Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 7 January 2026

T. Jensen

D. Redekop
Adam Networks Inc.

J. Todd
Quad9

6 July 2025

# Associated IP Prefixes for Domain Names
## draft-tdj-dnsop-associated-prefixes-for-domains-00

## Abstract

RFC9000 defines a mechanism that allows servers to migrate clients to another IP address without name resolution. The new address may not be discoverable as A/AAAA records for that domain name. This draft defines a mechanism that allows a client to get advance notice of associated IP addresses for a domain name as part of the DNS query.

## 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 7 January 2026.

Table of Contents



## 1. Introduction

It is common for services to be associated with domain names even if not every IP address used by the service are represented in A and AAAA records for the service's domain names. One example is teleconferencing, which often uses a media streaming protocol whose peer addresses are negotiated within a connection, such as the use of WebRTC. Another is QUIC's use of preferred addresses, defined in Section 9.6 of [RFC9000], which allows a server to migrate a client to another server IP address which may or may not have been resolvable from the domain name resolved to initiate the original connection.

This document defines a mechanism for domain name owners to advertise any IP address prefixes that are associated with a domain name. This allows client peers to predict which IP addresses may end up in use when contacting a given service.

## 2. Conventions and Definitions

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

## 3. Terminology

TBD

## 4. Mechanism

When a client wishes to discover the IP prefixes that are associated with a given domain name, it SHOULD issue a SVCB query for the domain name (which it may already be doing for other reasons). If the "cidrs" key is present, then it SHOULD issue another query of type CIDRS to retrieve the associated IP prefixes.

When a server wishes to provide the associated IP prefixes for a given name, it SHOULD create CIDRS records as well as a SVCB record with the "cidrs" key set. While the mechanism is simple, the Operational Considerations section contains further normative implementation guidance.

### 4.1. CIDRS record format

The CIDRS RR type is designed to convey an IP prefix, an associated port range, and a protocol number as defined in the Assigned Internet Protocol Numbers IANA registry. The port range can be defined as 0-65535 and the protocol number as 255 if these fields are not applicable for the record owner (because only the IP prefix is interesting).

#### 4.1.1. Wire format

CIDRS RR data on the wire uses the following format (TODO: ASCII art of the fields below):

* IP family (4 bits): indicates which version of the Internet Protocol the prefix belongs to

* Reserved for future use (4 bits)

* Length of Prefix (8 bits): number of bits of the IP prefix.

* Prefix (ceil(Length of Prefix / 8) bytes): the IP prefix itself. The length of this field is always an integer of bytes (begins and ends on octent boundaries) just large enough to contain a prefix of length equal to the Length of Prefix

* Port Range Begin (16 bits): the lowest port number in the included range

* Port Range End (16 bits): the highest port number in the included range

* Protocol Number (8 bits): a value from the Assigned Internet Protocol Numbers IANA registry (255 if no associated protocol should be defined)

#### 4.1.2. Zone file format

TODO: will do after the idea survives initial contact with the WG

#### 4.1.3. Justification for new RR type

Because the list of associated IP addresses for a given name is likely to be somewhat large, at least relative to the typical A, AAAA, and SVCB/HTTPS queries a client is likely to make, it makes sense to create a new type that ensures DNS query issuers can opt into receiving this payload. Additionally, inclusion within SVCB or HTTPS records is a dubious reuse given there is no guarantee that services at associated IP addresses can authenticate a name (if the resulting protocol does not perform name verification).

#### 4.1.4. Verification

CIDRS records MUST be DNSSEC signed. This is because unlike A and AAAA records, there is no expectation that the resulting traffic the querying client will send to these IP addresses will be able to prevent attacker impersonation via secure peer validation such as that provided by TLS certificates. By definition, the addresses in these CIDRs are used in association with services that use the domain name but cannot validate claims of the domain name. DNSSEC validation will provide assurance that the IP addresses are those expected by the valid owner of the domain name.

## 5. Security Considerations

### 5.1. Suspiciously large IP address sets

Servers SHOULD avoid claiming very short prefixes in CIDRS records. It is not expected that a single domain name is legitimately associated with a short prefix. Clients MAY choose to ignore claims of association by CIDRS records with prefixes shorter than a preconfigured minimum length per IP version. This version of the text does not suggest defined values for minimum prefix lengths, though as a reasonable rule of thumb, domain name owners SHOULD NOT have prefixes in CIDRS records that include multiple ASNs to avoid clients considering their prefixes to be too short.

Note that associated IP addresses SHOULD be restricted to IP addresses which a server reasonably expects a client will need to interact with the functionality provided by the service which uses the domain name. For example, name owners SHOULD NOT create CIDRS records that include all IP ranges owned by a company for the company’s primarily recognizable domain name (example-company.example. having a CIDRS record listing every IP address owned by Example Company would be inappropriate).

### 5.2. Large payload sizes

It is expected that in most common use cases, CIDRS records will need more than one CIDR value, possibly many (balancing this against guidance given in Section 6.2). Servers SHOULD truncate responses to avoid creating risk of effective DDoS attacks, even if the CIDRS record would fit in a single UDP packet. This means in effect that CIDRS records SHOULD NOT ever be sent using unencrypted DNS over UDP.

## 6. Operational Considerations

### 6.1. TTLs for SVCB and CIDRS records

Servers SHOULD NOT have a TTL value for SVCB records containing the "cidrs" key that are less than the TTL values for CIDRS records for the same domain name. This would cause the client to re-check for the presence of CIDRS records that it will still have cached.

The TTL values of CIDRS records SHOULD NOT be any shorter than the expected lifetime of traffic flows of typical service usage. Doing so would encourage DNS stub resolvers and the processes calling DNS sutb resolver APIs to ignore TTL values in favor of supporting performant user experiences.

### 6.2. Complex messages

[RFC3123] defines an experimental mechanism by which the APL RR type can both define CIDRs that are somehow associated with a domain name as well as negate subsets of the CIDR. This is specifically not supported for the CIDRS RR type because of the added complexity this creates for implementors and the support for defining non-consecutive subnets for the same purpose. Implementors SHOULD reduce the number of CIDRs needed for a given domain name rather than have many long prefixes that cannot be grouped under fewer CIDRs without needing to define the non-included gaps within them.

### 6.3. Transition period workarounds

Before deployment of this document is common, clients will frequently run into the problem of wanting to discover associated IP addresses for a given zone, but the zone owner does not yet support this document. During this ramp up of deployment, administrators might turn to workarounds, such as creating a zone they control to distribute CIDRs the administrator knows to be associated with services they depend on but do not yet support this document.

How to accomplish such a mapping is left to implementors as a non-standard mechanism. This is out of scope for this document, which only defines advertisement of IP addresses associated with a given name directly.

## 7. Security Considerations

TODO Security

## 8. IANA Considerations

TODO: new DNS RR type "CIDRS"

TODO: new SVCB key "cidrs"

## 9. References

### 9.1. 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/rfc/rfc2119>.

[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/rfc/rfc8174>.

### 9.2. Informative References

[RFC3123] Koch, P., "A DNS RR Type for Lists of Address Prefixes (APL RR)", RFC 3123, DOI 10.17487/RFC3123, June 2001, <https://www.rfc-editor.org/rfc/rfc3123>.

[RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <https://www.rfc-editor.org/rfc/rfc9000>.

## Acknowledgments

TODO acknowledge.

## Authors' Addresses

Tommy Jensen
Email: tojens.ietf@gmail.com

David Redekop
Adam Networks Inc.
Email: david.ietf@adamnet.works

John Todd
Quad9
Email: jtodd@quad9.net