DNSEXT Working Group
INTERNET-DRAFT <draft-vixie-dnsext-dnsshadow-00.txt>
Creation Date: 2010-02-26
Intended Status: Full Standard

P. Vixie, ISC

Use of DNS to Carry Configuration Metadata Concerning Automatic Replication of Zones

Abstract

Whenever it is desireable to exactly replicated the content of a DNS zone into one or more other DNS zones so that the content is reachable by multiple names at different zone apexes, it is likewise desireable that this behaviour be automated so that cooperating primary and secondary nameservers can generate and serve the entire set of shadows without human intervention and in an open multivendor manner. This document describes a new CLONE resource record for a zone apex which can guide such cooperation.

Status of this Memo

1. Introduction

1.1. Replication of DNS content so that it is reachable by multiple names at different zone apexes requires multiple delegation NS RRsets which might be in multiple parent zones (consider VIX.COM vs. VIXIE.SF.CA.US). Parent zone administrators (including registrars and registries) and zone administrators (including registrants) can use existing tools to maintain these delegations, and the resulting NS RRsets will flow through the existing authority servers using existing mechanisms such as DNS NOTIFY and DNS IXFR for propagation.

1.2. Zones replicated in this way must propagate through a multivendor network of primary and secondary name servers, even when the replication target list changes over time (for example, more brands under trademark), in an automated fashion. New coordinated human effort across a network of authority servers every time a zone is replicated to a new target namespace should be avoided.

1.3. The term "Amber zone" will be used to describe the original zone whose content is being replicated to new namespaces under different apexes. The term "Shadow zone" will be used to describe the replicated zone content as it appear in new namespaces under different apexes.

2. Data Model

2.1. DNS RDATA for RR types not explicitly named in [RFC1035] may be opaque to all secondary and recursive servers, and to stub resolvers. Only the primary master and the far-end application are required to understand an RDATA. Since some of these RDATA may contain domain names relative to the zone apex, the replication of Amber zone data toward Shadow zones must be performed on the primary master server. Such replication must occur every time a new Shadow target becomes known, and the Shadow zones must be updated or regenerated every time the corresponding Amber zone is changed.

2.2. Shadow zone generation and replication might lag slightly behind the corresponding Amber zone, but nominally the SOA SERIAL should be identical across an Amber zone and its Shadows, and the only differences in zone content will be where relative names were used in the Amber zone's content and were therefore qualified differently in the Shadow zones. When generating a Shadow zone, the primary master will not copy the apex SHADOW RRset, in order to prevent Shadow zones from being incorrectly treated as Amber zones.

2.2.1. For example, if an Amber zone at VIX.COM has a master file which describe an apex MX RR with an unqualified MX EXCHANGE domain name such as MAIL1, then the MX target in the Amber zone will be MAIL1.VIX.COM. A Shadow of this zone whose apex is VIXIE.SF.CA.US will show this MX EXCHANGE as MAIL1.VIXIE.SF.CA.US. This may require configuration changes to supporting applications such as SMTP servers. This behaviour can be prevented by using fully qualified names wherever the name of the Amber zone, and not the name of its various Shadow zones, is to be published in the RDATA. Such prevention is likely to be important for NS NSDNAMEs whose names are within the zone itself, and where the creation of per-Shadow nameserver names is an explicit non-goal of the zone administrator.

2.3. Shadow zone content is propagaged through the authority server network using existing DNS protocols such as DNS NOTIFY and DNS IXFR, and is retrieved and consumed using existing DNS verbs such as QUERY. There are no CNAMEs. any other indication that the Shadow names are not real first class names. As a result, names within Shadow zones can be used as MX EXCHANGE names or NS NSDNAME names or anywhere else within DNS that a domain name is the target of an RDATA whose target must be a canonical name rather than an alias name.

3. Details

3.1. In the Amber zone, a SHADOW RRset will enumerate the other zone apexes at which it's desired that the zone's content be replicated. For example:

```
$ORIGIN vix.com.
 ...
@ IN SHADOW vixie.com.
@ IN SHADOW vixie.sf.ca.us.
 ...
```

This RRset will be propagated normally through the authority server network, and will thus be part of the authoritative local data for

this zone as held by the primary master server and all secondary servers. This RRset will not appear in any Shadow of this Amber zone.

3.2. The primary master server will keep track of what zones contain apex SHADOW RRsets and will treat such zones as Amber zones. Upon startup or upon any change to a SHADOW RRset, the primary master server will maintain a Shadow zone for each apex SHADOW RR in each Amber zone. Configuration of a Shadow zone will be a copy, along with any access-control or other local information, of the corresponding Amber zone's configuration. Content of a Shadow zone will be generated by parsing the Amber zone's master file using a different default $ORIGIN. DNS NOTIFY messages will be sent for each Shadow zone as and when such content generation process completes. Changes to the Amber zone's master file will cause regeneration of each associated Shadow zone. Changes to a master file that involve adding or deleting apex SHADOW RRs will cause corresponding changes to the list of Shadows of that zones.

3.3. A secondary nameserver will, upon startup and upon receiving a new version of a zone, keep track of what zones contain apex SHADOW RRsets, and will treat such as Amber zones. For each Shadow zone, the zone configuration including any access-control or other local information will be copied from the associated Amber zone. This means a Shadow zone's master server list will automatically be the same as the associated Amber zone's master server list. Changes to an Amber zone that involve adding or deleting apex SHADOW RRs will cause corresponding changes to the list of Shadows of that zones. There is no other special processing required by secondary server -- once a Shadow zone has been transferred in the normal way it will be served in the normal way, including downstream DNS NOTIFY messages if the DNS IXFR/AXFR dependency graph is deep and if this would be done for the associated Amber zone.

3.4. UPDATE messages received by the primary master server whose ZONE section or whose implied zone apex is a Shadow zone, shall be rejected with error code 9 (NOTAUTH). This is to avoid the need to modify the UPDATE message to change fully qualified names under the Shadow zone's apex to be under the Amber zone's apex instead, which would be ambiguous since some such names might be intentionally within the Shadow zone, and the update may contain new DNSSEC signatures for new or changed RRsets. An update to an Amber zone will cause regeneration of each associated Shadow zone.

3.5. If an Amber zone is signed with DNSSEC, then the signature generation process must be available within the context of the primary master server. Thus, whenever it's necessary to generate or regenerate a Shadow zone, a normal DNSSEC signing procedure will also be done on the resulting Shadow. This requires that the Shadow zones be signed online, with no offline keys or other offline processing.

3.6. Parent zones must be maintained using existing tools, and do not benefit from the new metadata described here. NS RRsets and DS RRsets must be inserted and edited through the normal communication channels used by each parent zone (which could involve action by registries, registrars, and/or registrants, if a TLD or similar shared parent zone is involved).

3.7. No changes are required for recursive nameservers, stub resolvers, or applications.

4. Open Issues

4.1. Using a different default origin and then not touching fully-qualified names is weak. It plays especially poorly with fully dynamic zone content or database back ends. We either need to tail-replace the Amber apex with each Shadow apex, or we need to add new signalling somehow.

4.2. Not allowing updates on shadows is weak. We should probably just outlaw the use of Shadow names within zone content, and do the substitution of Shadow apex by Amber apex. Note that this gets messy if the update came with its own DNSSEC metadata for the new or modified RRsets.

4.3. Doing full Shadow regeneration after each UPDATE is weak. We need to figure out some way to, um, shadow the updates. This gets messy if DNSSEC is involved and if the signer is external to the primary master server (like if there are offline keys).

4.4. The Security Considerations section is empty, which seems wrong.

5. Security Considerations

5.1. Discussion needed.

IANA Considerations

IANA would have to allocated an RR type code for SHADOW if this goes forward.

Normative References

[RFC1035] P. Mockapetris, "Domain Names - Implementation and Specification," RFC 1035, USC/Information Sciences Institute, November 1987.

Authors' Addresses

Paul Vixie (text)

Internet Systems Consortium
Redwood City, CA, USA
EMail: vixie@isc.org