

v6ops  
Internet-Draft  
Intended status: Informational  
Expires: 29 November 2026

J. Palet Martinez  
The IPv6 Company  
A. D'Egidio  
Cabase - C  mara Argentina de Internet  
28 May 2026

464XLAT/MAT-T Optimization  
draft-ietf-v6ops-464xlat-optimization-05

## Abstract

IP/ICMP Translation Algorithm (SIIT) can be used to provide access for IPv4-only hosts or applications to IPv4-only or dual-stack destinations over IPv6-only infrastructure. In that case, the traffic flows are translated twice: first from IPv4 to IPv6 (stateless NAT46 at the ingress point to the IPv6-only infrastructure) and then from IPv6 back to IPv4 (stateful NAT64, at the egress point). When the destination is IPv6-enabled, the second translation might be avoided. This document describes a possible optimization to 464XLAT and MAP-T to avoid translating IPv6 flows back to IPv4 if the destination is reachable over IPv6. The proposed solution would significantly reduce the NAT64 utilization in the operator's network, increasing the performance.

## 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 <https://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 29 November 2026.

## Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

## Table of Contents

1. Introduction . . . . .	3
2. Requirements Language . . . . .	3
3. Possible Optimization . . . . .	4
4. Problem Statement Summary . . . . .	6
5. Solution Approaches . . . . .	8
5.1. Approach 1: DNS/Routing-based Solution . . . . .	8
5.2. Approach 2: NAT46/CLAT/DNS-proxy-EAM-based Solution . . . . .	9
5.2.1. Optimization enabling . . . . .	10
5.2.2. Detection of IPv4-only hosts . . . . .	10
5.2.3. Detection of IPv6-enabled service . . . . .	11
5.2.4. CE DNS proxy responses . . . . .	11
5.2.5. Creation of EAMT entries . . . . .	11
5.2.6. Forwarding path via stateful NAT for existing EAMT entries . . . . .	13
5.2.7. Maintenance of the EAMT entries . . . . .	14
5.2.8. Usage example . . . . .	14
5.2.9. Behavior in case of multiple A/AAAA RRs . . . . .	15
5.2.10. Behavior in presence/absence of DNS64 . . . . .	15
5.2.11. Behavior when using literal addresses or non IPv6-compliant APIs . . . . .	15
5.2.12. Behavior in case of Foreign DNS . . . . .	15
5.2.13. False detection of a dual-stack host as IPv4-only . . . . .	16
5.2.14. Behavior in presence of Happy Eyeballs . . . . .	17
5.2.15. Troubleshooting Implications . . . . .	18
5.3. Approach 3: NAT46/CLAT-provider-EAM-based Solution . . . . .	19
6. IPv6-only Services become accessible to IPv4-only devices/ apps . . . . .	20
7. Conclusions . . . . .	20
8. Security Considerations . . . . .	21
9. IANA Considerations . . . . .	21
10. Acknowledgements . . . . .	21
11. References . . . . .	22
11.1. Normative References . . . . .	22
11.2. Informative References . . . . .	23
Authors' Addresses . . . . .	25

## 1. Introduction

Different transition mechanisms, typically in the group of the so-called IPv6-only with IPv4aaS (IPv4-as-a-Service), such as 464XLAT ([RFC6877]) or MAP-T ([RFC7599]), allow IPv4-only hosts or applications to connect with IPv4 services in Internet over IPv6-only infrastructure, by means of a stateless NAT46 SIIT (IP/ICMP Translation Algorithm) as described by [RFC7915].

This is done by the implementation of SIIT at the CE (Customer Edge) Router or sometimes at the end-device, for example, the UE (User Equipment) in cellular networks. This functionality is the CLAT (Customer Translator) in the case of 464XLAT, while in the case of MAP-T is called NAT46.

The NAT46/CLAT (WAN side) is connected by IPv6-only to the operator network, which in turn, will have a reverse translation, the stateful NAT64 ([RFC6146]), known as PLAT (Provider Translator) in the case of 464XLAT. This allows to translate the IPv6 flow back to IPv4, in order to forward it to Internet.

In both cases (NAT46 and NAT64), the translation of the packet headers is done using the IP/ICMP translation algorithm defined in [RFC7915]. Translation between IPv4 and IPv6 addresses is done as per [RFC6052]. The NAT64 prefix should be discovered by the CE by one or more of the existing mechanisms ([RFC7225], [RFC8781] or [RFC7050]), and sometimes it is pre-configured at the CE to the WKP.

In addition to that, the concept of IPv6-Mostly [I-D.ietf-v6ops-6mops] has been introduced, by taking advantage of a new DHCP option [RFC8925] so that networks may cooperate with host that have a CLAT and don't need IPv4, can use IPv6-Only. In some cases this could mean that some enterprise networks may also incorporate internally stateful NAT64 functions.

## 2. 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 BCP14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

### 3. Possible Optimization

In the case of 464XLAT, a DNS64 ([RFC6147]) is (optionally) in charge of the synthesis of AAAA records from the A records, so the NAT64 can be used without the need of doing a double-translation by means of the NAT46/CLAT.

However, the DNS64 is not useful for the IPv4-only hosts or applications in the LANs, as they will not be able to use the AAAA records, so they are always forced to use the double-translation.

This is the expected behavior, as the original design of NAT64 was targeted to connect IPv6-only devices (using DNS) to IPv4-only services. 464XLAT expanded the solution to also allow IPv4-only devices (even if not using DNS) to connect to IPv4-only services by means of IPv6-only access networks.

The optimization solutions presented by this document try to avoid this double-translation, in the cases when the services, are already IPv6-enabled. So, in those cases, if the NAT46 already translated the IPv4 flow to IPv6, it doesn't look necessary to translate this back to IPv4.

A typical 464XLAT deployment is depicted in Figure 1.

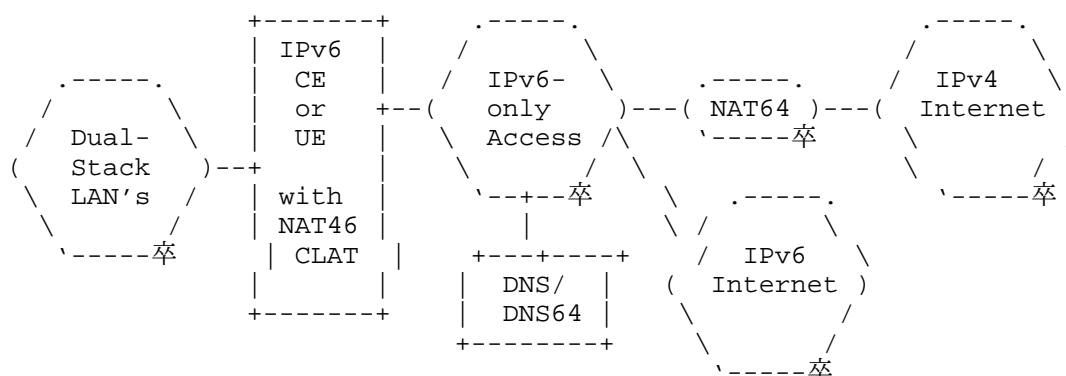


Figure 1: Typical 464XLAT Deployment

Examples of a topology shown on the above picture includes:

- \* An IPv6-only residential access network where the CE Router (with NAT46/CLAT) supports Dual-Stack in the customer LANs.

- \* An IPv6-only cellular network where a UE uses the NAT46/CLAT for dual-stack internal applications and other hosts connected via tethering.

If the operator is providing direct access, for example, to Content Delivery Networks (CDNs), caches, or other resources, and they are dual-stacked, the situation can be described as shown in Figure 2.

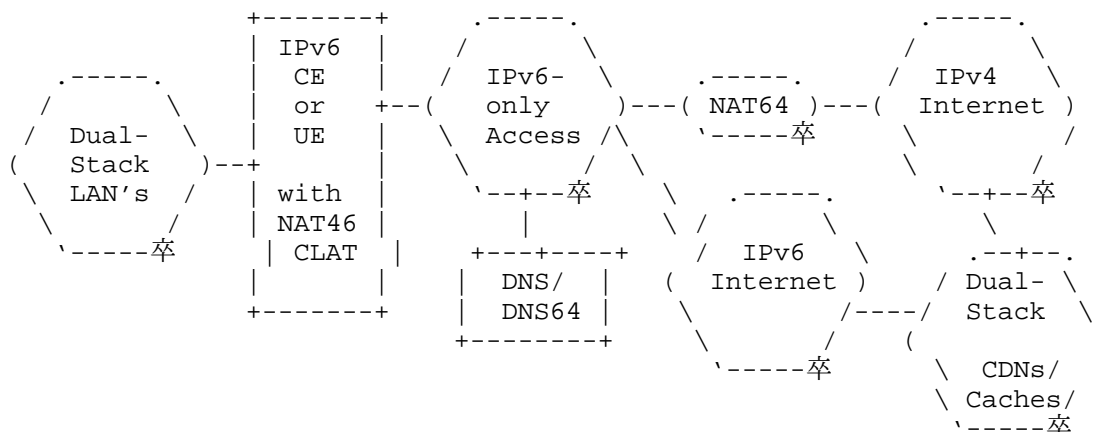


Figure 2: Typical 464XLAT Deployment with CDNs/Caches

In this case if the flows initiated in the LANs come from IPv4-only hosts or applications, even if the destination resources are IPv6-enabled, the double-translation is enforced, which has the following consequences:

- \* More traffic needs to pass thru the NAT64 devices.
- \* More NAT64 devices may be needed to handle the additional traffic.
- \* Additional usage of IPv4 addresses.
- \* Additional state at the NAT64 devices.
- \* Additional logging, its relevant storage and processing resources.
- \* Increasing of delay and reduction of traffic performance.
- \* Unnecessary point(s) of failure for that traffic.

Clearly, all those aspects have impact in both, CapEx and OpEx. This is extremely important when considering that most of the time, the contents stored in CDNs, caches, and so on, is there for a good

reason: It is frequently accessed resources and/or big. Examples such as video, audio, software and updates, are very common. So, this optimization can be highly impacting in many networks.

#### 4. Problem Statement Summary

If the hosts or applications in the customer LAN are IPv6-capable, then the access to the CDNs, caches or other resources, will be made in an optimized way, by means of IPv6-only, not using the NAT64, as depicted in Figure 3.

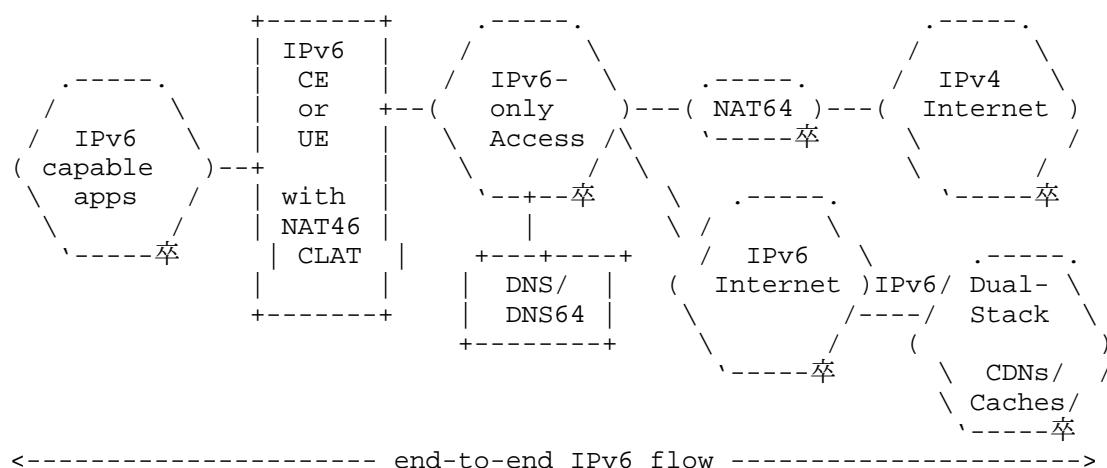


Figure 3: 464XLAT access to CDNs/Caches by IPv6-capable apps

However, if the hosts or applications are IPv4-only, for example, many SmartTVs and Set-Top-Boxes available today, a non-optimal double translation will occur (NAT46 at the CLAT and NAT64 at the PLAT), as illustrated in Figure 4.

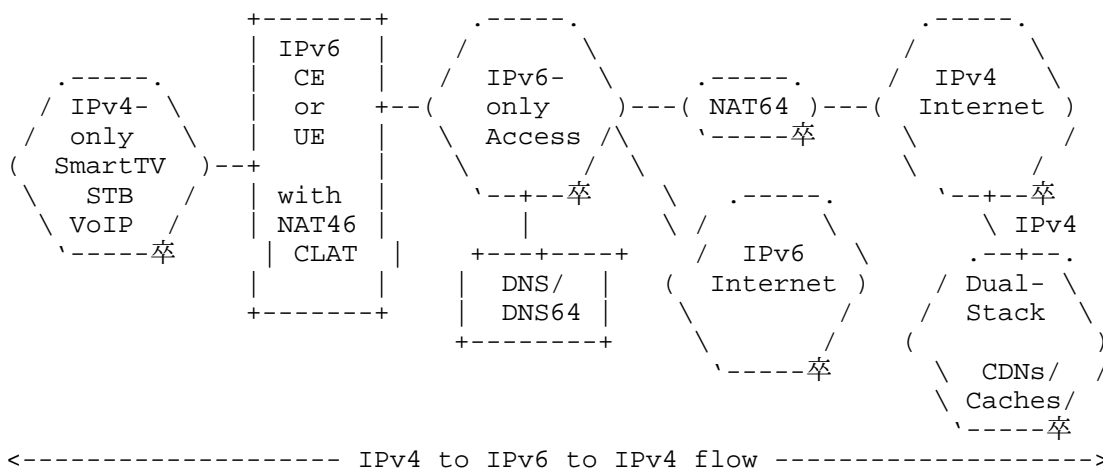


Figure 4: 464XLAT access to CDNs/Caches by IPv4-only apps

Clearly, this is a non-optimal situation, as it means that even if there is a dual-stack service, the NAT46/CLAT translated IPv4 to IPv6 traffic flow, is unnecessarily translated back to IPv4, traversing the stateful NAT64. This has a direct impact in the need to scale the NAT64 beyond what will be actually needed, if possible solutions, in order to keep using the IPv6 path towards those services, are considered.

As shown in the Figure 4, this is also the case for many other services, not just CDNs or caches, such as VoIP access to the relevant operator infrastructure, which may be also dual-stack. This is true as well for many other dual-stack or IPv6-enabled services, which may be directly reachable from the operator infrastructure, even if they are not part of it, for example peering agreements, services in IXs, etc. In general, this will become a more frequent situation for many other services, which are not yet dual-stack.

Moreover, it is already well-known that many CDN/caches, already offer in case of IPv4 via CGN, routing mechanisms so the private IPv4 addresses behind the CGN can access directly the CDN/caches without requiring the CGN translation. That means that IPv6-only with IPv4aaS solutions are in a clear disadvantage without the optimization as suggested in this document.

For simplicity, across the rest of this document, references to CDNs/caches, should be understood, unless otherwise stated, as any dual-stacked resources.

This document looks into different possible solution approaches in order to optimize the IPv4-only SIIT translation providing a direct path to IPv6-capable services, as depicted in Figure 5.

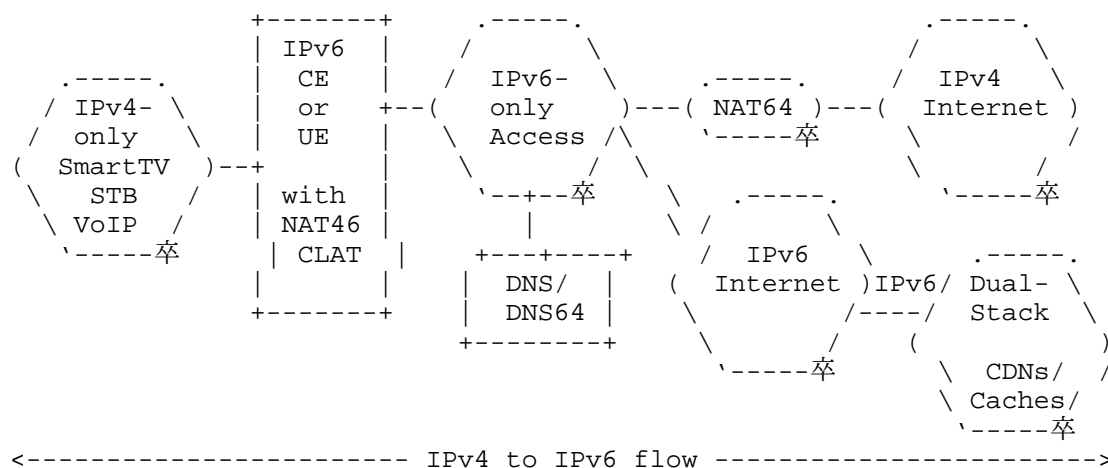


Figure 5: Optimized 464XLAT access to CDNs/Caches by IPv4-only apps

## 5. Solution Approaches

### 5.1. Approach 1: DNS/Routing-based Solution

Because the IPv4-only devices will not be able to query for AAAA records, the NAT46/CLAT/CE will translate the IPv4 addresses from the A record for the CDN/cache destination, using the WKP or NSP, as configured by the operator.

If the CDN/cache provider is able to configure, in the relevant interfaces of the CDN/caches, the same IPv6 addresses that will naturally result as the translated destination addresses for the queried A records, preceded by the WKP or NSP, then having more specific routing prefixes, will result in traffic to those destinations being directly forwarded towards those interfaces, instead of needing to traverse the NAT64.

For example, let's suppose a provider using the WKP (64:ff9b::/96) and a SmartTV querying for www.example.com:

www.example.com	A	192.0.2.1
NAT46/CLAT translated to		64:ff9b::192.0.2.1
CDN IPv6 interface must be		64:ff9b::192.0.2.1
Operator must have a specific route to		64:ff9b::192.0.2.1

Note: Examples using text representation as per Section 2.3 of [RFC6052] and IPv4 documentation addresses following [RFC5737].

It should be remarked that this approach requires that the path to the destination is configured in such way (i.e., more specific routing prefixes), that doesn't traverse the NAT64 devices.

Because the WKP is non-routable, this solution will only be possible if the CDN/cache is in the same ASN as the provider network, or somehow interconnected without routing thru Internet.

This solution has the additional drawback of the operational complexity/issues added to the operation of the CDN/cache, and the need to synchronize any IPv4 interface address changes with the relevant IPv6 ones, and possibly with routing.

However, is well know than several CDN/cache providers already do this when using CGN (NAT444), using non-global routable addresses such as those defined by [RFC1918] and [RFC5735]. Consequently, it should be possible to use the same approach to avoid the stateful NAT64 translation, by means of a proper DNS/Routing configuration.

The approach of using non-global routable addresses ([RFC1918] and [RFC5735]) for the synthesys of AAAA RRs is not possible with current specifications of [RFC6052], when using the WKP, as it doesn't allow to use non-global IPv4 addresses. However, this has been discussed as inconvenient, is not frecuently restricted in stateful NAT64 implementations and is being reversed by NAT64 WKP [I-D.ietf-v6ops-nat64-wkp-1918].

## 5.2. Approach 2: NAT46/CLAT/DNS-proxy-EAM-based Solution

If the NAT46/CLAT/CE, as commonly is the case, is also a DNS proxy/stub resolver, it is possible to modify the behavior and create an "internal" interaction among both of them.

This approach uses the existing IPv4 and IPv6 addresses in the A and AAAA records, respectively, so no additional complexity/issues added to the CDN/caches operations.

The following sub-sections detail this approach and provide a step-by-step example case.

### 5.2.1. Optimization enabling

This optimization **MUST** be enabled by default, but only when the WAN link is IPv6-only and the NAT46/CLAT is being used. Note that the NAT64/CLAT can only be used if the NAT64 prefix has been discovered. This allows the users to get CEs from the retail and take advantage of the optimization, without requiring any configuration.

The CE **MUST** support a way to completely disable the optimization, in order to allow the operator to turn it off in case it is required.

It is expected as well, that the NAT64/CLAT is only enabled if the WAN link is IPv6-only.

### 5.2.2. Detection of IPv4-only hosts

The goal is to ensure that only IPv4-only hosts are optimized. Towards that, the CE **MUST** use ARP and ND (and their relevant caches, if the information of a hosts has been already learnt) each time a new host starts a connection. If it is possible to bind the same MAC address to both an IPv4 and IPv6 address, then the host is not IPv4-only (it may be IPv6-only or dual-stack), and **MUST NOT** be optimized.

The CE **MUST** maintain a table of IPv4-only hosts and ensure that if any IPv4-only hosts become IPv6-enabled, it is properly handled.

This mechanism to detect IPv4-only hosts has two drawbacks:

1. If a subscriber has intermediate dual-stack routers in between the IPv4-only host and the NAT46/CLAT, the IPv4-only hosts will be detected as dual-stack, so no optimization will be performed. This is not the most common scenario, as typically the devices that are more relevant to the optimization (in terms of those that generate more IPv4-only traffic) are directly attached to the CE, or in bridged interfaces.
2. If a host is dual-stack, but has some IPv4-only applications, because the host will be detected as dual-stack, none of the applications will be optimized. This is a good trade-off, considering the most important traffic to optimize is typically coming from real IPv4-only devices such as old SmartTVs/STBs. Furthermore, this avoid breaking other mechanisms present only in dual-stack hosts, such as Happy Eyeballs [RFC8305] and simplifies troubleshooting.

It needs to be remarked that, if the detection of the IPv4-only host is done incorrectly (either not detecting it or by a false detection), the goal is that no harm is caused. In the worst case, optimization MUST NOT be performed.

#### 5.2.3. Detection of IPv6-enabled service

In the case of an IPv4-only detected host, the DNS proxy/stub resolver MUST actually perform an additional AAAA query, unless the information is already present in the Additional Section, as per Section 3 of [RFC3596].

Note that the NAT46/CLAT MUST already know the WKP or NSP being used in that network. If the response contains at least one IPv6 address not using the WKP/NSP, it means that the destination is IPv6-enabled (because at least one of the IPv6 addresses is not synthesized). This means that it is possible for the NAT46/CLAT, to create an Explicit Address Mapping ([RFC7757]).

#### 5.2.4. CE DNS proxy responses

In the case of an IPv4-only detected host, the CE DNS proxy MUST only return the answer to an A query once any of the following happens:

1. An answer to the AAAA query has been received.
2. A SERVFAIL has been received.
3. The "Resolution Delay" has passed.

The Resolution Delay MUST be set to the same value (50 milliseconds) as indicated by Happy Eyeballs [RFC8305].

#### 5.2.5. Creation of EAMT entries

An EAM Table (EAMT used for short, across the rest of this document) MUST be created/maintained automatically by the NAT46/CLAT, which is responsible to prioritize any available entries in the EAMT, versus the use of any synthetic AAAA.

The EAMT entry MUST only be created if all the following conditions are met:

1. There is not already EAMT entry for the same pair of MAC/destination IPv4 address/prefix.
2. The source host is IPv4-only.

3. The DNS proxy is ready to return the A answer (according to Section 5.2.4).
4. There is at least one non-synthesized AAAA response.
5. If DNSSEC is available, the response has been locally validated or the AD bit has been set by a trusted resolver, as per [RFC4035].

This avoids a slight NAT64 overload and flapping between destination addresses (IPv4/IPv6), which may impact some applications, at the cost of a small extra delay for the initial communication setup, when the EAMT entry doesn't yet exist.

Each EAMT entry MUST contain, the fields already described in [RFC7757] and a few new extensions (as per section 3.1 of [RFC7757]):

1. ID: EAMT Entry Index (optional).
2. MAC address: Identify the host to which this EAMT entry belongs.
3. Destination IPv4 address/prefix: By default, the prefix length is 32 bits.
4. Destination IPv6 address/prefix: By default, the prefix length is 128 bits. Only non-synthesized addresses are allowed.
5. TTL: It MUST be set to the minimum value from the AAAA/A RR pair. In normal conditions the TTL for both A and AAAA records, of a given FQDN, should be the same, so this ensures a proper behavior if there is any DNS mismatch.
6. FQDN: The one that originated the A query for this EAMT entry. Required in order to ensure a correct detection of cases such as the use of reverse-proxy with a single IPv4 address to multiple IPv6 addresses.
7. Valid/Stale/Invalid: When set to Stale, means that this EAMT entry MUST NOT be used for new connections. When set to Invalid, means that this EAMT entry can be deleted, unless the Auto/Static bit is also set.
8. Auto/Static: When set to 1, means that this EAMT entry has been manually/statically configured, for example by means of an explicit configuration (GUI, CLI, provisioning system, etc.).

Note that allowing destination IPv4 and IPv6 prefixes, together with the Valid/Stale/Invalid and Auto/Static flags, allows manual explicit optimization and non-optimization configuration for specific parts of Internet.

When a new EAMT entry is first automatically created, it is flagged as "Valid" and "Auto". If a subsequent A query, with a different FQDN, results in an IPv4 address that has already an EAMT entry and a different IPv6 address, it means that some reverse-proxy or similar functionality is being used by the IPv6-enabled service. In this case, the existing EAMT entry will be marked as "Stale". No new EAMT entry is created for that IPv4 address. Otherwise, the optimization will only allow to access the first set of IPv4/IPv6/FQDN, which may break the access to other FQDN that share the same IPv4 address and different IPv6 addresses.

In this case the EAMT entry will still become "Invalid" according the TTL, which allows to re-enable optimization if a new query for the A record has changed the situation. For example, maybe the reverse-proxy has been removed, or there is now only a single app using it, so at the time being, the optimization is again possible without creating troubles to other hosts.

An EAMT entry marked as "Stale" or "Invalid", only affects the relevant host. Other hosts may have their own EAMT entries or they are using the regular NAT46/CLAT+NAT64 path (without the optimization).

#### 5.2.6. Forwarding path via stateful NAT for existing EAMT entries

Following this approach, if there is a valid EAMT entry, for a given pair of source-MAC-address/IPv4-destination, the IPv6-native path pointed by the IPv6 address of that EAMT entry, will take precedence versus the NAT64 path, so the traffic will not be forwarded to the NAT64.

However, this is not sufficient to ensure that individual applications are able to keep existing connections. In many cases, audio and video streaming may use a single TCP connection lasting from minutes to hours. Instead, the CDN TTLs may be configured in the range from 10 to 300 seconds in order to allow new resolutions to switch quickly and to handle large recursive resolvers (with hundreds of thousands of clients behind them).

Consequently, the EAMT entries MUST NOT be used directly to establish a forwarding path, but instead, MUST be used to create a stateful NAT entry for the 4-tuple for the duration of the session/connection. This means that to implement the optimization the NAT46 MUST be

stateful. Typically, stateful NAT46 are implemented by means of a stateful NAT44 (which often maybe hardware off-loaded), followed by a stateless NAT46. If the SoC/code is able to do stateless NAT46, this still could be used when the optimization is disabled.

#### 5.2.7. Maintenance of the EAMT entries

An EAMT entry with the Auto/Static bit set, MUST NOT be automatically modified.

An EAMT entry with the Auto/Static bit clear, MUST be set to Stale in case of:

1. TTL time-out.
2. Or the conditions for creation of the EAMT entry (Section 5.2.5) aren't longer met.

Entries in Stale state MUST be set to Invalid once existing connections time-out.

The Invalid EAMT entries MUST be deleted, unless the Auto/Static bit is set. This allows users/operators to set explicit rules for diagnostics or resolution of issues in special situations.

#### 5.2.8. Usage example

Using the same example as in the previous approach:

www.example.com	A	192.0.2.1
	AAAA	2001:db8::a:b:c:d
EAMT entry	192.0.2.1	2001:db8::a:b:c:d
NAT46/CLAT translated to		2001:db8::a:b:c:d
CDN IPv6 interface already is		2001:db8::a:b:c:d
Operator already has a specific route to		2001:db8::a:b:c:d

The following is an example of the CE behavior after the previous case has already created an EAMT entry and a reverse-proxy is detected:

1. A query for www.example.net A RR is received
2. www.example.net A 192.0.2.1
3. www.example.net AAAA 2001:db8::e:e:f:f
4. A conflict has been detected

5. The existing EAMT entry for 192.0.2.1 is set to Stale

#### 5.2.9. Behavior in case of multiple A/AAAA RRs

If multiple A/AAAA RRs are available, any of them could be chosen and the optimization will not present any different result to the hosts compared with a situation where the optimization is not used.

Existing DNS proxy/stub resolvers already implement mechanisms for DNS Load Balancing ([RFC1794]). This don't need to be modified to implement the optimization if some degree of randomness is already secured.

To secure sufficient randomness, a possible algorithm shall ensure that different EAMT entries (for different hosts) are permuted randomly among different A/AAAA records on the A/AAAA RR set.

#### 5.2.10. Behavior in presence/absence of DNS64

464XLAT can be deployed/used with and without a DNS64. However, as indicated in Section 5.2.3, the EAMT entry is only created when the service is IPv6-enabled, because the optimization is only relevant for destinations which already have AAAA records. In those cases DNS64 is not relevant.

#### 5.2.11. Behavior when using literal addresses or non IPv6-compliant APIs

Because the EAMT entries are only created when the NAT46/CLAT/CE proxy/stub DNS is being used, any hosts or applications that don't use DNS, will not create the relevant entries.

They will not be optimized unless EAMT entries are statically configured.

#### 5.2.12. Behavior in case of Foreign DNS

Hosts or applications may use DNS servers from other networks. For a complete description of reasons for that, refer to Section 4.4 of [RFC8683]. In the case the DNS is modified, or some hosts or applications use other DNS servers, the possible scenarios and the implications are:

- a. Devices configured to use a DNS proxy/resolver which is not the CE/NAT46/CLAT (as it is the case for some consumer electronics). In this case this optimization will not work, because the EAMT entry will not be created based on their own flows. Nevertheless, the EAMT entry may be manually created. However,

the lack of EAMT entry, will not impact negatively in the user's hosts/applications (the optimization is not performed). It should be noticed that users commonly, don't change the configuration of devices such as SmartTVs or STBs (if they do, some other functionalities, such as CDN/caches optimizations may not work as well), so this only happens typically if the vendor is doing it on-purpose and for good well-known reasons.

- b. DNS privacy/encryption. Hosts or applications that use mechanisms for DNS privacy/encryption, such as DoT ([RFC7858], [RFC8094]), DoH ([RFC8484]) or DoQ ([RFC9250]), will not make use of the stub/proxy resolver, so the same considerations as for the previous case applies.
- c. Users that modify the DNS in their Operating Systems. This is quite frequent, however commonly Operating Systems are dual-stack, so aren't part of the problem statement described by this document and will not be adversely affected.
- d. Users that modify the DNS in the CE. This is less common. In this case, this optimization is not adversely affected, because it doesn't depend on the operator DNS, it works only based on the internal CE interaction between the NAT46/CLAT and the stub/proxy resolver. Note that it may be affected if the operator offers different "DNS views" or "split DNS", however this is not related to this optimization and will anyway impact in the other possible operator optimizations (i.e. CDN/cache features).
- e. Combinations of the above ones. No further impact is observed, beyond the ones already described.

#### 5.2.13. False detection of a dual-stack host as IPv4-only

If a dual-stack host is being detected as IPv4-only, is because it is not responding to the CE ND messages, so by all means, should be considered, at the time being, as IPv4-only, and consequently EAMT entries will be created and traffic will be optimized for IPv4 flows.

However, if this hosts suddenly become IPv6-enabled (or dual-stack), the relevant EAMT entries must be flagged by the CE as "Stale". The host will be able to complete the connections and the entries will be marked as "Invalid" and deleted.

Anyway, those EAMT entries, while "Valid", may not be actually used by the dual-stack hosts, because those hosts or applications should prefer IPv6, so most probably was either a temporary failure or done on-purpose (user, troubleshooting). If the host is preferring IPv4 for connecting to the CDN/cache or IPv6-enabled service, it will be

actually using the NAT46/CLAT, including the EAMT entry and consequently IPv6, so this mechanism will be correcting an undesirable behavior. This may be also a special case, which actually seems to be an incoherent host or application implementation.

#### 5.2.14. Behavior in presence of Happy Eyeballs

Happy Eyeballs [RFC8305] is only enabled in dual-stack hosts. Consequently, it is not affected by this optimization because both, as the host should not be detected as IPv4-only, following Section 5.2.2.

If Happy Eyeballs triggers a fallback to IPv4 for a given host, it will be actually using IPv4 without the optimization, which in turn, uses the IPv6-only WAN link of the NAT46/CLAT/CE. So even if Happy Eyeballs is present, IPv4 is expected to be slower than native IPv6 itself due to delays added by the NAT46/CLAT+NAT64 translations. This optimization reduces those delays by eliminating the second translation (NAT64) for IPv4-only detected hosts.

However, there may be cases where this may be understood as problematic. The possible reasons why Happy Eyeballs may trigger an IPv4 fallback, in the case of IPv6-only access networks with IPv4aaS, in general, can be classified as:

1. Failure at the CE or customer LANs. It may happen that the CE or other devices in the customer LANs are showing erratic behaviors or malfunctions. It is difficult to believe that this happens only with IPv6, and if that's the case Happy Eyeballs will not resolve the issue, because IPv4 is provided as a service on top of IPv6.
2. Complete failure of the IPv6-only link or IPv6-only operator's infrastructure (up to the NAT64). In this case, IPv4 will not work for that subscriber. Happy Eyeballs will not resolve the issue, and instead will only be adding some extra delay (the attempt to fallback to IPv4 before timing-out).
3. Complete failure of both IPv4 and IPv6 links behind the operator's NAT64 towards the destination. In this case, typically both, IPv4 and IPv6 will fail (in many cases, they are dual-stack links, not different links). Again, Happy Eyeballs will also fail to resolve the issue.
4. Complete failure only in the IPv6 links behind the operator's NAT64 towards the destination. This is less frequent, but still miss-configured AAAA RRs, or diverse paths for IPv4 and IPv6

together with outages or IPv6-only routing issues, could generate this problem. In this case, Happy Eyeballs could resolve the issue, however. It will work because the optimization is not enabled for dual-stack hosts.

5. Partial failure: Slower IPv6 vs IPv4 path end-to-end. In general, the added delay of the IPv4 translations and NATs across the path, increases the chances that IPv4 is slower than IPv6. However, it may happen that there is some IPv6 specific link congestion or packet dropping, that generates the reverse situation, so IPv4 become faster than IPv6. Because the optimization is not being used for dual-stack hosts, Happy Eyeballs will be resolving the issue.

In summary, the optimization will not change the Happy Eyeballs behaviour. Furthermore, it should be observed that IPv6 failures will also impact other operators (even if not using the optimization), and especially those using only NAT64+DNS64 instead of 464XLAT, or even more, any IPv6-only hosts and applications in any operator network across the entire Internet. It looks like it is very important to make sure that, as IPv6 is more prevalent, there is a better monitoring and failures are detected ASAP, instead of being hidden by Happy Eyeballs, specially in IPv6-only networks. It should be noticed also that in IPv6-only with IPv4aaS, the chances of troubles in the IPv4 paths seem to be higher than in the IPv6, as there are more translations, more devices, more delays, while the optimization will precisely reduce them.

#### 5.2.15. Troubleshooting Implications

When there is a need to troubleshoot IPv4 from the CE LANs, it may happen that there is an EAMT entry forcing the flow to a given destination(s) to use IPv6, which will distort the results, unless the host being used is dual-stack (which is the most common situation).

This can be avoided, using a CLI/GUI or provisioning procedure, to either completely disable the optimization during the troubleshooting, or create specific static EAMT entries, using the Valid/Stale/Invalid and Auto/Static flags, as described in Section 5.2.5.

Consequently, the CE MUST allow both, disabling the optimization and the setup of manual/static EAMT entries.

### 5.3. Approach 3: NAT46/CLAT-provider-EAM-based Solution

Instead of using the DNS proxy/stub resolver to create the EAMT entries, the operator may push this table (or parts of it) into the CE/NAT46/CLAT, by using configuration/management mechanisms.

This solution has the advantage of not being affected by any DNS changes from the user (the EAMT is created by the operator) and ensures a complete control from the operator. However, it may impact the cases of devices with a DNS configured by the vendor.

In general, most of the considerations from the previous approach will apply.

One more advantage of this solution is that the EAMT pairs doesn't need to match the "real" IPv4/IPv6 addresses available in the A/AAAA records, as shown in the next example.

www.example.com	A	192.0.2.1
	AAAA	2001:db8::a:b:c:d
EAMT pulled/pushed entry	192.0.2.1	2001:db8::f:e:d:c
NAT46/CLAT translated to		2001:db8::f:e:d:c
CDN IPv6 interface already is		2001:db8::f:e:d:c
Operator already has a specific route to		2001:db8::f:e:d:c

EAMT may contain TTLs which probably are derived from DNS ones, or alternatively, a global TTL for the full table.

An alternative way to configure the table, is that the CE is actually pulling the table (or parts of it) from the operator infrastructure. In this case it will be mandatory that the entries have individual TTLs, again probably derived from the DNS ones.

This approach has three major drawbacks:

1. CDNs are used to do frequent changes at the DNS level, so unless the CDNs offer an API or equivalent convenient solution to keep updated the EAMT, the operator will need to cache the most frequent FQDNs being resolved in their own DNS and based on the TTLs, update the EAMT.
2. It requires a new protocol, or an extension to existing ones, in order to push or pull the EAMT updates.
3. It may generate additional bandwidth utilization in the WAN links for every CE when the EAMT needs to be update, even when a CE reboots.

## 6. IPv6-only Services become accessible to IPv4-only devices/apps

One of the issues with the IPv6 deployment, is that those services which become IPv6-only in Internet, aren't reachable by IPv4-only hosts and applications. This means that new content providers must support dual-stack even for new services, even while IPv4 public addresses aren't available.

If NAT46/CLAT/DNS-proxy-EAM approach (Section 5.2) is chosen, it also offers the chance to resolve this issue. This is possible if IPv6-only services get configured with an A resource record pointing to a well-known IPv4 address despite they aren't actually connected to IPv4. This is out of scope for this document, as it will require further work and a reservation by IANA, This will mean that those services will work fine if there is a NAT46/CLAT supporting the optimization. This A RR has no negative impact if the NAT46/CLAT doesn't exist, or it is not optimized, because is not reachable via IPv4-only, so is not a different situation compared with not having an A RR.

The result of this is equivalent to the approach taken by MAP-T (Section 12.3 of [RFC7599]). However, it has the advantage that the MAP-T approach is restricted to services in the MAP-T domain.

In fact, it may become an incentive for the IPv6 deployment in Internet services, as it provides the option to use a well-known IPv4 address (maybe anycast) for the "non-valid" A RR, that points, in case of port 80/443 to a web page or service that returns a warning such as "This service is only available if the network is properly connected to Internet with IPv6".

## 7. Conclusions

NAT46/CLAT/DNS-proxy-EAM approach (Section 5.2) seems the right solution for optimizing the access to dual-stack services, whether they are located inside or outside the ISP. It is also the only approach which has no additional requirements for the network operators (both ISPs and CDN/cache operators).

Having this type of optimization facilitates and increases the usage of IPv6, even for IPv4-only hosts and applications, at the same time that decreases the use of the NAT64.

SIIT already has a SHOULD for EAM support, so it is not a high additional burden the support required for existing implementations to be updated for this optimization.

However, this approach requires a modification of the NAT46/CLAT/CE code, which we know it can take a long time to be not only in the market, but also into firmware upgrades applied to existing CEs. Instead DNS/Routing-based Solution approach (Section 5.1), is already being used for CGN (NAT444) deployments, so it makes sense that those deployments, when the CDN/caches, already are dual-stacked, and have an internal routing path make use of the WKP with non-global routable IPv4 addresses.

Note that this approach can also be used in IPv6-Mostly deployments when the stateless NAT64 is routed internally in the enterprise network.

In summary, we shall recommend the DNS/Routing-based Solution approach (Section 5.1).

## 8. Security Considerations

This document does not have any new specific security considerations, besides the ones already known for DNS64.

Spoofed DNS responses could generate incorrect EAMT entries. However, this seems not different than if the optimization is not in place and the spoofed DNS responses are cached by the CE DNS proxy/stub resolver or even by hosts in the CE LANs. It very much depends on how and where the attack is actually performed.

In both cases, 464XLAT and MAP-T, the CE device should contain a DNS proxy/stub resolver, which is also required for the optimization. Nevertheless, it is common that the user change DNS settings. If it happens, in the case of MAP-T, the port-set is restricted for an efficient public IPv4 address sharing, so the entropy of the source ports is significantly lowered. In this case, theoretically MAP-T is less resilient against cache poisoning ([RFC5452]) compared with 464XLAT. However, an efficient cache poisoning attack requires that the subscriber operates its own caching DNS server and the attack is performed in the service provider network, so the chances of a successful exploitation of this vulnerability are low.

## 9. IANA Considerations

This document does not have any new specific IANA considerations.

## 10. Acknowledgements

The authors would like to acknowledge the inputs of Erik Nygren, Fred Baker, Martin Hunek, Chongfeng Xie, Fernando Gont, Fernando Frediani, Jen Linkova, Eduard Vasilenko and Philip Homburg.

## 11. References

### 11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", STD 88, RFC 3596, DOI 10.17487/RFC3596, October 2003, <<https://www.rfc-editor.org/info/rfc3596>>.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005, <<https://www.rfc-editor.org/info/rfc4035>>.
- [RFC5452] Hubert, A. and R. van Mook, "Measures for Making DNS More Resilient against Forged Answers", RFC 5452, DOI 10.17487/RFC5452, January 2009, <<https://www.rfc-editor.org/info/rfc5452>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<https://www.rfc-editor.org/info/rfc6052>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <<https://www.rfc-editor.org/info/rfc6146>>.
- [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, DOI 10.17487/RFC6147, April 2011, <<https://www.rfc-editor.org/info/rfc6147>>.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <<https://www.rfc-editor.org/info/rfc6877>>.

- [RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <<https://www.rfc-editor.org/info/rfc7050>>.
- [RFC7225] Boucadair, M., "Discovering NAT64 IPv6 Prefixes Using the Port Control Protocol (PCP)", RFC 7225, DOI 10.17487/RFC7225, May 2014, <<https://www.rfc-editor.org/info/rfc7225>>.
- [RFC7599] Li, X., Bao, C., Dec, W., Ed., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", RFC 7599, DOI 10.17487/RFC7599, July 2015, <<https://www.rfc-editor.org/info/rfc7599>>.
- [RFC7757] Anderson, T. and A. Leiva Popper, "Explicit Address Mappings for Stateless IP/ICMP Translation", STD 103, RFC 7757, DOI 10.17487/RFC7757, February 2016, <<https://www.rfc-editor.org/info/rfc7757>>.
- [RFC7915] Bao, C., Li, X., Baker, F., Anderson, T., and F. Gont, "IP/ICMP Translation Algorithm", RFC 7915, DOI 10.17487/RFC7915, June 2016, <<https://www.rfc-editor.org/info/rfc7915>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8305] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, December 2017, <<https://www.rfc-editor.org/info/rfc8305>>.
- [RFC8781] Colitti, L. and J. Linkova, "Discovering PREF64 in Router Advertisements", RFC 8781, DOI 10.17487/RFC8781, April 2020, <<https://www.rfc-editor.org/info/rfc8781>>.

## 11.2. Informative References

- [I-D.ietf-v6ops-6mops] Buraglio, N., Caletka, O., and J. Linkova, "IPv6-mostly Networks: Deployment and Operations Considerations", Work in Progress, Internet-Draft, draft-ietf-v6ops-6mops-07, 2 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-6mops-07>>.

- [I-D.ietf-v6ops-nat64-wkp-1918]  
Kumari, W. and J. Linkova, "NAT64 WKP", Work in Progress, Internet-Draft, draft-ietf-v6ops-nat64-wkp-1918-02, 16 May 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-nat64-wkp-1918-02>>.
- [RFC1794] Brisco, T., "DNS Support for Load Balancing", RFC 1794, DOI 10.17487/RFC1794, April 1995, <<https://www.rfc-editor.org/info/rfc1794>>.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC5735] Cotton, M. and L. Vegoda, "Special Use IPv4 Addresses", RFC 5735, DOI 10.17487/RFC5735, January 2010, <<https://www.rfc-editor.org/info/rfc5735>>.
- [RFC5737] Arkko, J., Cotton, M., and L. Vegoda, "IPv4 Address Blocks Reserved for Documentation", RFC 5737, DOI 10.17487/RFC5737, January 2010, <<https://www.rfc-editor.org/info/rfc5737>>.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC8094] Reddy, T., Wing, D., and P. Patil, "DNS over Datagram Transport Layer Security (DTLS)", RFC 8094, DOI 10.17487/RFC8094, February 2017, <<https://www.rfc-editor.org/info/rfc8094>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.
- [RFC8683] Palet Martinez, J., "Additional Deployment Guidelines for NAT64/464XLAT in Operator and Enterprise Networks", RFC 8683, DOI 10.17487/RFC8683, November 2019, <<https://www.rfc-editor.org/info/rfc8683>>.
- [RFC8925] Colitti, L., Linkova, J., Richardson, M., and T. Mrugalski, "IPv6-Only Preferred Option for DHCPv4", RFC 8925, DOI 10.17487/RFC8925, October 2020, <<https://www.rfc-editor.org/info/rfc8925>>.

[RFC9250] Huitema, C., Dickinson, S., and A. Mankin, "DNS over Dedicated QUIC Connections", RFC 9250, DOI 10.17487/RFC9250, May 2022, <<https://www.rfc-editor.org/info/rfc9250>>.

#### Authors' Addresses

Jordi Palet Martinez  
The IPv6 Company  
Molino de la Navata, 75  
28420 La Navata - Galapagar Madrid  
Spain  
Email: [jordi.palet@theipv6company.com](mailto:jordi.palet@theipv6company.com)  
URI: <https://www.theipv6company.com/>

Alejandro D'Egidio  
Cabase - Cテ。mara Argentina de Internet  
Suipacha, 128 - 3ツーF  
C1008AAD Ciudad Autテウ noma de Buenos Aires  
Argentina  
Email: [adegidio@cabase.org.ar](mailto:adegidio@cabase.org.ar)  
URI: <https://www.cabase.org.ar/>