academicJournals

Vol. 9(12), pp. 566-570, 30 June, 2014
DOI: 10.5897/SRE2013.5741
Article Number: 2B9B29C45695
ISSN 1992-2248 © 2014

http://www.academicjournals.org/SRE

Scientific Research and Essays

Full Length Research Paper

Seamless transition of domain name system (DNS) authoritative servers

Zheng Wang

China Organizational Name Administration Center (CONAC) Jia 31, Guangximen Beili, Xibahe, Chaoyang District, Beijing, 100028, P. R. China.

Received 18 November 2013; Accepted 26 May, 2014

The domain name system (DNS) resolution service often migrates from the one set of authoritative servers to another. The basic requirements for such transition are to ensure zero down time and minimize the transition delay. The optimum transition schemes are proposed favoring seamless and fast DNS resolution service migration. The transition of DNS authoritative servers may take place horizontally or vertically. For the horizontal case, the delegated authority is handovered from the old set of authoritative servers to the new one. For the vertical case, a zone cut is initiated from the parent zone to a newly delegated set of authoritative servers. For the DNSSEC signed zones, the DNSSEC-aware transition schemes are proposed to ensure the continuity of the trust chain. The transition delays as well as how to optimize them are discussed.

Key words: Domain name system (DNS), authoritative server, service migration

INTRODUCTION

The domain name system (DNS) is a fundamental component of the modern Internet, providing a critical link between human users and Internet routing infrastructure by mapping host names to IP addresses.

The DNS uses a tree (or hierarchical) name structure. The top of the tree is the root node followed by the top-level domains (TLDs), then the Second-Level Domains (SLD) and any number of lower levels. Each node within the domain name hierarchy is assigned to an authority - an organization or person responsible for the management and operation of that node. Such an organization or person is said to administer the node authoritatively. The authority for a particular node can in turn delegate authority for lower levels of that node within the domain name hierarchy. When a parent zone delegates part of its namespace to a child zone, the parent zone stores a list of NS resource records for the authoritative servers of the child zone. This list of NS resource records are kept both at the parent and the child zone. As shown in Figure 1, com. zone delegates example.com. zone to a child zone. The authoritative servers of the child zone are listed in a set of NS resource records. And normally the same set of NS resource records are also contained in the zone file of the child zone- example.com. zone here.

The DNS resolution service often migrates from the one set of authoritative servers to another. The basic requirements for the transition are to ensure zero down time and minimize the transition time. Service continuity is the key objective of the transition of DNS authoritative

Email: zhengwang09@126.com Tel: +86 10 52035185

```
$ORIGIN com.
example.com.    IN  NS  ns1.example.com.
example.com.    IN  NS  ns2.example.com.
example.com.    IN  NS  ns3.example.com.
$ORIGIN example.com.
example.com.    IN  NS  ns1.example.com.
example.com.    IN  NS  ns2.example.com.
example.com.    IN  NS  ns3.example.com.
```

Figure 1. DNS delegation example.

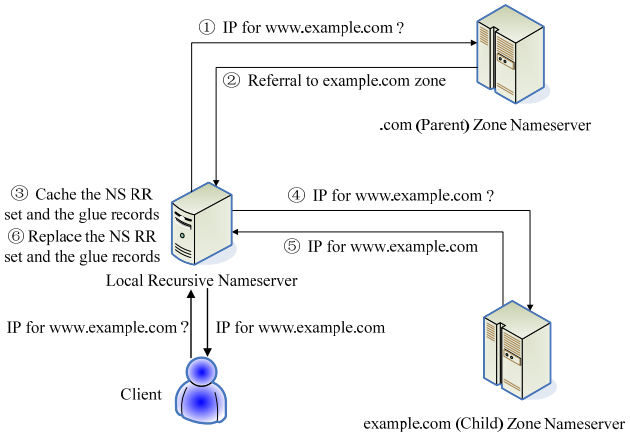

Figure 2. Recursive resolution procedure.

servers, and any resolvers should be served with response in compliance with DNS specifications during the transition. For the efficiency consideration, the transition time should be minimized in order to reduce the cost of simultaneous service of the predecessor and the successor.

The service migration problem has been addressed in the past particularly in the area of generic networks. Oikonomou and Stavrakakis (2010) proposed to determine the optimal location of a service facility in a way that is both scalable and deals inherently with network dynamicity. Shayani et al. (2010) applied techno-economic analysis to model and study the service migration between platforms. Gabner et al. (2011) investigated service component migration between the mobile client and the infrastructure-based cloud as a means to avoid service failures and improve service performance. Vanbever et al. (2011) proposed router grafting, where parts of a router are seamlessly removed from one router and merged into another, allowing a network operator to rehome a customer with no disruption. To improve the reliability and efficiency of a system in the pervasive computing domain, Cai et al. (2013) proposed a service-oriented intelligent seamless migration (SOISM) mechanism and algorithm. However, all of them cannot be directly applied to the problem of seamless transition of authoritative servers, which requires specific DNS protocol compliance.

With the introduction and deployment of DNSSEC, sustaining trust chain in parallel with the transition of DNS authoritative servers is non-trivial for the DNSSEC signed zones. An overview of challenges and potential pitfalls of DNSSEC was presented in Herzberg and Shulman (2013). Yang et al. (2011) provided a systematic examination of the design, deployment, and operational challenges encountered by DNSSEC. While key rollover was discussed as a component of DNSSEC service transition, the authoritative server transition has not been examined in combination with trust chain transition in previous works.

This work provides the following two major contributions. 1) The seamless transition of authoritative servers is analyzed and solved based on two categories: delegation transition and authority transition. 2) The DNSSEC solution for secure seamless transition of authoritative servers is proposed.

DNS RECURSIVE RESOLUTION

Figure 2 illustrates the recursive resolution procedure. The client's browser uses a resolver and queries a local recursive server for a name (say example.com). The query may miss the DNS cache in this server, that is there is no cached A records for "www.example.com". Moreover, if the NS record set for the queried domain also expired at this time (otherwise, the server can go to the authoritative server of the "example.com" zone directly), the recursive server has to request the parent zone of "example.com" by contacting the authoritative server of ".com" zone. The ".com" authoritative server answers with a referral to the servers responsible for the example.com domain. This is in the form of NS records of servers in the authority section of the DNS message. Though technically we asked only for the NS records, the servers also give us the IP address of each in the additional section of the DNS message: this is known as "glue" and is provided to avoid "query loop" and save us from having to look it up. The recursive server chooses one of the authoritative servers and sends off the same query: "what's the A record for www. example.com?". The authoritative server's reply message contains the A record in the answer section, the NS records and glue records in authority and additional section respectively.

TRANSITION SCHEME OF DNS AUTHORITATIVE SERVERS

Delegation transition

The typical transition of DNS authoritative servers is

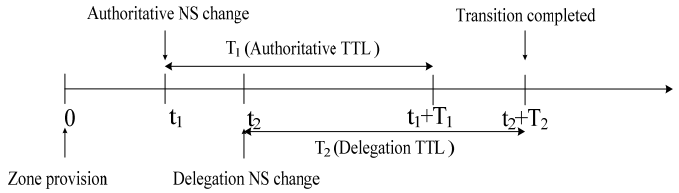

Figure 3. The time line of delegation transition.

```
$ORIGIN com.
www.example.com.  IN A   128.0.0.1
ns1.example.com.   IN   A   128.0.0.2
ns2.example.com.   IN   A   128.0.0.3
ns3.example.com.   IN   A   128.0.0.4
```

Figure 4. The original parent zone records prior to the delegation.

migration from one set of authoritative servers to another. The delegation relationship between the parent zone and the child zone does not change but the authoritative servers of the child zone changes.

According to the analysis above, the NS RRset cached by the resolvers can be the authoritative one from the child zone or the delegating one from the parent zone. This is dependent on the implementations of the child zone's name servers and the resolvers. Considering the diversity of DNS implementations, the transition mechanism should fit both cases for the guarantee of consistent service. That is, the resolvers have to wait enough time for the expiration of the old NS RRset from the cache.

Technically, migrating the delegation in the parent zone is enough for the authoritative server transition since the resolvers have to contact the parent zone for the referral information after the relevant NS RRset in the cache expires. For the child zone's nameservers which include the apex NS RRset in responses, migrating the apex NS RRset also allows for the speeding of the authoritative server transition. This is due to the possibility that some resolvers may follow the migrated apex NS RRset to reach the new authoritative server before they have the opportunity to refresh their referral information.

Let the TTL of the authoritative NS RRset in the child zone be T_1 and the TTL of the delegation NS RRset in the parent zone be T_2. Let the zone provision of the new DNS authoritative servers be launched at time 0. Let the authoritative NS RRset in the child zone be changed at t_1 and the delegation NS RRset in the parent zone be changed at t_2.

During the transition, there are three repositories of NS RRset, the parent zone, the old child zone and the new child zone. We discuss them respectively as follows.

If the old NS RRset is fetched from the parent zone by the resolvers, the resolver will get the new delegation NS RRset after t_2+T_2, which is expiration time of the old NS RRset from the cache. If the old NS RRset is fetched from the old child zone by the resolvers, the resolver will get the new authoritative NS RRset after t_1+T_1, which is expiration time of the old NS RRset from the cache.

After t_1 or t_2, the new NS RRset may be fetched from the new child zone by the resolvers. This is due to either the new delegation NS RRset at the parent zone or the new authoritative NS RRset at the old child zone.

In summary, after $\max\{t_1+T_1,t_2+T_2\}$, all resolvers have their caches refreshed by the new NS RRset. Since that time, all resolvers will not send DNS requests to the old child zone. So the Old DNS Authoritative Servers can come to the end of service. The time line is shown in Figure 3.

Authority transition

In the case of authority transition, some name space, once resolvable in the zone, is delegated as its child zone. It is the generation of a new child zone from its parent zone. Compared with the delegation transition, the migration does not move horizontally, but vertically. The original zone records related to the example.com subzone is illustrated in Figure 4. In the transition, the parent zone adds delegation records and at the same time removes all authoritative records of the delegated zone. The parent zone records and the child zone records posterior to the delegation is shown in Figure 5. If the delegated zone is provisioned at the new DNS authoritative servers prior to the delegation, any afterward requests for the delegated zone arriving at the parent zone is answered with the referral information directing to the new DNS authoritative servers. Let the zone provision of the new DNS authoritative servers start at time 0 and the delegation records adding and authoritative records of the delegated zone removing in the parent zone happen at time t_1.

DNSSEC TRANSITION SCHEME

DNSSEC provides acryptographic solution to the original DNS specifications. Public/private key pairs are used for the authentication of each zone. The public keys are stored in DNSKEY RRset, and all the sig-natures are stored in RRSIG RRset. In response to a query, an authoritative server returns both the requested data and its associated RRSIGRRset.

```
$ORIGIN com.
example.com.   IN  NS   ns1.example.com.
example.com.   IN  NS   ns2.example.com.
example.com.   IN  NS   ns3.example.com.
$ORIGIN example.com.
example.com.   IN  NS   ns1.example.com.
example.com.   IN  NS   ns2.example.com.
example.com.   IN  NS   ns3.example.com.
www.example.com.  IN A   128.0.0.1
ns1.example.com.   IN   A   128.0.0.2
ns2.example.com.   IN   A   128.0.0.3
ns3.example.com.   IN   A   128.0.0.4
```

Figure 5. The parent zone records and the child zone records posterior to the delegation.

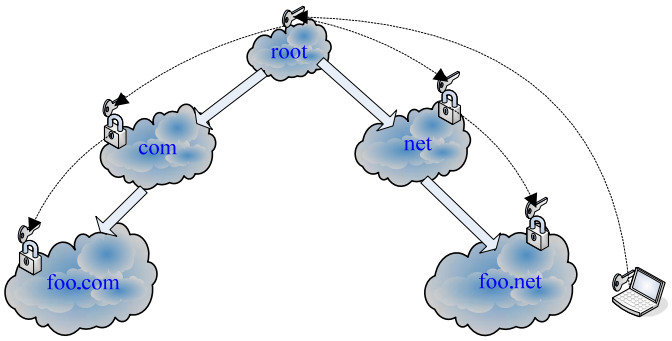

Figure 6. Record verification through the chain of trust by resolvers.

A resolver that has learned the DNSKEY of the requested zone can verify the origin, authenticity and integrity of the reply data. To resist replay attacks, each signature carries a definitive expiration time.

In order to authenticate the DNSKEY for a given zone, say www.foo.com, the resolver needs to construct a chain of trust that follows the DNS hierarchy from a trusted root zone key down to the key of the zone in question (this is shown in Figure 6). In the ideal case, the public key of the DNS root zone would be obtained offline in a secure way and stored at the resolver, so that the resolver can use it to authenticate the public key of com.; the public key of com. would then be used to authenticate the public key of foo.com.

A parent zone must encode the authentication of each of its child zone's public keys in the DNS. To accomplish this, the parent zone creates and signs a Delegation Signer (DS) RR that corresponds to a DNSKEYRR at the child zone, and creates an authentication link from the parent to child. It is the child zone's responsibility to request an update to the DSRR every time the child's DNSKEY changes.

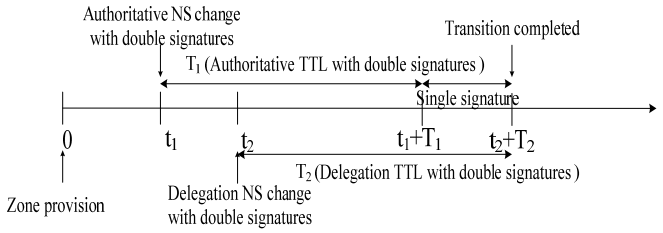

Figure 7. The time line of DNSSEC-aware delegation transition.

For the signed zone, the main objective of transition is to maintain any records verifiable through the chain of trust.

DNSSEC-aware delegation transition

In the case of authority transition, some name space, once resolvable in the zone, is delegated as its child zone. It is the generation of a new child zone from its parent zone. Compared with the delegation transition, the migration does not move horizontally, but vertically. The original zone records related to the example.com subzone is illustrated in Figure 4. The delegation relationship between the parent zone and the child zone does not change but the authoritative servers of the child zone changes.

When the authoritative NS records or the delegation NS records change, their signatures or RRSIG records should be generated by resigning the NS records with the DNSKEY. But the previously fetched NS records are still retained in the cache until they expire from the cache according to the TTL. The resolver follows the cached NS records to request the old authoritative servers for the NS records' signatures. This makes it necessary to maintain the signatures of old NS records together with those of the new ones. Otherwise, the cached NS records would lose their signatures and fail verifications because the replied RRSIG RR set only contains the RRSIG for the new NS records. Therefore, the double signatures, for both new and old NS records, should be kept in the old authoritative servers and parent servers for the TTL of the NS records. However, for the new authoritative servers, the double signature scheme is unnecessary. When a query is sent to the new authoritative servers, the cached NS records in the resolver must be the ones referring to the new authoritative servers. So the signature for the new NS records is enough for the successful verification. The time line is shown in Figure 7.

DNSSEC-aware authority transition

Compared with DNSSEC-oblivious authority transition,

the zone provision of the new DNS authoritative servers should be conducted along with the submission of DS records to the parent zone and zone signing. The parent zone should add the submitted DS records and sign them with its DNSSEY to establish a chain of trust linking the new child zone.

TRANSITION DELAY AND ITS OPTIMIZATION

For the delegation transition, the transition delay is max{ t_1+T_1, t_2+T_2}. To accelerate the transition, the authoritative NS RRset in the child zone and the delegation NS RRset in the parent zone should change as soon as possible after the zone provision of the new DNS authoritative servers. Minimizing the TTL of the authoritative NS RRset in the child zone and the TTL of the delegation NS RRset in the parent zone also helps to speed up the transition.

For the authority transition, the transition delay is t_1. To accelerate the transition, the authoritative NS RRset in the child zone and the delegation NS RRset in the parent zone should change as soon as possible after the zone provision of the new DNS authoritative servers. Minimizing the TTL of the authoritative NS RRset in the child zone and the TTL of the delegation NS RRset in the parent zone also helps to speed up the transition.

For the DNSSEC-aware delegation transition, the transition delay is max{ t_1+T_1, t_2+T_2}, which determines the resolution service duration of the old authoritative servers. Note that if $t_1+T_1<t_2+T_2$, a single signature time window emerges for the old authoritative severs. In that period, the old authoritative severs only need to keep the RRSIG records for the new NS records because the old NS records have already expired from the cache. But the old authoritative severs should continue its resolution service because the old delegation NS records have not expired from the cache. In the scenario, the resolver may still send its queries to the old authoritative servers following the references in its cache (the old delegation NS records). So if the old authoritative severs is unresponsive, the queries will get the failure response. For the DNSSEC-aware authority transition, the transitiondelay is t_1.

Conclusions

DNS operators are under increasing pressure to make their resolution service highly reliable and continuous to avoid service disruptions. But operators often need to change the authoritative servers to upgrade faulty equipment, deploy new servers, or transfer services. Unfortunately, unexamined authoritative server changes may cause disruptions. In this paper, seamless transition schemes are presented allowing an operator to migrate DNS authoritative servers with no disruption. In addition, the transition schemes are examined in the DNSSEC cases aiming at sustaining trust chains. The transition delay and its optimization are discussed.

Conflict of interests

The author(s) have not declared any conflict of interests.

ACKNOWLEDGEMENTS

This work was supported by the National Key Technology R & D Program of China (No. 2012BAH16B00), the National Science Foundation of China (No. 61003239), Beijing Institution of Higher Learning “Young Talents Plan”, Beijing Natural Science Foundation (No. 4144084), and the National Science Foundation for Young Scholars of China (No. 61102057).

REFERENCES

Cai H, Chao P, Robert HD, Linhua J (2013). A novel service-oriented intelligent seamless migration algorithm and application for pervasive computing environments. Future Gener. Comput. Syst. 29(8):1919-1930. http://dx.doi.org/10.1016/j.future.2013.02.008

Gabner R, Hans-Peter S, Karin AH, Gunter H (2011). Optimal Model-Based Policies for Component Migration of Mobile Cloud Services. In Proceedings of the 2011 IEEE 10th International Symposium on Network Computing and Applications (NCA '11). pp.195-202. http://dx.doi.org/10.1109/NCA.2011.33

Herzberg A, Shulman H (2013). DNSSEC: Interoperability Challenges and Transition Mechanisms. In Proceedings of the 2013 International Conference on Availability, Reliability and Security (ARES '13). pp. 398-405. http://dx.doi.org/10.1109/ARES.2013.53; PMid:23186151

Oikonomou K, Stavrakakis I (2010). Scalable service migration in autonomic network environments. IEEE J. Sel. A. Commun. 28(1):84-94. http://dx.doi.org/10.1109/JSAC.2010.100109

Shayani D, Machuca CM, Jager M (2010). A techno-economic approach to telecommunications: the case of service migration. IEEE Trans. on Netw. Serv. Manage. 7(2):96-106.

Vanbever L, Stefano V, Cristel P, Pierre F, Olivier B (2011). Seamless network-wide IGP migrations. SIGCOMM Comput. Commun. Rev. 41(4):314-325. http://dx.doi.org/10.1145/2043164.2018473

Yang H, Eric O, Dan M, Songwu L, Lixia Z (2011). Deploying Cryptography in Internet-Scale Systems: A Case Study on DNSSEC. IEEE Trans. Dependable Secur. Comput. 8(5):656-669. http://dx.doi.org/10.1109/TDSC.2010.10