practical difference between the two — the root zone would contain glue for every TLD server address under either policy.

As with other changes to server naming, switching to in-domain server names only requires changes to the DNS data, no operational changes. The physical servers and IP addresses can remain the same, and there are no restrictions on the ability to have servers operated by third parties, in different networks, and in different geographic locations.

When a TLD switches from shared to non-shared server names, using names within the TLD rather than a domain belonging to the server operator has two main advantages:

- It reduces the risks of transitive trust, discussed in detail in the following section
- Assuming the name is within the TLD zone and not within a delegated subdomain, the change can be made by changing the data in only two zones (TLD and root) rather than three.

One potential risk of in-domain server names is that they may cause problems for resolvers using a "delegation-only" configuration; this issue is discussed in section 3.6.2.

With regards to the question of IANA's possible duty to notify TLDs about the renumbering of shared name servers, discussed in section 2.5, the use of in-domain server names pointing to shared IP addresses presents some additional challenges. First, notification of renumbering becomes increasingly important, because unlike the case of renumbering a server shared by name, where configuration changes at the TLD may not even be needed, the renumbering of a server shared by address always involves changes to the in-domain address record as well as an active request by the TLD to update the corresponding root glue. Second, when a TLD submits a such a root glue update request, it is difficult for IANA to determine whether the change also affects other TLDs, because a glue change request involving a shared IP address could mean either that the shared server is being renumbered (affecting all TLD sharing it), or that the TLD submitting the change request is simply switching to a different server.

3.6.1. Transitive Trust

The use of out-of-domain server names has long been known to be a potential security risk because it creates a situation where the compromise of an authoritative server can be exploited to forge DNS data not only in the domains directly using that server, but also in other domains using NS records pointing to server names within those first domains, and so on (Bernstein 2000). This effectively causes the trust relationship that a domain places in its servers to cascade transitively through multiple domains, hence the term *transitive trust* (Ramasubramanian and Sirer 2005).

Attacks based on this vulnerability are practical, as recently shown in a step-by-step demonstration of how the compromise of one server in the US could lead to the injection of forged data for the .FR domain into a caching server, by exploiting a transitive trust relationship spanning four domains (Dempsky 2009). The .FR domain is no longer vulnerable to this problem as it is now using in-domain server names, but other domains still using out-of-domain server names are vulnerable to similar attacks. Experiments performed as part of this study indicate that the attack is not limited to the dnscache caching server used in the original demonstration, but with minor modifications, is also effective against caching servers running a current version of BIND.

3.6.2. Problems with "delegation-only" Configurations

In response to the controversial introduction of a wildcard A record into the COM and NET zones by VeriSign for their Site Finder service in 2003 and similar use of wildcards in other TLDs, some caching server implementations now support a configuration option to suppress the effect of such wildcard records by replacing authoritative (i.e., non-delegation) responses from TLD servers by name error (NXDOMAIN) responses.

For example, BIND supports a configuration option called *delegation-only*, which applies such a transformation to a specific list of TLDs, as well as the option *root-delegation-only*, which applies it to the root domain and all TLDs except those listed in an optional exclusion list.

As pointed out in Nominet's response to (ICANN 2006), when a domain switches from out-of-domain to in-domain name servers, the authoritative address records for those servers can become subject to these transformation, causing them to be effectively ignored by caching servers that have either configured the TLD as a *delegation-only* domain, or have enabled root-*delegation-only* and failed to include the TLD in the exception list (Nominet 2007).[3]

This type of configuration error can go undetected for a long time, because resolution will initially rely on glue provided by the root servers. Problems will arise only when the glue expires, or if lookups of the authoritative address records are triggered by explicit queries sent by clients, possibly as a deliberate denial of service attack.

In principle, this problem should only apply to server names that are in-zone, and placing the server names in a delegated subdomain should be an effective work-around. However, the BIND *delegation-only* documentation mentions that there is a possibility of "false positives" in the case of servers serving both a parent and child, which implies that it may also be necessary to serve that delegated subdomain from different servers than the TLD itself; other implementations may also have similar problems. Further research on this issue is needed.

A possible solution to this problem would be for caching server vendors to change the implementation of the *delegation-only* functionality such that the rewriting of non-delegation responses into NXDOMAINs would only apply to external queries from clients, not to internally generated lookups of name server addresses. For example, data from non-delegation responses could be cached without transformation but flagged, and any flagged cache entry would be transformed into an NXDOMAIN when requested by a client, but not when used internally by the server.

3.7. Use TLDs as Proxies for the Server Operators

DENIC's response to (ICANN 2006) proposed a procedure involving a threshold scheme similar to that of section 3.1, together with the idea that the subset of the affected TLDs consulted by IANA would not be asked to approve of the glue change as such, but to act as "witnesses", verifying that the renumbering has taken place by communicating with the server operator (DENIC 2007).

This proposal could be seen as a way to achieve much of the same effect as the proposal to introduce the server operators as direct participants into the process (section 3.3), and therefore shares much of the appeal of that proposal, while avoiding the need to establish direct relationships between IANA and the server operators, instead using TLDs as proxies.

[3] The comment actually stated that the vulnerability applies to glue with more than two labels in the TLD zone, but it is not clear why the problem would be related to the number of labels, and tests performed as part of this study indicate that the problem also affects in-zone server names having exactly two labels.

4. Automated Glue Updates

The question has been raised whether the current procedure for handling root glue changes could be replaced by a fully automated process to synchronize the root zone glue records with the corresponding authoritative data.

In theory, it should be possible to handle glue updates completely automatically, as the glue should always be an exact copy of the corresponding authoritative address records, and the authoritative address records can be determined through the automatic process of DNS resolution. A similar opportunity for automation also exists with regards to updating the delegating NS records in the root zone, which should be an exact copy of the corresponding authoritative NS records in the TLD zone.

Such a fully automated procedure would effectively replace both the existing procedure for submitting root glue change requests and the existing procedure of verifying the validity of the request with the zone contacts. It would also be possible to implement partially automated procedures where only the submission phase or only the verification phase is automated; in the following sections, these two aspects will be considered separately.

4.1. Automated Submission of Root Glue Change Requests

Server renumbering events could be discovered by means of automated cross-checks of the root glue with the corresponding authoritative data. Such cross-checks could be run either periodically or triggered by requests from affected parties, and could either completely replace the existing process of manually submitted root glue change requests, or operate in parallel with it, serving as a backup procedure to ensure that address changes are discovered even if the affected TLDs fail to manually submit a glue update request.

In the existing process, root glue update requests are unauthenticated, and the validity of a request is ensured through separate verification steps. Assuming this security model is retained when the submission process is automated, the automated submission procedure is not a security critical component and could be designed with an emphasis on simplicity and efficiency rather than maximum resilience against attack; if it generates an incorrect glue update request, it will be caught at the later verification stage.

Because the current procedures allow root glue update requests to be submitted by anyone, it would in principle be possible even today for a third party to begin performing automated cross-checks and submitting update requests to IANA, without any change to the IANA procedures.

4.2. Automated Verification of Root Glue Change Requests

Whether the need to update root glue is discovered through the manual submission of a change request or through an automated periodic cross-check as discussed in the previous section, IANA must verify that the requested update is correct and will not harm the stability of the DNS.

The current procedure does this through a combination of technical checks (IANA 2009) and authorization by the contacts of the affected TLDs. A fully automated procedure would have to rely entirely on technical checks. This raises the question of whether either the current set of technical checks or some extended set of checks can determine the authoritative server address with a sufficient level of certainty, and ensure that the glue update will not harm the stability of the DNS. The following sections discuss some of the issues that need to be considered in making that determination.

4.3. Security Considerations

A major concern in relying on automated mechanisms to authorize a root glue update is that an attacker might be able to mislead them into accepting an incorrect server IP address, which would then be automatically published as glue in the root zone. Since the process of determining the authoritative server IP address is essentially a DNS resolution, it is subject to the same risks facing other resolvers, such as brute-force spoofing attacks and the issues of transitive trust discussed in section 3.6.1.

The argument can be made that introducing automated glue updates would not in itself introduce any new risks, because assuming the resolver used by IANA is at least as resilient to attack as those used to answer client queries in production, any attack against the resolver used for verifying the updates could already be used directly against production resolvers. On the other hand, a successful attack against the glue update system could have a far greater impact than one against any given production resolver, as it would affect the entire Internet and not just the users of that single resolver. In other words, the glue update system represents a high-value target that can be used to amplify the effects of existing attacks.

4.3.1. Spoofing Attacks

As the glue update system would represent a high-value target for spoofing attacks, it is important that it is highly resilient to such attacks. If or when DNSSEC is deployed in all zones containing TLD server addresses, it will offer a defense against these attacks, but at present, other means need to be employed.

In practice, the checks currently specified in the technical requirements for TLD servers (IANA 2009) are probably sufficient, provided they are implemented correctly and with security in mind. These checks involve making multiple queries that must yield consistent results, and the queries are made to different servers, over TCP as well as UDP, and originating from different geographical locations. Even though an attacker may be able to inject spoofed responses to individual queries, it is highly unlikely that he would able to do so consistently for all these queries, especially those made over TCP. To ensure that there are no common points of failure affecting multiple queries, the queries should be made entirely using special-purpose tools that perform the necessary DNS resolution autonomously, without relying on previously cached data or general-purpose caching resolvers.

4.3.2. Compromise of Trusted Servers

As discussed in section 3.6.1, a TLD has a trust relationship not only with its own servers, but also with the domains containing the names of its servers, and consequently with the servers serving those domains. Although the possibility that the latter servers could be compromised already poses a risk to

the TLD, an automated glue update mechanism could be exploited to increase the impact of such compromise.

One TLD that clearly illustrates these trust relationships is the BB domain. BB is served by the servers NS1.BARBADOSDOMAINS.NET and NS2.BARBADOSDOMAINS.NET; for brevity, these will be collectively referred to as NS[12].BARBADOSDOMAINS.NET. The BARBADOSDOMAINS.NET domain is in turn served by three servers, NS[123].MDNSSERVICE.COM.

To automatically update the root glue records for NS[12].BARBADOSDOMAINS.NET, the glue update system would query one or more of NS[123].MDNSSERVICE.COM for the address records of NS[12].BARBADOSDOMAINS.NET, and then install those address records as the new root zone glue.

The servers NS[123].MDNSSERVICE.COM are already trusted by the BB TLD in the sense of transitive trust discussed in section 3.6.1, and therefore a party controlling those servers is already in a position to mount forgery attacks against BB, but the consequences of those attacks still fall short of a complete control of all DNS traffic for the BB domain. With an automated glue update system, complete control could be achieved. If the BARBADOSDOMAINS.NET zone file hosted on NS[123].MDNSSERVICE.COM were maliciously modified to make the address records for NS[12].BARBADOSDOMAINS.NET point at different servers, the automated root glue update system would copy those records into the root zone as glue, with the result that all DNS traffic for BB would be directed to the new servers.

4.3.3. Delayed Recovery from Compromise

If an attacker were to momentarily gain control over the contents of a TLD zone that uses in-zone server names, and changed the server addresses in the zone to point to a different set of servers under his own control, the automated root glue update mechanism would quickly reflect this change in the root glue. Even if the legitimate owner of the TLD quickly regained control of the zone contents and changed the addresses back, the automated mechanism would not change the root glue back, because it would now be querying the attacker's servers. Instead, the legitimate TLD owner would have to go through administrative channels to regain control of the root glue. The end effect of this is to turn a short-lived compromise of TLD zone into a longer-lived one.

4.4. Operational Considerations

4.4.1. Accidental Misconfiguration

The administrator of the authoritative address records for a server name could accidentally configure them to point at a server that is not correctly configured as authoritative for all the TLDs using that server name, or perhaps not even responding at all. In such a situation, it would be inadvisable to update the root glue to match the new, incorrect authoritative data; the misconfiguration of the authoritative records can of course alone be sufficient to cause problems, but mirroring the error in the root glue can only make the situation worse. Thankfully, the existing technical requirements for name servers already include a requirement that the servers answer authoritative for the SOA record, which should be sufficient to catch most errors of this nature and stop the automatic glue update from proceeding.

4.4.2. Reliance on Old Glue to Verify New Glue

When the server being renumbered is in-domain, the process of verifying its authoritative address through resolution can involve a chicken-and-egg situation where that resolution involves sending queries to that same server, to the address we are trying to determine. The standard way to resolve this dilemma is to rely on glue, but in this case, the reason we are doing the resolution in the first place is that the glue in question is believed to be incorrect.

If the verification was triggered by manually submitted update request specifying a new server address, it would be tempting to make the resolution use the address given in the request (which is supposedly current) rather than the one in the root glue (which is supposedly out of date), but given that update requests are not authenticated, doing so would be insecure, allowing the trivial takeover of a server address by means of a fraudulent change request. Instead, the resolution must rely on the existing root glue despite the fact that it is believed to be out of date, which in practice means that the old server addresses are queried to determine the new server addresses. In a properly staged renumbering where both the old and new addresses are responding during a transition period, this will work fine, but if the old servers are no longer responding or not responding correctly, the verification will fail and manual intervention will be required.

4.4.3. Handling of Failures of Technical Checks

Although the technical checks to verify that the newly renumbered servers meet IANA's technical requirements can be automated, it is less clear what will happen in the case where the servers fail to meet the requirements. To some extent it may be possible to deal with such situations by automatically notifying the affected TLDs of the failure, but if the problem persists despite such notification, manual review and communication with the partied involved is needed.

To make a rough estimate of the magnitude of this problem, measurements were made by querying the TLD servers that are currently missing root glue for one or more of their authoritative addresses. This would be the same set of servers for which the initial run of an automated glue updating procedure would attempt to add new glue. Each such server was sent an SOA query for each TLD served, and the success or failure of the SOA query was recorded. Out of 98 servers with missing glue, 23 servers failed to respond to at least one of the SOA queries.

A similar measurement querying the servers that currently have inconsistent glue (i.e., where glue is present, but differs from the authoritative address) showed that out of 12 such servers, four failed to respond to at least one of the SOA queries.

Since correctly answering an SOA query is just one of several technical requirements checked by IANA, there are likely to be additional failures when using the full battery of checks.

4.4.4. Limited Automation of Subsequent Processing Steps

After the submission and verification of the root glue update request, the current procedure involves the further steps of U.S. Government authorization and implementation in the root zone by VeriSign. It

is unclear whether these steps can also be fully automated, or whether they will still require manual processing.

4.4.5. Lack of Operational Experience

Although fully automated updating of glue records would be applicable not only to the root zone but also to the TLDs and other parts of the DNS tree, I am not aware of any existing deployments of such systems. It would seem prudent to test the technique in a TLD or other lower-level domain to gather operational experience and gain confidence in it before deploying it in the root.

4.5. Self-service Automation

As an alternative to the fully automated glue update systems outlined above, root glue updates could also benefit from a different form of automation, namely the automated processing of manually submitted, authenticated change requests. This is the method typically used by second-level domain administrators to update glue records in the TLDs, with password-protected HTTPS web forms serving as a common method of authenticated manual submission.

Work is already underway within IANA to introduce similar authenticated mechanisms for submitting root zone change requests. For glue updates where approval by multiple parties is currently required, the authenticated submission of a request by a TLD could serve as an implicit approval by that TLD, but the other approvals will still require manual processing. If the need for multiple approvals could be eliminated, for example through a move to a narrow glue policy or the voluntary transition of TLDs to in-domain servers, the process could become mostly automatic.

5. Other Approaches

In the course of preparing this report, a number of additional options have come to mind. I am presenting them here for consideration, without any implied criticism of or opposition to the options already presented.

5.1. Allow Approval by Either Server Name Owner or All Affected TLDs

One possible procedure for verifying root glue changes would be to apply either the current procedure of obtaining approval from each affected TLD, or, at IANA's discretion, verifying the change with the contacts for the zone containing the server name.

This procedure aims to combine the advantages of the proposal to introduce name server operators as participants in the process (section 3.3) with those of the current policy of requiring approval by all affected TLDs.

As discussed in section 3.3, having IANA establish contacts with the maintainers of every zone containing a TLD server name would be burdensome. On the other hand, most server names are used by only one or a few TLDs, and the existing procedure of verifying the change with each affected TLD works well in these cases because it will involve communication with only a small number of parties. By establishing

new relationships with the zone maintainers for a small number of server names used by large numbers of TLDs, either in advance or when changes to those names are requested, IANA can use those relationships to verify change requests involving many TLDs, while continuing to use the current procedures for requests involving few TLDs. This way, IANA could achieve a large reduction in the overall amount of communication needed to verify a change in return for a modest increase in the number of relationships maintained.

5.2. Use Separate Policies for Glue Addition and Deletion

Adding a new glue record to the root zone must be subject to careful checks because an incorrect or unauthorized glue record could compromise DNS integrity. On the other hand, the deletion of a glue record will at most impact efficiency, or only in extreme cases such as the removal of all glue for a domain with only in-domain servers, availability, whereas the *failure* to delete a glue record can impact integrity, as discussed in section 2.4.

Therefore, it would make sense to apply different policies to the cases of glue addition and deletion, with more extensive checking and approval required in the addition case.

For example, in the event of a lengthy dispute regarding a root glue update, IANA should have the option of deleting the disputed glue until the dispute has been resolved. Similarly, if there is a discrepancy between the root zone glue and the authoritative data, and the server pointed to by a root glue record does not meet IANA's technical requirements, but the one pointed to by the authoritative address record does, there should be no obstacle to the removal of the glue from the root zone even if the addition of a replacement record is subject to delays.

5.3. Use In-Parent Server Names

Although server names are usually either in the domain they serve or in a domain belonging to the server operator, there is also a third option that should be mentioned in the interest of completeness, namely placing them in the parent domain. This configuration is seldom or never seen in practice even though it has several desirable properties:

- Only one copy of the server address record needs to be maintained
- The server address can always be included in referrals
- No glue is involved; the sole copy of the server address is authoritative data

In the case of the TLD delegations, this option would mean placing the TLD server names in the root zone.

There are many possible naming schemes for such an arrangement. One would be to reserve a non-delegated name in the root for this purpose, e.g., TLD-SERVERS, and create further non-delegated subdomains under it for each TLD and its servers. With such an arrangement, the delegation for EXAMPLE might involve the following records, all part of the root zone:

```
EXAMPLE. NS A.EXAMPLE.TLD-SERVERS.
EXAMPLE. NS B.EXAMPLE.TLD-SERVERS.
```

```
A.EXAMPLE.TLD-SERVERS. A 192.0.2.1
B.EXAMPLE.TLD-SERVERS. A 192.0.2.2
```

No server address records would be needed in the TLD itself.

One potential problem with this scheme is that when the root zone is DNSSEC signed, the address records are signed as part of it, causing referral responses to grow significantly in size as DNSSEC signatures are added to the additional section. This could possibly be worked around by making the TLD-SERVERS domain a separate, insecurely delegated zone and making all root servers authoritative for it. Another potential issue is that caching servers configured to do *delegation-only* rewriting for the root zone could be affected by the problems discussed in section 3.6.2.

Interesting as this option may be, considerable operational experience would have to be gathered in lower-level domains before it could be seriously considered for use in the root zone.

6. Conclusion and Recommendations

A large number of different solutions have been proposed for the problem of excessive delays in the IANA processing of root glue updates, and many of them appear viable from a technical perspective.

As discussed in sections 2.5 and 2.6, IANA's primary concern should be to ensure the correctness of the glue in the root zone, not to notify the TLDs of changes to it or to seek their approval. In principle, verifying the correctness of glue could be done entirely through automated technical checks, as discussed in section 4. However, I would not go so far as to recommend that IANA rely on technical checks as the only means of verification.

Instead, I would recommend a adopting a procedure where the technical name server checks are considered the primary means of verification for glue updates, but some minimal amount of verification through communication with TLD or other contacts is retained as a secondary safeguard. One reasonable choice for such a minimal administrative verification procedure, but by no means the only possible choice, would be the one suggested by DENIC (section 3.7).

Independently of the above, I would strongly encourage TLDs currently using out-of domain server names to migrate to in-domain ones, while keeping in mind the *delegation-only* problem discussed in section 3.6.2. This is primarily motivated by the security issues related to transitive trust (section 3.6.1), but moving to in-domain server names would also alleviate the need for large numbers of approvals under the current IANA procedures and pave the way for future self-service automation of glue change requests.

References

Bernstein, Dan. "The 200 trusted .com servers." January 2000.
http://seclists.org/bugtraq/2000/Jan/336

Dempsky, Matthew. "DNSTrust - 28 queries later: an example attack on .fr." March 2009. http://shinobi.dempsky.org/~matthew/dnstrust/example.html

DENIC. "DENIC's Comments on DNS Root Zone Glue Policy." January 2007. http://forum.icann.org/lists/root-glue-comments/msg00006.html

IANA. "Technical requirements for authoritative name servers." June 2009. http://www.iana.org/procedures/nameserver-requirements.html

ICANN. "Comments Sought on DNS Root Zone Glue Policy." December 2006. http://www.icann.org/en/announcements/announcement-3-05dec06.htm

Koch, Peter. "DNS Glue RR Survey and Terminology Clarifiaction" (work in progress). November 2007. http://tools.ietf.org/html/draft-koch-dns-glue-clarifications-03

Mockapetris, P. "Domain Names - Concepts and Facilities." November 1987. http://www.ietf.org/rfc/rfc1034.txt

Nominet. "Nominet's Comments on DNS Root Zone Glue Policy." January 2007. http://forum.icann.org/lists/root-glue-comments/msg00007.html

Ramasubramanian, Venugopalan, and Emin Gün Sirer. "Perils of Transitive Trust in the Domain Name Systems." *Proceedings of the 5th Conference on Internet Measurement.* 2005, p. 379-384. http://www.cs.cornell.edu/people/egs/papers/dnssurvey.pdf