

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 10, 2015

P. Koch
DENIC eG
March 9, 2015

DNS Glue RR Survey and Terminology Clarification
draft-koch-dns-glue-clarifications-05

Abstract

This document presents a survey of the use of the term "glue record" in DNS related RFCs and proposes a terminology for the various glue policies seen in different top level domains.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 10, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	2
1.1.	Purpose of this Document	3
1.2.	RFC Survey	3
2.	Terms	3
3.	Name Server Naming Strategies	4
4.	Glue Policies	5
5.	Special Considerations	6
5.1.	Root Server "Glue" in the Root Zone File	6
5.2.	Using Glue records in responses	6
5.3.	Registry considerations for maintaining Glue RRs	7
5.4.	Glue RRs for multihomed name servers	8
5.5.	Grandchild Glue	8
5.6.	DNSSEC Considerations	8
5.7.	IPv6 Considerations	9
5.8.	Open Issues	9
6.	Security Considerations	9
7.	IANA Considerations	9
8.	References	9
8.1.	Normative References	10
8.2.	Informative References	10
Appendix A.	Extended RFC Survey	13
Appendix B.	Document Revision History	13
B.1.	Changes from -04 to -05	13
B.2.	Changes from -03 to -04	13
B.3.	Changes from -02 to -03	13
B.4.	Changes from -01 to -02	14
B.5.	Changes from -00 to -01	14
	Author's Address	14

[1. Introduction](#)

When delegating zones from a top level domain (TLD) or other DNS zone, some additional information is needed when resolving a name server's (as per an NS RR) address would involve the particular name

Koch

Expires September 10, 2015

[Page 2]

server itself. Such a dependency on itself, direct or indirect, may effectively shadow a part of a zone's NS RRSets, reducing redundancy, or even render the zone completely unresolvable. This additional information is an amendment to the delegation in the form of glue address records. In the real life DNS multiple strategies to determine necessity or acceptance of glue records co-exist. This document lists a subset of those approaches.

This document also tries to clarify when and where to call an address record "glue record".

Comments should be directed at the author.

Domain names and IP addresses herein are for explanatory purposes only and should not be expected to lead to useful information in real life [[RFC2606](#)], [[RFC5735](#)]).

1.1. Purpose of this Document

XXX gather intelligence, clarify terminology, list name server naming schemes and glue policies, collect glue related operational and protocol considerations.

1.2. RFC Survey

To find out more about early motivations and strategies for DNS glue records, all existing RFCs were automatically searched for the term "glue" (case insensitive, no word boundaries) and those matching were inspected on a case by case basis. Whenever the term was used in a DNS context, the RFC was added to the list which can be found in the References section. It turned out that, while all early RFCs are consistent in using "glue" only for type A address records for NS RR targets, they apply slightly different logic as to when a glue A RR should be present.

2. Terms

When the term "glue record" was introduced in [[RFC0973](#)], it was meant to denominate both data origin and purpose. Data origin is related to the zone, although the glue records do not belong to the authoritative zone data. The purpose is constrained to providing address information for name servers mentioned in NS RRs, which would otherwise not be resolvable. Glue records are address information accompanying a delegation (in the delegating zone).

There is sometimes confusion when data in a DNS response is also called "glue" data, e.g. [[RFC2010](#)] starts speaking of "fetching" glue. In a DNS response packet (answer or referral) the address

information for name servers is carried in the additional section. This address information might have originated from glue records but might also come from cached or authoritative data. DNS data in response packets should only be called "glue data" when it is certain and needs to be emphasized that it originates from glue records.

[RFC4472] introduces the concept of 'critical' and 'courtesy' additional data.

3. Name Server Naming Strategies

The DNS refers to name servers by their name in NS resource records. The servers' names can have different relations to both the delegating zone and the zone delegated to them. The following categories describe the naming scheme from the perspective of the delegated zone. They will be illustrated with the help of the "example" TLD zone, which delegates the zone "del.example" to the respective sets of name servers.

"in domain": A name server can be within the delegated domain, which includes the name of the delegated domain itself, e.g., "del.example", "dns.del.example", or "dns.sub.del.example". In all but the first case the name might be part of a subzone of "del.example". This naming scheme is sometimes also called "in-bailiwick".

"sibling domain": A name server's name can be within the delegating but outside the delegated domain, e.g., "dns.other.example". There are two sub-cases here, depending on the treatment of the "other.example" domain:

"authoritative in sibling zone": When the name server's name is within the parent domain, but in a separate zone, this is a sibling zone of the delegated zone. XXX Strictly speaking, see "grandchild glue", it might be a nephew/niece zone.

"authoritative in parent zone": A name server's name may also appear authoritative in the parent zone.

"unrelated": A name server's name may not share its parent with the delegated domain, e.g., "dns.example.org".

The discussion of the advantages and disadvantages of the three categories or any combination is beyond the scope of this document.

4. Glue Policies

In the DNS tree different policies are applied with respect to registering glue with the delegating zone. "Registering" in this case means that the respective glue information is accepted, requested or required and then attached to the zone data so that it is available at all authoritative servers, i.e. the glue travels with the zone data by AXFR, IXFR or other means. However, it does not make the glue data part of the zone's authoritative data [[RFC5936](#)].

This is a list of existing glue policies:

"never" or "null": Glue RRs are never registered. This currently applies to larger parts of the IN-ADDR.ARPA reverse tree.

XXX but this interacts with the naming schemes to be applied.

"narrow": Glue RRs are registered if and only if the name server resides within or below the delegated (child) zone (that is, within the delegated domain). This was suggested by [[RFC1034](#)].

"wide": Glue RRs are registered if and only if the name server resides below the delegating (parent) zone. If the name server's name belongs to the parent zone, it is part of the authoritative data and thus there is no need to register glue RRs. This was suggested by [[RFC1033](#)]. It is used for the root zone.

"case by case": Glue RRs are registered following the "narrow" policy except where there are (circular) dependencies that demand additional glue RRs.

"mandatory": Glue RRs are always registered for all name servers. This was suggested by [[RFC0973](#)].

"other": Combinations of the above may exist, e.g. if a registry runs multiple sibling domains and decides to register glue RRs whenever a name server resides in or below one of the siblings. This category would also include other policies like "random" or "arbitrary".

Glue RRs are needed only in the delegating zone, regardless of glue policy. See [Section 5.1](#) for a discussion of root zone issues.

Various RFCs have identified extraneous glue RRs as sources of error and confusion ([[RFC1713](#)], [[RFC1912](#)]).

5. Special Considerations

5.1. Root Server "Glue" in the Root Zone File

As said before, Glue is meant to be present in the delegating zone only. The only exception seems to be root zone which also contains the address records for its authoritative name servers. However, with the current setup the root servers also serve the ARPA domain and with the root zone's "wide" glue policy this means that there should be glue RRs for this particular set of nameservers, but only in their capacity as ARPA TLD servers.

[The position of the A RRs in the root zone file (which has just editorial value) as well as their TTLs suggest that historically there will have been a different reason].

[The upcoming change in the ARPA NS RRSet in the spirit of [[RFC5855](#)] renders this consideration void.]

Also, per operational practice, all root servers are authoritative for the zone they reside in and that is today also reflected by the official delegation for the "root-servers.net" zone. So, they have the authoritative data present and do not need to rely upon the data transported with the root zone.

[To have a complete trust chain available at the root servers leading to their own names, it would be useful to have them configured authoritative for all intermediate zones. It has been suggested before to move the root server's names to a distinct TLD. Another option would be to move their names to e.g. ROOT-SERVERS.ARPA instead.]

5.2. Using Glue records in responses

Some implementations use Glue information not only during additional section processing, but also in the answer section of responses.

Given an excerpt of the "example" TLD zone file,

```
one.example.      NS dns.one.example.  
                  NS dns.two.example.  
dns.one.example.  A  192.0.2.53
```

what should a name server authoritative for the example TLD do when asked for the A RR for dns.one.example? Some implementations will put the A RR in the answer section of the response, others will respond with a referral and only copy the glue A RR into the

additional section (the handling of `dns.two.example`'s A RR is not considered here).

Step 4 of the algorithm in 4.3.2 of [\[RFC1034\]](#) suggests that after copying the NS RRs into the authority section (in step 3b) the cache should be consulted and used to fill the answer section. Depending on whether or not Glue data is considered to reside in the cache (it is definitely not authoritative), one or the other response type will be preferred.

With DNSSEC an A RRSets response originating from glue data will always miss the appropriate signature, because neither does the delegating zone sign the glue RRSets nor does a glue RRSIG (child's signature covering the address RRSets) exist in that delegating zone.

[discuss levels of indirection and operational reasons that lead to the "gluepot response"]

5.3. Registry considerations for maintaining Glue RRs

As explained in [\[RFC5936\]](#), Glue address records are "registered with" a zone, but since they fall underneath (or, sometimes, onto) the next zone cut, they are not part of that zone. Depending on the data model and glue policy in use for a TLD ([\[RFC5731\]](#), [\[RFC5732\]](#)), different side effects may be the consequence of undelegating a domain. The standard zone file format does not allow for the explicit dedication of address records as glue information. Instead, the distinction is made based on the presence or absence of zone cuts. If, for example, the `del.example` domain was delegated to, amongst others, the `dns.del.example` name server, an address record for `dns.del.example` in the `example` TLD zone will be interpreted as glue record. After deleting the `del.example` NS RRSets (the delegation) from the `example` zone, the corresponding address record would have to be deleted, as well. Should it remain, it would be elevated to authoritative data, since there no longer is a zone cut. Such an effect might be highly undesirable and should be avoided.

Custom name server software may be able to keep delegation and glue data separate from the authoritative data so that `dns.del.example` would still exist but would not be elevated to authoritative status. However, this effect is similarly undesirable since the address might be used to fill the additional section for referrals containing NS RRs pointing to `dns.del.example`, as if "wide" glue policy was in effect.

With the "wide" glue policy, a glue address record registered with some delegation might not even be related to the delegation of its own second level domain, i.e., the corresponding name server does not

have to be part of the NS RRSets for that domain. Therefore, broader checks have to be applied to avoid the aforementioned undesired effects.

5.4. Glue RRs for multihomed name servers

Some name server names resolve to A or AAAA RRSets consisting of more than one record, i.e. they have multiple addresses. It is recommended that these RRSets be consistent between the child and the parent.

Research is needed to evaluate the effective difference between multiple names and multiple addresses for a name server. These effects heavily depend on server selection algorithms in resolvers.

5.5. Grandchild Glue

When a name server resides within the delegated domain, the delegation needs a glue record with both the "wide" and the "narrow" glue policy. However, the server does not necessarily have its name within the delegated zone since it may belong to a child or grandchild zone of the delegated one.

This is a delegation in the example TLD:

```
one.example.      NS one.example.  
                  NS dns.one.example.  
                  NS dns.deep.one.example.
```

Only the first name server is known to have its name in the delegated zone, where the second and third could both be in separate zones. NB: even dns.one.example. could be a zone delegated from one.example.

As a consequence, it cannot be concluded that any such name server is able to authoritatively serve its own name, e.g., if it does not serve the grandchild zone.

5.6. DNSSEC Considerations

DNSSEC signatures do not cover glue records [[RFC3833](#)], [[RFC4033](#)].

Using the gluepot to fill the answer section is discouraged with DNSSEC, see [Section 5.2](#).

5.7. IPv6 Considerations

While this document makes no explicit statements about AAAA RRs, similar logic applies except in cases where A and AAAA glue RR interaction requires specific consideration (response packet size, TTL consistency, namespace fragmentation).

The specification of the A6 RR [RFC2874] contains, in [section 5.1.2](#), a detailed discussion of glue issues due to the variable representation of IPv6 addresses in A6.

5.8. Open Issues

Future versions of this document might expand on these topics:

- o Software issues when following NS RRs
[[I-D.minda-dnsop-using-in-bailiwick-nameservers](#)]
- o Mixed IPv4 and IPv6 environments, following the example of
[[RFC4472](#)].
- o TTL considerations: glue data vs. authoritative data as well as NS RRSets vs A RRSets
- o Circular dependencies and "narrow" glue policy

6. Security Considerations

This section needs more work

7. IANA Considerations

This document does not request any IANA action.

The considerations of the DNS root zone glue policy are explanatory only and do not strive to change this policy. The details of ensuring that root name servers can always respond to priming queries not only with the root NS RRSets, but also a sufficient number of root name server addresses, are covered in [[I-D.ietf-dnsop-resolver-priming](#)].

[This section needs more work]

8. References

8.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC2606] Eastlake, D. and A. Panitz, "Reserved Top Level DNS Names", [BCP 32](#), [RFC 2606](#), June 1999.
- [RFC5735] Cotton, M. and L. Vegoda, "Special Use IPv4 Addresses", [RFC 5735](#), January 2010.
- [RFC5936] Lewis, E. and A. Hoenes, "DNS Zone Transfer Protocol (AXFR)", [RFC 5936](#), June 2010.

8.2. Informative References

- [I-D.ietf-dnsop-resolver-priming]
Koch, P. and M. Larson, "Initializing a DNS Resolver with Priming Queries", [draft-ietf-dnsop-resolver-priming-04](#) (work in progress), February 2014.
- [I-D.ietf-dnsop-respsize]
Vixie, P., Kato, A., and J. Abley, "DNS Referral Response Size Issues", [draft-ietf-dnsop-respsize-15](#) (work in progress), February 2014.
- [I-D.minda-dnsop-using-in-bailiwick-nameservers]
Minda, M., "Using In-Bailiwick Nameservers in .ARPA", [draft-minda-dnsop-using-in-bailiwick-nameservers-01](#) (work in progress), July 2005.
- [RFC0973] Mockapetris, P., "Domain system changes and observations", [RFC 973](#), January 1986.
- [RFC1033] Lottor, M., "Domain administrators operations guide", [RFC 1033](#), November 1987.
- [RFC1207] Malkin, G., Marine, A., and J. Reynolds, "FYI on Questions and Answers: Answers to commonly asked "experienced Internet user" questions", [RFC 1207](#), February 1991.
- [RFC1386] Cooper, A. and J. Postel, "The US Domain", [RFC 1386](#), December 1992.

- [RFC1537] Beertema, P., "Common DNS Data File Configuration Errors", [RFC 1537](#), October 1993.
- [RFC1637] Manning, B. and R. Colella, "DNS NSAP Resource Records", [RFC 1637](#), June 1994.
- [RFC1713] Romao, A., "Tools for DNS debugging", [RFC 1713](#), November 1994.
- [RFC1912] Barr, D., "Common DNS Operational and Configuration Errors", [RFC 1912](#), February 1996.
- [RFC2010] Manning, B. and P. Vixie, "Operational Criteria for Root Name Servers", [RFC 2010](#), October 1996.
- [RFC2065] Eastlake, D. and C. Kaufman, "Domain Name System Security Extensions", [RFC 2065](#), January 1997.
- [RFC2136] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", [RFC 2136](#), April 1997.
- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", [RFC 2181](#), July 1997.
- [RFC2535] Eastlake, D., "Domain Name System Security Extensions", [RFC 2535](#), March 1999.
- [RFC2672] Crawford, M., "Non-Terminal DNS Name Redirection", [RFC 2672](#), August 1999.
- [RFC2845] Vixie, P., Gudmundsson, O., Eastlake, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", [RFC 2845](#), May 2000.
- [RFC2874] Crawford, M. and C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering", [RFC 2874](#), July 2000.
- [RFC2931] Eastlake, D., "DNS Request and Transaction Signatures (SIG(0)s)", [RFC 2931](#), September 2000.
- [RFC3363] Bush, R., Durand, A., Fink, B., Gudmundsson, O., and T. Hain, "Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System (DNS)", [RFC 3363](#), August 2002.

- [RFC3364] Austein, R., "Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)", [RFC 3364](#), August 2002.
- [RFC3375] Hollenbeck, S., "Generic Registry-Registrar Protocol Requirements", [RFC 3375](#), September 2002.
- [RFC3658] Gudmundsson, O., "Delegation Signer (DS) Resource Record (RR)", [RFC 3658](#), December 2003.
- [RFC3731] Hollenbeck, S., "Extensible Provisioning Protocol (EPP) Domain Name Mapping", [RFC 3731](#), March 2004.
- [RFC3732] Hollenbeck, S., "Extensible Provisioning Protocol (EPP) Host Mapping", [RFC 3732](#), March 2004.
- [RFC3833] Atkins, D. and R. Austein, "Threat Analysis of the Domain Name System (DNS)", [RFC 3833](#), August 2004.
- [RFC3845] Schlyter, J., "DNS Security (DNSSEC) NextSECure (NSEC) RDATA Format", [RFC 3845](#), August 2004.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", [RFC 4033](#), March 2005.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", [RFC 4034](#), March 2005.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", [RFC 4035](#), March 2005.
- [RFC4183] Warnicke, E., "A Suggested Scheme for DNS Resolution of Networks and Gateways", [RFC 4183](#), September 2005.
- [RFC4472] Durand, A., Ihren, J., and P. Savola, "Operational Considerations and Issues with IPv6 DNS", [RFC 4472](#), April 2006.
- [RFC4697] Larson, M. and P. Barber, "Observed DNS Resolution Misbehavior", [BCP 123](#), [RFC 4697](#), October 2006.
- [RFC4871] Allman, E., Callas, J., Delany, M., Libbey, M., Fenton, J., and M. Thomas, "DomainKeys Identified Mail (DKIM) Signatures", [RFC 4871](#), May 2007.

- [RFC4931] Hollenbeck, S., "Extensible Provisioning Protocol (EPP) Domain Name Mapping", [RFC 4931](#), May 2007.
- [RFC4932] Hollenbeck, S., "Extensible Provisioning Protocol (EPP) Host Mapping", [RFC 4932](#), May 2007.
- [RFC5731] Hollenbeck, S., "Extensible Provisioning Protocol (EPP) Domain Name Mapping", STD 69, [RFC 5731](#), August 2009.
- [RFC5732] Hollenbeck, S., "Extensible Provisioning Protocol (EPP) Host Mapping", STD 69, [RFC 5732](#), August 2009.
- [RFC5855] Abley, J. and T. Manderson, "Nameservers for IPv4 and IPv6 Reverse Zones", [BCP 155](#), [RFC 5855](#), May 2010.

[Appendix A.](#) Extended RFC Survey

At the time of this writing, 80 RFC documents contained the term "glue", XX of which were DNS related. These latter RFCs were individually inspected to establish a historic record of the use of the term "glue", as outlined in the following subsections. Any comments are solely intended to put the respective RFC into context and to point out similarities and differences. This is not an attempt to update the documents under consideration.

[Appendix B.](#) Document Revision History

This section is to be removed should the draft be published.

\$Id: [draft-koch-dns-glue-clarifications.xml](#),v 1.12 2015/03/09 20:04:22 pk Exp \$

[B.1.](#) Changes from -04 to -05

Added appendix with extended RFC survey

[B.2.](#) Changes from -03 to -04

Adjusted boilerplate text

Added section about registry considerations

Expanded RFC survey, maintained references

[B.3.](#) Changes from -02 to -03

Added text about name server naming

Maintenance of references, minor edits

B.4. Changes from -01 to -02

Added text about grandchild glue

Maintenance of references, minor edits

B.5. Changes from -00 to -01

Mentioned RFC survey

Added text about root server glue

New text for using glue in responses

Author's Address

Peter Koch
DENIC eG
Kaiserstrasse 75-77
Frankfurt 60329
DE

Phone: +49 69 27235 0
Email: pk@DENIC.DE

