

Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 30, 2020

M. Kuehlewind
Z. Sarker
Ericsson
January 27, 2020

Discovery Mechanism for QUIC-based, Non-transparent Proxy Services
draft-kuehlewind-quic-proxy-discovery-01

Abstract

Often an intermediate instance (such as a proxy server) is used to connect to a web server or a communicating peer if a direct end-to-end IP connectivity is not possible or the proxy can provide a support service like, e.g., address anonymisation. To use a non-transparent proxy a client explicitly connects to it and requests forwarding to the final target server. The client either knows the proxy address as preconfigured in the application or can dynamically learn about available proxy services. This document describes different discovery mechanisms for non-transparent proxies that are either located in the local network, e.g. home or enterprise network, in the access network, or somewhere else on the Internet usually close to the target server or even in the same network as the target server.

This document assumes that the non-transparent proxy server is connected via QUIC and discusses potential discovery mechanisms for such a QUIC-based, non-transparent proxy.

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 July 30, 2020.

Internet-Draft QUIC Non-transparent Proxy Discovery January 2020

Copyright Notice

Copyright (c) 2020 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 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.	Using DHCP for Local Discovery	3
3.	Using IPv6 Neighbor Discovery for Local Discovery	5
3.1.	Using PVDs	6
4.	DNS Service Discovery (DNS-SD)	7
4.1.	Local discovery using mDNS	7
4.2.	Discovery for Remote Domains	8
5.	Using PCP options	8
6.	Using Anycast address	9
7.	IANA Considerations	9
8.	Security Consideration	10
9.	Contributors	10
10.	Acknowledgments	10
11.	References	10
11.1.	Normative References	11
11.2.	Informative References	12
	Authors' Addresses	12

[1.](#) Introduction

QUIC is a new transport protocol that was initially developed as a way to optimize HTTP traffic by supporting multiplexing without head-of-line-blocking and integrating security directly into the transport. This tight integration of security allows the transport and security handshakes to be combined into a single round-trip exchange, after which both the transport connection and authenticated

encryption keys are ready.

Often an intermediate instance (such as a proxy server) is used to connect to a web server or a communicating peer if a direct end-to-end IP connectivity is not possible or the proxy can provide a

support service like, e.g., address anonymization. QUIC's ability to multiplex, encrypt data, and migrate between network paths makes it ideal for solutions that need to tunnel or proxy traffic.

Existing proxies that are based on TCP and HTTP are often transparent. That is, they do not require the cooperation of the ultimate connection endpoints, and are often not visible to one or both of the endpoints. If QUIC provides the basis for future tunneling and proxying solutions, it is expected that this relationship will change. At least one of the endpoints will be aware of the proxy, explicitly connect to it, and coordinate with it. This makes the proxy and tunneling non-transparent to at least most often the client. This allows client hosts to make explicit decisions about the services they request from proxies (for example, simple forwarding or more advance performance-optimizing services), and to do so using a secure communication channel between itself and the proxy. [[I-D.kuehlewind-quic-substrate](#)] describes some of the use cases for using QUIC for proxying and tunneling.

To use a non-transparent proxy service, a client explicitly connects to it and requests forwarding to the final target server. The client either knows the proxy address as preconfigured in the application or can dynamically learn about available proxy servers. This document describes different discovery mechanisms for proxies that are either located in the local network, e.g. home or enterprise network, in the access network, or somewhere else on the Internet usually close to the target server or even in the same network as the target server. For the rest of the document the work "proxy" refers to a non-transparent proxy.

The discovery mechanisms proposed in this document cover a range of approaches based on IETF protocols and commonly used mechanisms, however, other mechanisms in more specialized networks are possible as well. For 5G networks, the 3GPP specifies an extended exposure framework that potentially can also be used for proxy discovery and routing support.

After discovery a client can connect to the proxy and request a proxy service, e.g. using the MASQUE protocol [[I-D.schinazi-masque](#)], to instruct the proxy forward traffic to a target server as well as negotiate and request proxy capabilities and parameters.

2. Using DHCP for Local Discovery

DHCP [[RFC2131](#)] can be used to announce the IP address of local proxy server in IPv4 networks, as well DHCPv6 [[RFC8415](#)] in IPv6 networks. New options for both protocols are specified below and as shown in Figure 1 and Figure 2. In both cases the option can contain one or

more IP addresses (but of course IPV4 and IOv6 address respectively) of QUIC-based proxy servers (indicated by the Q flag). All of the addresses in one option share the same Lifetime value. If it is desirable to have different Lifetime values, multiple options can be used.

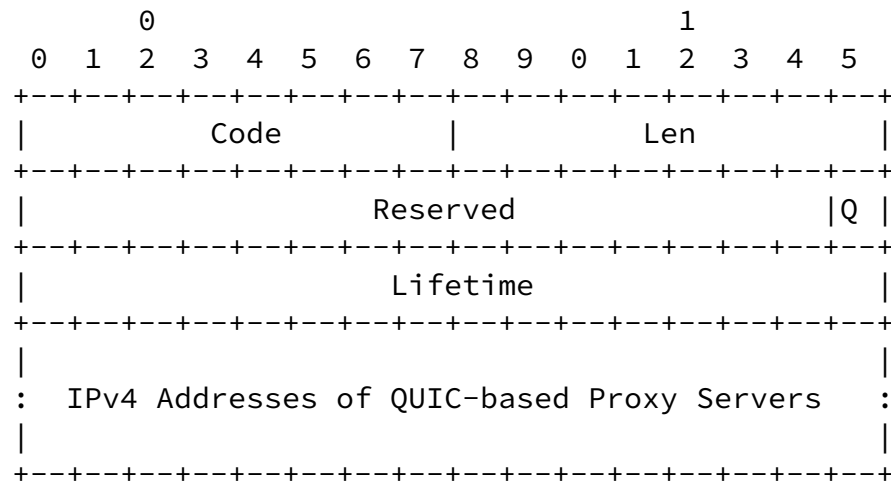


Figure 1: IPv4 Proxy Discovery DHCP option format

Code: Proxy Discovery option code (TBD) (8 bit)

Len: length of the option (without the Code and Len fields) in units of octets. The minimum value is 8 if one IPv4 address is contained in the option. Every additional IPv4 address increases the length by 4. (8-bit unsigned integer)

Q: is set to one if proxy supports QUIC on port 443 (1 bit)

Lifetime: maximum time in seconds (relative to the time the packet is received) over which these IPv4 addresses can be used for proxy discovery. A value of all one bits (0xffff) represents infinity. A value of zero means that the proxy addresses SHOULD no longer be used. (16-bit unsigned integer)

IPv4 Addresses of QUIC-based Proxy Servers: one or more 64-bit IPv4 addresses of QUIC-based proxy servers. The number of addresses is determined by the Length field. That is, the number of addresses is equal to $(\text{Length} - 4) / 4$.

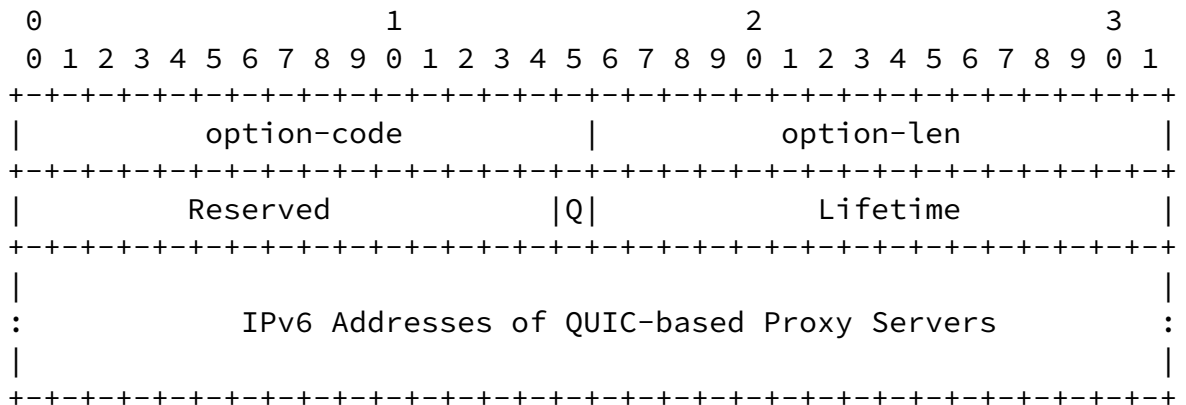


Figure 2: IPv6 Proxy Discovery DHCP option format

option-code: Proxy Discovery option code (TBD) (16 bit)

option-len: length of the option (without the Type and Length fields) in units of octets. The minimum value is 20 if one IPv6 address is contained in the option. Every additional IPv6 address increases the length by 16. (16-bit unsigned integer)

Q: is set to one if proxy supports QUIC on port 443 (1 bit)

Lifetime: maximum time in seconds (relative to the time the packet is received) over which these IPv6 addresses can be used for proxy discovery. A value of all one bits (0xffff) represents infinity. A value of zero means that the proxy addresses SHOULD no longer be used. (16-bit unsigned integer)

IPv6 Addresses of QUIC-based Proxy Servers: one or more 128-bit IPv6 addresses of QUIC-based proxy servers. The number of addresses is determined by the Length field. That is, the number of addresses is equal to (Length - 4) / 16.

3. Using IPv6 Neighbor Discovery for Local Discovery

If a proxy is located in the local network, information to discover a proxy service can be provided in a new Router Advertisement (RA) Option [[RFC4861](#)], the Proxy Discovery option.

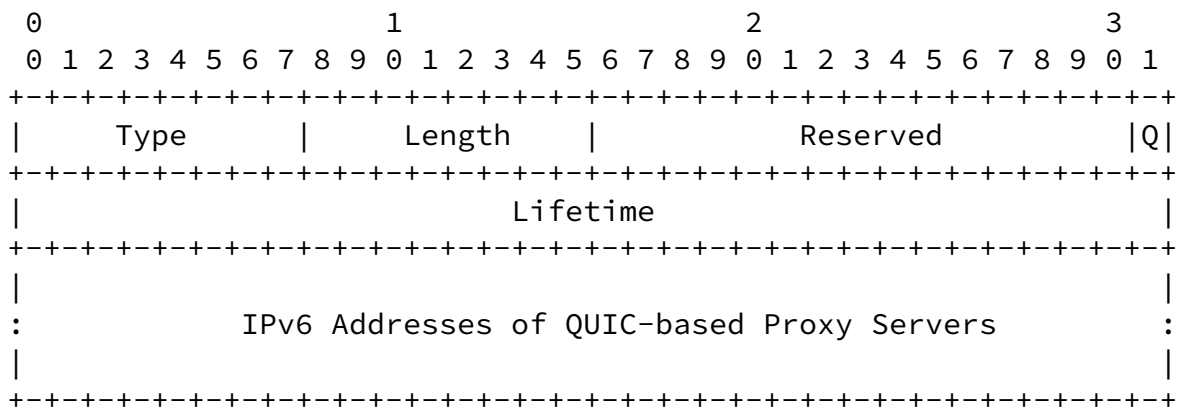


Figure 3: Proxy Discovery RA option format

Type: Proxy Discovery option type (TBD) (8 bit)

Length: length of the option (including the Type and Length fields) in units of 8 octets. The minimum value is 3 if one IPv6 address is contained in the option. Every additional IPv6 address increases the length by 2. (8-bit unsigned integer)

Q: is set to one if proxy supports QUIC on port 443 (1 bit)

Lifetime: maximum time in seconds (relative to the time the packet is received) over which these IPv6 addresses can be used for proxy discovery. A value of all one bits (0xffffffff) represents infinity. A value of zero means that the proxy addresses SHOULD no longer be used. (32-bit unsigned integer)

IPv6 Addresses of QUIC-based Proxy Servers: one or more 128-bit IPv6 addresses of QUIC-based proxy servers. The number of addresses is determined by the Length field. That is, the number of addresses is equal to $(\text{Length} - 1) / 2$.

[3.1. Using PVDs](#)

If the local network provides configuration with an Explicit Provisioning Domain (PvD) [[I-D.ietf-intarea-provisioning-domains](#)], the RA defined above can be used with the PvD Option or alternatively proxy information can be retrieved in the additional information JSON files associated with the PvD ID. The endhost resolves the URL provided in the PvD ID into an IP address using the local DNS server that is associated with the corresponding PvD (see also section 3.4.4 of [[I-D.ietf-intarea-provisioning-domains](#)]). If a QUIC-based proxy services is provided the additional information JSON file contains the key "QuicProxyIP". It can then optionally also contain more information about the specific proxy services offered using the "ProxyService" key. Or the client can connect directly to the proxy

over QUIC on port 443 and request information about the proxy service directly from the proxy server.

For remote network a Web PvD might be available that contains proxy information. If provided, the PvD JSON configuration file retrievable at the URI with the format:

`https://<Domain>/well-known/pvd`

4. DNS Service Discovery (DNS-SD)

[RFC6763] describes the use of SRV records to discover the available instances of a type of service. To get a list of names of the available instance for a certain service a client requests records of type "PTR" (pointer from one name to another) in the DNS namespace [RFC1035] for a name containing the service and domain.

As specified in [RFC6763] the client can perform a PTR query for a list of available proxy instance in the following way:

```
_quicproxy._udp.<domain>
```

here the <domain> portion is the domain name where the service is registered. The domain name can be obtained via DHCP options or preconfigured.

The result of this PTR lookup is a set of zero or more PTR records giving Service Instance names. Then to contact a particular service, the client can query for the SRV [RFC2782] and TXT records of the selected service instance name. The SRV record contains the IP address of the proxy service instance as well as the port number. The port number of QUIC-based proxy is usually expected to be 443 but may differ. The TXT can contain additional information describing the kind of proxy services that is offered.

4.1. Local discovery using mDNS

[RFC6762] defines the use of ".local." for performing DNS like operations on the local link. Any DNS query for a name ending "local." will be sent to a predefined IPv4 or IPv6 link local multicast address.

To discover QUIC-based proxy services locally, the client request the PTR record for the name:

```
_quicproxy._udp.local.
```

The result of this PTR lookup is a set of zero or more PTR records

giving Service Instance Names of the form:

<Instance>._quicproxy._udp.local.

Editors' Note: Or _masque._udp ? Or _proxy._quic._udp or _quicproxy._http._udp ...? However in the later case the proxy should probably also actually offer a webpage...

4.2. Discovery for Remote Domains

If a client wants to discover a QUIC-based proxy server for a remote domain, this domain has to be known by the client, e.g. being preconfigured in the application.

5. Using PCP options

Port Control Protocol (PCP), described in [RFC6887], defines mechanism to do packet forwarding for different types of IPv4/Ipv6 Network Address Translators (NAT) or firewalls. Usual deployments of PCP include Carrier-Grade NAT (CGN), Customer-premises Equipment (CPE), or residential NATs. When PCP is used to control address translation and forwarding, the PCP server can also be used to announce the existence of a QUIC-based proxy to the client.

PCP allows options to be included in the PCP request and response header. To announce information from the PCP server to the client, information about who to find a the QUIC-based proxy can be included in the response header as an option. As [RFC6887] describes, the client will ignore any options that it does not understand. A new PCP option carrying QUIC-based proxy information is specified below.

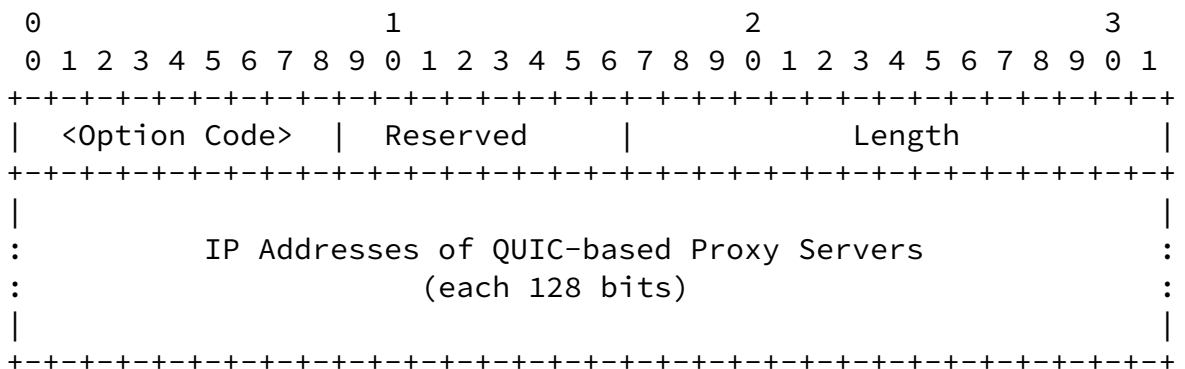


Figure 4: Proxy Discovery PCP option format

The fields are described below -

Option Code: 8 bits. The most significant bit indicates if this option is mandatory (0) or optional (1) to process.

Reserved: 8 bits. MUST be set to 0 on transmission and MUST be ignored on reception.

Option Length:

:16 bits. Indicates the length of the enclosed data, in octets. Options with length of 0 are allowed. Options that are not a multiple of 4 octets long are followed by one, two, or three 0 octets to pad their effective length in the packet to be a multiple of 4 octets. The Option Length reflects the semantic length of the option, not including any padding octets.

IP Addresses of QUIC-based Proxy Servers: one or more IPv6 addresses and/or IPv4 addresses of QUIC-based proxy servers. As specified in [section 5 of \[RFC6887\]](#) all addresses use fixed-size 128-bit fields. When the address field holds an IPv4 address, an IPv4-mapped IPv6 address [[RFC4291](#)] is used (::ffff:0:0/96). The number of addresses is determined by the Length field. That is, the number of addresses is equal to Length/16.

6. Using Anycast address

Well-known IP anycast addresses can be used to start communicating with QUIC proxy or to discovery any or a list of unicast address of a QUIC proxy. When the proxy receives the request for proxy functionalities, it can either decide to respond to the client with the anycast address as source address or it can send back a list of unicast address with a redirect command.

TODO: complete the description

7. IANA Considerations

IANA is requested to assign two DHCP options, one for IPv4 and one for IPv6, in the "BOOTP Vendor Extensions and DHCP Options" registry (<http://www.iana.org/assignments/bootp-dhcp-parameters>), as specified in [[RFC2939](#)], and the "Option Codes" registry under DHCPv6 parameters (<http://www.iana.org/assignments/dhcpv6-parameters>), respectively, as well a new value for the Proxy Discovery Option in the IPv6 Neighbor Discovery Option Formats registry.

This document adds a key to the "Additional Information PvD Keys" registry, defined by [[I-D.ietf-intarea-provisioning-domains](#)].

Internet-Draft QUIC Non-transparent Proxy Discovery January 2020

JSON key	Description	Type	Example
QuicProxyIP	IP address for QUIC-based proxies	Array of Strings	"["2001:db8:::1", "2001:db8:::2"]"
ProxyService	IDs identifying a specific service	Array of Strings	"["Forwarding", "DNSResolution"]"

Further, IANA is requested to register a new service name "quicproxy" in the "Service Name and Transport Protocol Port Number Registry" (<https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml>).

8. Security Consideration

Discovery mechanisms that are not authenticated provide no guarantees about the proxy configuration information provided. In some scenarios a client may decide to use this information anyway, as either the local environment that the discovery was performed in is trusted, or the client has means to authenticate the identify of the proxy when connecting using QUIC and only uses the discovery to dynamically detect an IP address.

Further even if the proxy is not trusted, simple forwarding or other network-based services may be used by the client if the forwarded traffic itself is end-to-end encrypted. In this case the trust level should not be assumed to be higher than in the connectivity case without proxy usage. Also note that even when the proxy is assumed to be untrusted, an attacker could still use the opportunity to redirect traffic over a specific node in order to more easily observe the traffic. However, in this case the client is at least aware of the use of the proxy and therefore has means to potentially even identify the proxy provider, e.g. based on the IP or certificate.

For further discussion of the security of each discovery mechanism, see also the security consideration section of these specifications.

9. Contributors

[10.](#) Acknowledgments

[11.](#) References

[11.1.](#) Normative References

[I-D.ietf-intarea-provisioning-domains]

Pfister, P., Vyncke, E., Pauly, T., Schinazi, D., and W. Shao, "Discovering Provisioning Domain Names and Data", [draft-ietf-intarea-provisioning-domains-10](#) (work in progress), January 2020.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.

[RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", [RFC 2782](#), DOI 10.17487/RFC2782, February 2000, <<https://www.rfc-editor.org/info/rfc2782>>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", [RFC 6762](#), DOI 10.17487/RFC6762, February 2013,

<<https://www.rfc-editor.org/info/rfc6762>>.

- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", [RFC 6763](#), DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.
- [RFC6887] Wing, D., Ed., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", [RFC 6887](#), DOI 10.17487/RFC6887, April 2013, <<https://www.rfc-editor.org/info/rfc6887>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 8415](#), DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.

[11.2.](#) Informative References

- [I-D.kuehlewind-quic-substrate]
Kuehlewind, M., Sarker, Z., Fossati, T., and L. Pardue, "Use Cases and Requirements for QUIC as a Substrate", [draft-kuehlewind-quic-substrate-02](#) (work in progress), November 2019.
- [I-D.schinazi-masque]
Schinazi, D., "The MASQUE Protocol", [draft-schinazi-masque-02](#) (work in progress), January 2020.
- [RFC2939] Droms, R., "Procedures and IANA Guidelines for Definition of New DHCP Options and Message Types", [BCP 43](#), [RFC 2939](#), DOI 10.17487/RFC2939, September 2000, <<https://www.rfc-editor.org/info/rfc2939>>.

Authors' Addresses

Mirja Kuehlewind
Ericsson

Email: mirja.kuehlewind@ericsson.com

Zaheduzzaman Sarker
Ericsson

Email: zaheduzzaman.sarker@ericsson.com