

Behave WG
Internet-Draft
Intended status: Standards Track
Expires: October 06, 2013

T. Savolainen
Nokia
J. Korhonen
Nokia Siemens Networks
D. Wing
Cisco Systems
April 04, 2013

Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis
draft-ietf-behave-nat64-discovery-heuristic-17.txt

Abstract

This document describes a method for detecting the presence of DNS64 and for learning the IPv6 prefix used for protocol translation on an access network. The method depends on the existence of a well-known IPv4-only fully qualified domain name "ipv4only.arpa". The information learned enables nodes to perform local IPv6 address synthesis and to potentially avoid NAT64 on dual-stack and multi-interface deployments.

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 October 06, 2013.

Copyright Notice

Copyright (c) 2013 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

Internet-Draft

Pref64::/n Discovery

April 2013

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.

Table of Contents

1.	Introduction	2
2.	Requirements and Terminology	3
2.1.	Requirements	3
2.2.	Terminology	3
3.	Node Behavior	4
3.1.	Validation of Discovered Pref64::/n	6
3.1.1.	DNSSEC Requirements for the Network	6
3.1.2.	DNSSEC Requirements for the Node	7
3.2.	Connectivity Check	8
3.2.1.	No Connectivity Checks Against ipv4only.arpa	9
3.3.	Alternative Fully Qualified Domain Names	9
3.4.	Message Flow Illustration	10
4.	Operational Considerations for Hosting the IPv4-Only Well-Known Name	11
5.	Operational Considerations for DNS64 Operator	12
5.1.	Mapping of IPv4 Address Ranges to IPv6 Prefixes	12
6.	Exit Strategy	13
7.	Security Considerations	13
8.	IANA Considerations	14
8.1.	Domain Name Reservation Considerations	14
8.2.	IPv4 Address Allocation Considerations	15
8.3.	IAB Statement Regarding This .arpa Request	16
9.	Acknowledgements	16
10.	References	16
10.1.	Normative References	16
10.2.	Informative References	17
Appendix A.	Example of DNS Record Configuration	18
Appendix B.	About the IPv4 Address for the Well-Known Name	19
Authors' Addresses		19

[1.](#) Introduction

As part of the transition to IPv6, NAT64 [[RFC6146](#)] and DNS64 [[RFC6147](#)] technologies will be utilized by some access networks to provide IPv4 connectivity for IPv6-only nodes [[RFC6144](#)]. DNS64

utilizes IPv6 address synthesis to create local IPv6 addresses for peers having only IPv4 addresses, hence allowing DNS-using IPv6-only nodes to communicate with IPv4-only peers.

However, DNS64 cannot serve applications not using DNS, such as those receiving IPv4 address literals as referrals. Such applications could nevertheless be able to work through NAT64, provided they are able to create locally valid IPv6 addresses that would be translated to the peers' IPv4 addresses.

Additionally, DNS64 is not able to do IPv6 address synthesis for nodes running validating DNSSEC-enabled DNS resolvers, but instead the synthesis must be done by the nodes themselves. In order to perform IPv6 synthesis, nodes have to learn the IPv6 prefix(es) used on the access network for protocol translation. A prefix, which may be a Network Specific Prefix (NSP) or Well-Known Prefix (WKP) [[RFC6052](#)], is referred to in this document as Pref64::/n [[RFC6146](#)].

This document describes a best-effort method for applications and nodes to learn the information required to perform local IPv6 address synthesis. The IPv6 address synthesis procedure itself is out-of-scope of this document. An example application is a browser encountering IPv4 address literals in an IPv6-only access network. Another example is a node running a validating security-aware DNS resolver in an IPv6-only access network.

The knowledge of IPv6 address synthesis taking place may also be useful if DNS64 and NAT64 are used in dual-stack enabled access networks or if a node is multi-interfaced [[RFC6418](#)]. In such cases, nodes may choose to prefer IPv4 or an alternative network interface in order to avoid traversal through protocol translators.

It is important to note that use of this approach will not result in a system that is as robust, secure, and well-behaved as an all-IPv6 system would be. Hence it is highly recommended to upgrade nodes' destinations to IPv6 and utilize the described method only as a transition solution.

[2.](#) Requirements and Terminology

[2.1.](#) Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

[2.2.](#) Terminology

NAT64 FQDN: a fully qualified domain name (FQDN) for a NAT64 protocol translator.

Pref64::[/n](#): a 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 [RFC 6052](#) [[RFC6052](#)].

Secure Channel: a communication channel a node has between itself and a DNS64 server protecting DNS protocol related messages from interception and tampering. The channel can be, for example, IPsec-based virtual private network (VPN) tunnel or a link layer utilizing data encryption technologies.

Well-Known IPv4-only Name (WKN): the fully qualified domain name, "ipv4only.arpa", well-known to have only A record(s).

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.](#) Node Behavior

A node requiring information about the presence (or absence) of NAT64, and one or more Pref64::[/n](#) used for protocol translation SHALL send a DNS query for AAAA resource records of the Well-Known IPv4-only Name (WKN) "ipv4only.arpa". The node MAY perform the DNS query in both IPv6-only and dual-stack access networks.

When sending a DNS AAAA resource record query for the WKN, a node MUST set the "Checking Disabled (CD)" bit to zero [[RFC4035](#)], as otherwise the DNS64 server will not perform IPv6 address synthesis ([Section 3 of \[RFC6147\]](#)) and hence would not reveal the Pref64::[/n](#)

used for protocol translation.

A DNS reply with one or more AAAA resource records indicates that the access network is utilizing IPv6 address synthesis. In some scenarios captive portals, or NXDOMAIN and NODATA hijacking, performed by the access network may result in a false positive. One method to detect such hijacking is to query a fully qualified domain name that is known to be invalid (and normally return an empty response or an error response) and see if it returns a valid resource record. However, as long as the hijacked domain does not result in AAAA resource record responses that contain well-known IPv4 address in any location defined by [RFC6052](#), the response will not disturb the Pref64::

A node MUST look through all of the received AAAA resource records to collect one or more Pref64::RFC6052] or one or more Network-Specific Prefixes. In the case of NSPs, the node SHALL determine the used address format by searching the received IPv6 addresses for the

WKN's well-known IPv4 addresses. The node SHALL assume the well-known IPv4 addresses might be found at the locations specified by [\[RFC6052\] section 2.2](#). The node MUST check on octet boundaries to ensure a 32-bit well-known IPv4 address value is present only once in an IPv6 address. In case another instance of the value is found inside the IPv6 address, the node SHALL repeat the search with the other well-known IPv4 address.

If only one Pref64::

If more than one Pref64::

If the well-known IPv4 addresses are not found within the standard

locations, it indicates that the network is not using a standard address format or unexpected IPv4 addresses were used in the AAAA resource record synthesis. In either case, the Pref64::

In case a node does not receive a positive DNS reply to the AAAA resource record query, the node MAY perform a DNS A resource record query for the well-known name. If the node receives a positive reply to the DNS A resource record query it means the used recursive DNS server is not a DNS64 server.

In the case of a negative response (NXDOMAIN, NODATA) or a DNS query timeout: a DNS64 server is not available on the access network, the access network filtered out the well-known query, or something went wrong in the DNS resolution. All unsuccessful cases result in a node being unable to perform local IPv6 address synthesis. In the case of timeout, the node SHOULD retransmit the DNS query like any other DNS query the node makes [[RFC1035](#)]. In the case of a negative response (NXDOMAIN, NODATA), the node MUST obey the Time-To-Live [[RFC1035](#)] of the response before resending the AAAA resource record query. The node MAY monitor for DNS replies with IPv6 addresses constructed from

the WKP, in which case if any are observed the node SHOULD use the WKP as if it were learned during the query for the well-known name.

To save Internet resources if possible, a node should perform Pref64::

[3.1.](#) Validation of Discovered Pref64::

If a node is using an insecure channel between itself and a DNS64 server, or the DNS64 server is untrusted, it is possible for an

attacker to influence the node's Pref64::

To mitigate against attacks, the node SHOULD communicate with a trusted DNS64 server over a secure channel, or use DNSSEC. NAT64 operators SHOULD provide facilities for validating discovery of Pref64::

It is important to understand that DNSSEC only validates that the discovered Pref64::

3.1.1. DNSSEC Requirements for the Network

If the operator has chosen to support nodes performing validation of discovered Pref64::

1. Have one or more fully qualified domain names for the NAT64 translator entities (later referred as NAT64 FQDN). In the case of more than one Pref64::- 2. Each NAT64 FQDN MUST have one or more DNS AAAA resource records containing Pref64::WKA (Pref64::

3. Each Pref64::WKA MUST have a PTR resource record that points to the corresponding NAT64 FQDN.
4. Sign the NAT64 FQDNs' AAAA and A resource records with DNSSEC.

3.1.2. DNSSEC Requirements for the Node

A node SHOULD prefer a secure channel to talk to a DNS64 server,

whenever possible. In addition, a node that implements a DNSSEC validating resolver MAY use the following procedure to validate discovery of the Pref64::

1. Heuristically find Pref64::Section 3. This will result in IPv6 addresses consisting of Pref64::- 2. Send a DNS PTR resource record query for the IPv6 address of the translator (for "ip6.arpa"), using the Pref64::WKA learned in step 1. CNAME and DNAME results should be followed according to the rules in [RFC 1034 \[RFC1034\]](#), [RFC 1034 \[RFC1035\]](#), and [RFC 6672 \[RFC6672\]](#). The ultimate response will include one or more NAT64 FQDNs.
- 3. The node SHOULD compare the domains of learned NAT64 FQDNs to a list of the node's trusted domains and choose a NAT64 FQDN that matches. The means for a node to learn the trusted domains is implementation-specific. If the node has no trust for the domain, the discovery procedure is not secure and the remaining steps described below MUST NOT be performed.
- 4. Send a DNS AAAA resource record query for the NAT64 FQDN.
- 5. Verify the DNS AAAA resource record contains Pref64::WKA addresses received at the step 1. It is possible that the NAT64 FQDN has multiple AAAA records, in which case the node MUST check if any of the addresses match the ones obtained in step 1. The node MUST ignore other responses and not use them for local IPv6 address synthesis.
- 6. Perform DNSSEC validation of the DNS AAAA response.

After the node has successfully performed the above five steps, the node can consider Pref64::

After learning a Pref64::

There are two main approaches to determine if the learned Pref64::

The node SHOULD use an implementation-specific connectivity check server and a protocol of the implementation's choice, but if that is not possible, a node MAY do a PTR resource record query of the Pref64::WKA to get a NAT64 FQDN. The node then does an A resource query of the NAT64 FQDN, which will return zero or more A resource records pointing to connectivity check servers used by the network operator. A negative response to the PTR or A resource query means there are no connectivity check servers available. A network operator that provides NAT64 services for a mix of nodes with and without implementation-specific connectivity check servers SHOULD assist nodes in their connectivity checks by mapping each NAT64 FQDN to one or more DNS A resource records with IPv4 address(es) pointing to connectivity check server(s). The Pref64::

In case of multiple connectivity check servers being available for use, the node chooses the first one, preferring implementation-specific servers.

The connectivity check protocol used with implementation-specific connectivity check servers is implementation-specific.

The connectivity check protocol used with connectivity check servers pointed to by the NAT64 FQDN's A resource records is ICMPv6 [[RFC4443](#)]. The node performing a connectivity check against these servers SHALL send an ICMPv6 Echo Request to an IPv6 address synthesized by combining discovered Pref64::RFC6052]. This will test the IPv6 path to the NAT64, the NAT64's operation, and the IPv4 path all the

way to the connectivity check server. If no response is received for the ICMPv6 Echo Request, the node SHALL send another ICMPv6 Echo Request, a second later. If still no response is received, the node SHALL send a third ICMPv6 Echo Request two seconds later. If an ICMPv6 Echo Response is received, the node knows the IPv6 path to the connectivity check server is functioning normally. If, after the three transmissions and three seconds since the last ICMPv6 Echo Request, no response is received, the node learns this Pref64::/n might not be functioning, and the node MAY choose a different Pref64::/n (if available), choose to alert the user, or proceed anyway assuming the failure is temporary or with the connectivity check itself. After all, the ICMPv6 is by design unreliable and failure to receive ICMPv6 responses may not indicate anything other than network failure to transport ICMPv6 messages through.

If no separate connectivity check is performed before local IPv6 address synthesis, a node MAY monitor success of connection attempts performed with locally synthesized IPv6 addresses. Based on success of these connections, and based on possible ICMPv6 error messages received (such as Destination Unreachable messages), the node MAY cease to perform local address synthesis and MAY restart the Pref64::/n discovery procedures.

[3.2.1.](#) No Connectivity Checks Against ipv4only.arpa

Clients MUST NOT send a connectivity check to an address returned by the ipv4only.arpa query. This is because, by design, no server will be operated on the Internet at that address as such. Similarly, network operators MUST NOT operate a server on that address. The reason this address isn't used for connectivity checks is that operators who neglect to operate a connectivity check server will allow that traffic towards the Internet where it will be dropped and cause a false negative connectivity check with the client (that is, the NAT64 is working fine, but the connectivity check fails because a server is not operating at "ipv4only.arpa" on the Internet and a server is not operated by the NAT64 operator). Instead, for the connectivity check, an additional DNS resource record is looked up and used for the connectivity check. This ensures that packets don't unnecessarily leak to the Internet and reduces the chance of a false negative connectivity check.

[3.3.](#) Alternative Fully Qualified Domain Names

Some applications, operating systems, devices, or networks may find it advantageous to operate their own DNS infrastructure to perform a function similar to "ipv4only.arpa", but using a different resource

record. The primary advantage is to ensure availability of the DNS infrastructure and ensure the proper configuration of the DNS record

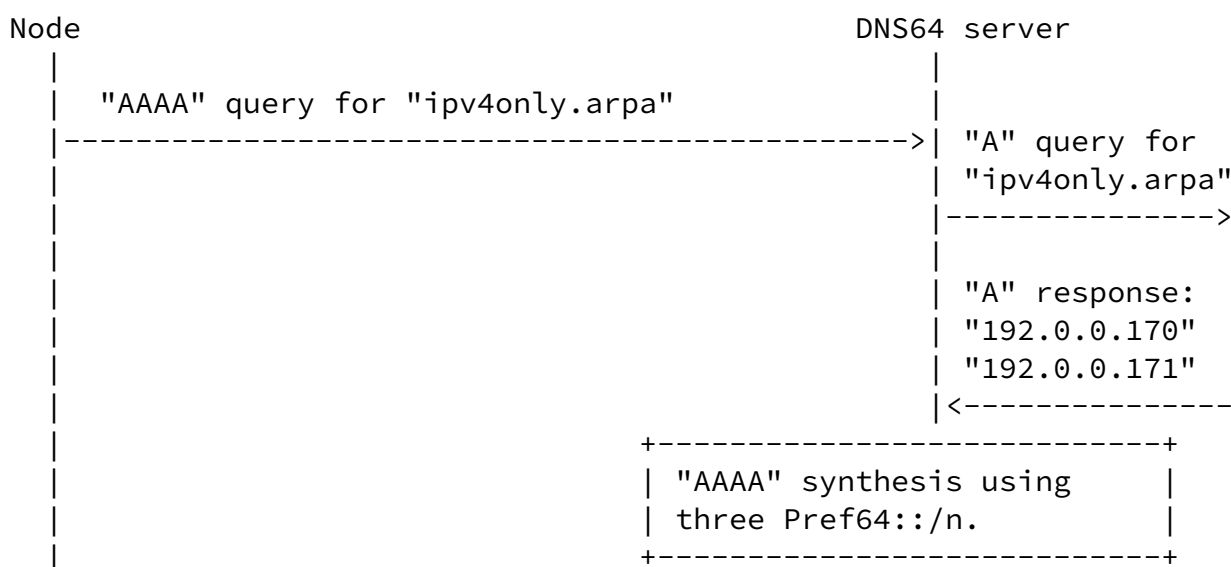
itself. For example, a company named Example might have their application query "ipv4only.example.com". Other than the different DNS resource record being queried, the rest of the operations are anticipated to be identical to the steps described in this document.

3.4. Message Flow Illustration

The figure below gives an example illustration of a message flow in the case of prefix discovery utilizing Pref64::

In this example, three Pref64::

The validation is not done for the WKP, see [Section 3.1](#).



```

| "AAAA" response with:
| "2001:db8:42::192.0.0.170"
| "2001:db8:43::192.0.0.170"
| "64:ff9b::192.0.0.170"
| <-----
+-----+
| If Pref64::/n validation is not performed, a
| node can fetch prefixes from AAAA responses
|

```

```

| at this point and skip the steps below.
+-----+
|
| "PTR" query #1 for "2001:db8:42::192.0.0.170"
| ----->
| "PTR" query #2 for "2001:db8:43::192.0.0.170"
| ----->
|
| "PTR" response #1 "nat64_1.example.com"
| <-----
| "PTR" response #2 "nat64_2.example.com"
| <-----
|
+-----+
| Compare received domains to a trusted domain
| list and if matches are found, continue.
+-----+
|
| "AAAA" query #1 for "nat64_1.example.com"
| ----->
| "AAAA" query #2 for "nat64_2.example.com"
| ----->
|
| "AAAA" resp. #1 with "2001:db8:42::192.0.0.170"
| <-----
| "AAAA" resp. #2 with "2001:db8:43::192.0.0.170"
| <-----
|
+-----+
| Validate AAAA responses and compare the IPv6
| addresses to those previously learned.
+-----+

```

```

|
+-----+
| Fetch the Pref64::

```

Pref64::

4. Operational Considerations for Hosting the IPv4-Only Well-Known Name

The authoritative name server for the well-known name SHALL have DNS record Time-To-Live (TTL) set to at least 60 minutes in order to improve effectiveness of DNS caching. The exact TTL value will be determined and tuned based on operational experiences.

The domain serving the well-known name MUST be signed with DNSSEC. See also [section 7](#).

5. Operational Considerations for DNS64 Operator

A network operator of a DNS64 server can guide nodes utilizing heuristic discovery procedures by managing the responses a DNS64 server provides.

If the network operator would like nodes to utilize multiple Pref64::Section 3 the nodes can then use them all.

There are no guarantees on which of the Pref64::

Besides choosing how many Pref64::

5.1. Mapping of IPv4 Address Ranges to IPv6 Prefixes

[RFC 6147](#) [[RFC6147](#)] allows DNS64 implementations to be able to map specific IPv4 address ranges to separate Pref64::RFC5735]. The example setup where this might be used is illustrated in Figure 1. The NAT64 "A" is used when accessing IPv4-only servers in the datacenter, and the NAT64 "B" is used for Internet access.




Figure 1: NAT64s with IPv4 Address Ranges

The heuristic discovery method described herein does not support learning of the possible rules used by a DNS64 server for mapping specific IPv4 address ranges to separate Pref64::

and connectivity establishment procedures. The operator of the NAT64 and the DNS64 ought to take this into account in the network design.

The network operators can help IPv6-only nodes by ensuring the nodes do not have to work with IPv4 address literals for which special mapping rules are used. That is, the IPv4-only servers addressed from the special IPv4 address ranges ought to have signed AAAA records, which allows IPv6-only nodes to avoid local address synthesis. If the IPv6-only nodes are not using DNSSEC, then it is enough if the network's DNS64 server returns synthetic AAAA resource records pointing to IPv4-only servers. Avoiding the need for IPv6-only nodes to perform address synthesis for IPv4 addresses belonging to special ranges is the best approach to assist nodes.

If the IPv6-only nodes have no other choice than using IPv4-address literals belonging to special IPv4 address ranges, and the IPv6-only node will perform local synthesis by using the discovered Pref64::

delivered to the correct NAT64. For example, a router in the path from an IPv6-only host to NAT64s can forward the IPv6 packets to the correct NAT64 as illustrated in Figure 2. The routing could be based on the last 32-bits of the IPv6 address, but the network operator can also use some other IPv6 address format allowed by [RFC 6052](#) [[RFC6052](#)], if it simplifies routing setup. This setup requires additional logic on the NAT64 providing connectivity to special IPv4 address ranges: it needs to be able to translate packets it receives that are using the Pref64::




Figure 2: NAT64s with Assisting Router

6. Exit Strategy

A day will come when this tool is no longer needed. A node SHOULD implement a configuration knob for disabling the Pref64::

7. Security Considerations

The security considerations follow closely those of [RFC 6147](#) [[RFC6147](#)]. The possible attacks are very similar in the case where an attacker controls a DNS64 server and returns tampered IPv6 addresses to a node and in the case where an attacker causes the node

to use tampered Pref64::

attacker).

Even though well-known name's DNS A resource record would not necessarily need to be protected with DNSSEC, as both the name and IPv4 addresses well-known, for the DNS AAAA resource record queries DNSSEC protection is required. Without DNSSEC, fake positive AAAA responses could cause hosts to erroneously detect Pref64::[RFC6147] Section 3 case 5, for example) to detect malicious AAAA resource records. Therefore, the zone serving the well-known name has to be protected with DNSSEC.

For Pref64::Section 3.1 to validate each discovered Pref64::

The procedure of [Section 3.1.2](#) requires a node using DNSSEC to validate discovery of Pref64::Section 3.1.2 an implementation might be tempted to ask for a user to add a received domain as trusted - temporarily or permanently. The history has shown that average users are unable to properly handle such queries, and tend to answer positively without thinking in an attempt to get quickly forward. Therefore, unless the DNSSEC-using implementation has a way to dynamically and reliably add trusted domains, it is better to fail the Pref64::

Lastly, the best mitigation action against Pref64::

[8.](#) IANA Considerations

[8.1.](#) Domain Name Reservation Considerations

According to procedures described in [\[RFC3172\]](#) and [\[I-D.cheshire-dnsext-special-names\]](#) this document requests IANA to delegate a new second level domain in the .ARPA zone for the well-known domain name "ipv4only.arpa". The intention is that there will

not be any further delegation of names below the ipv4only.arpa domain. The administrative and operational management of this zone is to be undertaken by IANA. The answers to seven questions of [[I-D.cheshire-dnsext-special-names](#)] are as follows:

1. No, although this is a domain delegated under the .arpa infrastructural identifier top level domain."
2. Yes. Any application attempting to perform NAT64 discovery will query the name.
3. Yes, to the extent the API or library is affected by NAT64.
4. No.
5. No.
6. This name has effects for operators of NAT64/DNS64, but otherwise is just another delegated .arpa domain.
7. The registry for .arpa is held at IANA, and only IANA needs to take action here.

[8.2.](#) IPv4 Address Allocation Considerations

The well-known name needs to map to two different global IPv4 addresses, which are to be allocated as described in [[RFC5736](#)] and [[I-D.bonica-special-purpose](#)]. The addresses are to be taken from the IANA IPv4 Special Purpose Address Registry [[RFC5736](#)], and 192.0.0.170 and 192.0.0.171 are to be assigned to this document with parameters shown below:

Attribute	Value
Address Block	192.0.0.170/32 192.0.0.171/32
Name	NAT64/DNS64 Discovery
RFC	RFC TBD Section 2.2.
Allocation Date	February 2013
Termination Date	N/A
Source	False
Destination	False
Forwardable	False

Global	False	
Reserved-by-protocol	True	
+-----+-----+-----+		

The Record for IPv4 address allocation for IPv4 Special Purpose Address Registry

The following two records should be added to .arpa to the name servers run by ICANN/IANA:

```
ipv4only IN A 192.0.0.170
```

```
ipv4only IN A 192.0.0.171
```

8.3. IAB Statement Regarding This .arpa Request

With the publication of this document, the IAB approves of the delegation of "ipv4only" in the .arpa domain. Under [[RFC3172](#)], the IAB is requesting IANA to delegate and provision ipv4only.arpa as written in this specification. However, the IAB does not take any architectural or technical position about this specification.

9. Acknowledgements

Authors would like to thank Dmitry Anipko, Cameron Byrne, Aaron Yi Ding Christian Huitema, Washam Fan, Peter Koch, Stephan Lagerholm, Zhenqiang Li, Simon Perreault, Marc Petit-Huguenin, Andrew Sullivan, and Dave Thaler, for significant improvement ideas and comments.

10. References

10.1. 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.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", [RFC 4035](#), March 2005.

Internet-Draft

Pref64::/n Discovery

April 2013

- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), March 2006.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", [RFC 6052](#), October 2010.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [RFC 6146](#), April 2011.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", [RFC 6147](#), April 2011.
- [RFC6672] Rose, S. and W. Wijngaards, "DNAME Redirection in the DNS", [RFC 6672](#), June 2012.

10.2. Informative References

- [I-D.bonica-special-purpose]
Cotton, M., Vegoda, L., Bonica, R., and B. Haberman, "Special-Purpose IP Address Registries", [draft-bonica-special-purpose-07](#) (work in progress), January 2013.
- [I-D.cheshire-dnsext-special-names]
Cheshire, S. and M. Krochmal, "Special-Use Domain Names", [draft-cheshire-dnsext-special-names-03](#) (work in progress), September 2012.
- [RFC3172] Huston, G., "Management Guidelines & Operational Requirements for the Address and Routing Parameter Area Domain ("arpa")", [BCP 52](#), [RFC 3172](#), September 2001.
- [RFC5735] Cotton, M. and L. Vegoda, "Special Use IPv4 Addresses", [BCP 153](#), [RFC 5735](#), January 2010.
- [RFC5736] Huston, G., Cotton, M., and L. Vegoda, "IANA IPv4 Special


```
2002050501 ; serial
100         ; refresh (1 minute 40 seconds)
200         ; retry (3 minutes 20 seconds)
604800      ; expire (1 week)
100         ; minimum (1 minute 40 seconds)
)
```

```
example.com. IN NS ns.example.com.
```

```
nat64.example.com.
```

```
IN AAAA 2001:db8:0:0:0:0:C000:00AA
IN AAAA 2001:db8:0:0:0:0:C000:00AB
IN A 192.0.2.1
```

To DNSSEC sign the records, the owner of the example.com zone would have RRSIG records for both the AAAA and A records for nat64.example.com. As a normal DNSSEC requirement, the zone and its parent also need to be signed.

[Appendix B](#). About the IPv4 Address for the Well-Known Name

The IPv4 addresses for the well-known name cannot be non-global IPv4 addresses as listed in the [Section 3 of \[RFC5735\]](#). Otherwise DNS64 servers might not perform AAAA record synthesis when the well-known prefix is used, as stated in [Section 3.1 of \[RFC6052\]](#). However, the addresses do not have to be routable or allocated to any real node, as no communications will be initiated to these IPv4 address.

Allocation of at least two IPv4 addresses improves the heuristics in cases where the bit pattern of the primary IPv4 address appears more than once in the synthetic IPv6 address (i.e., the NSP prefix contains the same bit pattern as the IPv4 address).

If no well-known IPv4 addresses would be statically allocated for this method, the heuristic would require sending of an additional A query to learn the IPv4 addresses that would be then searched from inside of the received IPv6 address.

Teemu Savolainen
Nokia
Hermiankatu 12 D
FI-33720 Tampere
Finland

Email: teemu.savolainen@nokia.com

Jouni Korhonen
Nokia Siemens Networks
Linnoitustie 6
FI-02600 Espoo
Finland

Email: jouni.nospam@gmail.com

Savolainen, et al.

Expires October 06, 2013

[Page 19]

Internet-Draft

Pref64::/n Discovery

April 2013

Dan Wing
Cisco Systems
170 West Tasman Drive
San Jose, California 95134
USA

Email: dwing@cisco.com

