IPv6 Operations (v6ops)
Internet-Draft
Intended status: Standards Track
Expires: September 11, 2019

Martin Hunek
Technical University of Liberec

March 11, 2019

NAT64/DNS64 detection via SRV Records
draft-ietf-v6ops-nat64-srv-00

Abstract

This document specifies the way of discovering the NAT64 pools in use as well as DNS servers providing DNS64 service to the local clients. The discovery is done via SRV records, which also allows asignment of priorities to the NAT64 pools as well as DNS64 servers. It also allows clients to have diferent DNS providers than NAT64 provider, while providing a secure way via DNSSEC validation of provided SRV records. This way, it provides DNS64 service even in case where DNS over HTTPS is used.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The simultaneous use of NAT64/DNS64 and DNSSEC outlined by [RFC7050], does not solve all the aspects of such use. Namely [RFC7050] does not allow assignment of NAT64 priorities in case when multiple network prefixes are in use. [RFC7050] also doesn’t work in the case when network operator and DNS operator are not the same subject, like in the case when the end node is using some public DNS resolvers. This document describes the way how to circumvent that limitation while maintaining added security provided by DNSSEC.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Terminology

End node: Either DNS stub resolver or the DNS recursive resolver serving a local area network or station.

Pref64::/n: an IPv6 prefix used for IPv6 address synthesis [RFC6146].

Pref64::WKA: an IPv6 address consisting of Pref64::/n and WKA at any of the locations allowed by [RFC6052].

Well-Known IPv4 Address (WKA): an IPv4 address that is well-known and present in an A record for the well-known name. Two well-known IPv4 addresses are defined for Pref64::/n discovery purposes: 192.0.0.170 and 192.0.0.171.

3. NAT64 service SRV record

This document specifies two new well-known SRV records. The one for NAT64 prefix which validation end node MUST implement:

```
nat64. ipv6.Name TTL Class SRV Priority Weight Port Target
```

The TTL, Class, Priority and Weight follows the same scheme as defined in [RFC2782] and have theirs standard meaning.

Port: IPv6 as L3 protocol doesn’t use port numbers. Because of that this field SHOULD be either set to zero, or SHOULD be used to indicate length of network prefix mask in both IPv6 and IPv4 protocol, used NAT64. In such case the port 16b integer MUST be constructed by directly appending IPv4 pool prefix mask after the

IPv6 prefix mask decadicaly. Usually this would mean 9632 stating that IPv6 prefix with mask /96 is translated into single IPv4 address (/32).

Target: MUST point to AAAA record formed from Pref64::/n prefix and WKA same way as in [RFC7050] (Pref64::WKA). Target MAY also point to A record, in which case it SHOULD point to IPv4 address used for NAT64 (or base address of the NAT64 IPv4 prefix).

4. DNS64 service SRV record

The second SRV record is for the discovery of DNS64 service. Support of this record is OPTIONAL but end node SHOULD implement it.

```
dns64.Protocol.Name TTL Class SRV Priority Weight Port Target
```

Record informs about location of DNS64 service. This might be used in case that network operator doesn't want to deploy DNS64 in their main DNS infrastructure. A DNS64 SRV record follows the rules specified by [RFC2782] and does not modify meaning of any field.

Server provided by this record SHOULD only be used for domain names which have returned NODATA for AAAA record.

5. Node Behavior

In early stage of end node connection to the network - after the end node is configured with IP address, the end node MUST get local domains used in the network. Method of obtaining such information is out of scope of this document, but it might contain one or more methods, like the SLAAC-DNSSL [RFC8106], the DHCPv6 - option 24 or a manual configuration. In case, when no local domain can be discovered, the end node SHOULD continue NAT64/DNS64 detection by other means, like [RFC7050].

After the list of local domains has been established, the end node MUST ask for NAT64 SRV record for every domain in the list. Result of such queries SHOULD be ordered by following the rules of [RFC2782]. In case when multiple records do have a same values of both priority and weight, the records SHOULD maintain the same order as its domain in the discovered domain list.

For every domain with NAT64 SRV record the end node SHOULD perform query for DNS64 SRV record. If such a record is obtained and the end node is not configured to make DNS64 synthesis itself, the end node SHOULD use preferred target of DNS64 SRV record to query for FQDN without AAAA record - when it received NODATA response.

If the end node is capable of validation of DNS records via DNSSEC, the end node MUST perform validation of NAT64/DNS64 SRV record. Default behavior of end node SHOULD be to ignore any NAT64/DNS64 SRV records which cannot be validated or did not pass the validation.

5.1. Example

The end node is a home router connected to the ISP network in which the NAT64/DNS64 is used and the ISP has the following SRV records in their zones:

```
- nat64.ipv6.example.com. IN SRV 5 10 9632 nat64-pool-1.example.com.
- nat64-pool-1.example.com. IN AAAA 2001:db8:64:ff9b:1::c000:aa
- nat64-pool-1.example.com. IN A 192.0.2.64
- nat64.ipv6.example.com.
                         IN SRV 10 10 9632 nat64-pool-2.example.com.
- nat64-pool-2.example.com. IN AAAA 2001:db8:64:ff9b:2::c000:aa
- nat64-pool-2.example.com. IN A 192.0.2.164
- nat64.ipv6.example.net. IN SRV 10 10 9624 nat64-pool.example.net.
- nat64-pool.example.net. IN AAAA 2001:db8:64:ff9b:abc::c000:aa
- nat64-pool.example.net. IN A 198.51.100.0
- nat64.ipv6.example.invalid.
                         IN SRV 10 10 9624 nat64-pool.example.org.
- nat64-pool.example.org. IN AAAA 2001:db8:64:ff9b:def::c000:aa
- nat64-pool.example.org. IN A 203.0.113.0
```

In addition the zones "example.net" and "example.invalid" has got DNS64 SRV records:

```
- dns64.tcp.example.net. IN SRV 5 10 53 dns64.example.net.
- dns64.udp.example.net. IN SRV 10 10 53 dns64.example.net.
- dns64.example.net. IN AAAA 2001:db8::53
- dns64.udp.example.invalid. IN SRV 10 10 53 dns64.example.org.
- dns64.example.org. IN AAAA 2001:db8:123::53
```

The zones "example.com" and "example.net" are secured and successfully validated by the DNSSEC. Domain "example.invalid" is either not secured by the DNSSEC or its validation failed. Domain "example.org" is DNSSEC secured but does not have any NAT64/DNS64 SRV records.

The end node has been supplied with the following list of domains via SLAAC-DNSSL:

1. example.net
2. example.invalid
3. example.com
4. example.org

The end node would fetch all available SRV records and its A and AAAA counterparts and sort it in following order:

pool	DNSSEC	priority	reason
nat64-pool-1.example.com.	yes	5	lowest priority field
nat64-pool.example.net.	yes	10	discovered first
nat64-pool-2.example.com.	yes	10	higher priority field
nat64-pool.example.org.	no	10	no valid DNSSEC chain

After sorting, the end node SHOULD graylist any record which cannot be validated by the DNSSEC. In this example it would be "nat64-pool.example.org." because it has been obtained from insecure domain "example.invalid". A such pool SHOULD NOT be used if it is not confirmed by other DNSSEC secured record.

If the end node is capable to act as recursive or caching DNS server and it is configured to provide the DNS64 service, it MUST provide this service using sorted list of NAT64 pools. For such end node a process of the NAT64/DNS64 ends here.

However, when the end node is not capable of record synthesis or it is not configured to provide DNS64 service, it SHOULD perform detection of DNS64 by querying for "ipv4only.arpa" like in the case of [RFC7050]. If the reply contains a pool listed in the NAT64 pool list, the corresponding entry is marked as having DNS64 provided by recursive DNS.

When the end node supports DNS64 SRV record and there is at least one non-graylisted NAT64 pool, which is not reachable by using the end node’s recursive DNS, the end node MUST make a sorted list of DNS64 servers from the DNS64 SRV records. The DNS64 sorted list would look like this:

server	proto	DNSSEC	priority	reason
dns64.example.net.	tcp	yes	5	lowest priority field
dns64.example.net.	udp	yes	10	higher priority field
dns64.example.org.	udp	no	10	no valid DNSSEC chain

Sorting is done in the same fashion as any other SRV record with the same exception of graylisting records without valid DNSSEC chain. Those SHOULD NOT be used when not confirmed by DNSSEC validated record and SHOULD be kept in the end of the list.

For example when ISP is providing DNS64 service in their main DNS infrastructure only for pools in the domains "example.com" and "example.org" and the pool "nat64-pool.example.net" is used only with corresponding DNS64 server. The final sorted list of NAT64 prefixes used by the end node in the ISP network would be:

pool	state	priority	reason
nat64-pool-1.example.com.	active	5	lowest priority field
nat64-pool-2.example.com.	backup	10	higher priority field
nat64-pool.example.net.	backup*	10	main DNS has priority
nat64-pool.example.org.	inactive	10	no valid DNSSEC chain

As the pool "nat64-pool.example.net" is used only with the server "dns64.example.net" this would effectively put this pool to the end of the list. Because it would be used only for FQDN for which the regular DNS infrastructure returns NODATA.

Now the end node has successfully identified NAT64 pools and the DNS64 servers in the ISP infrastructure. The discovered prefixes SHOULD be considered safe and DNSSEC validation of answers in these prefixes MUST be either disabled or performed by validating only the suffix.

6. IANA Considerations.

This document proposes a usage of "ipv6" in Proto field and two services "nat64" and "dns64" in Service field of SRV RR ([RFC2782]).

7. Security considerations

Method proposed by this document relies on security principles based on DNSSEC and secure discovery of local domain. In order to be secure, the network operator MUST deploy DNSSEC on at least one domain (advertised to end node) and establish secure channel to this advertisement.

8. References

8.1. Normative References

[RFC2119] S. Bradner. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119. RFC Editor, Mar. 1997, pp. 1-3. url: https://www.rfc-editor.org/rfc/rfc2119.txt.

[RFC2782] A. Gulbrandsen, P. Vixie, and L. Esibov. A DNS RR for specifying the location of services (DNS SRV). RFC 2782. RFC Editor, Feb. 2000, pp. 1-12. url: https://www.rfc-editor.org/rfc/rfc2782.txt.

[RFC6146] M. Bagnulo, P. Matthews, and I. van Beijnum. Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers. RFC 6146. RFC Editor, Apr. 2011, pp. 1-45. url: https://www.rfc-editor.org/rfc/rfc6146.txt.

[RFC7050] T. Savolainen, J. Korhonen, and D. Wing. Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis. RFC 7050. RFC Editor, Nov. 2013, pp. 1-22. url: https://www.rfc-editor.org/rfc/rfc7050.txt.

[RFC8174] B. Leiba. Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words. RFC 8174. RFC Editor, May 2017, pp. 1-4. url: https://www.rfc-editor.org/rfc/rfc8174.txt.

8.2. Informative References

[RFC6052] C. Bao et al. IPv6 Addressing of IPv4/IPv6 Translators. RFC 6052. RFC Editor, Oct. 2010, pp. 1–18. url: https://www.rfc-editor.org/rfc/rfc6052.txt.

[RFC8106] J. Jeong et al. IPv6 Router Advertisement Options for DNS Configuration. RFC 8106. RFC Editor, Mar. 2017, pp. 1–19. url: https://www.rfc-editor.org/rfc/rfc8106.txt.

Acknowledgments

This work has been supported by Student Grant Scheme (SGS 2019) at Technical University of Liberec.

Authors’ Addresses

Martin Hunek
Technical University of Liberec
Studentska 1402/2
Liberec, 46017 Czech Republic

phone: +420 485 35 3792
e-mail: martin.hunek@tul.cz