AVI Networks®

Extension Mechanisms for DNS (EDNS) Client Subnet Option Insertion

Avi Technical Reference (v17.1)

Extension Mechanisms for DNS (EDNS) Client Subnet Option Insertion

view online

As of release 17.1.3, Avi Vantage supports insertion of the ECS option in a DNS query if the query has no ECS option. As of 17.1.4, it supports updating of the ECS option if the DNS query already has an ECS option.

Background

Since it was first developed, the Domain Name System has required enhancements. Restrictions in the size of several flags fields, return codes and label types available in the basic DNS protocol has motivated extending DNS in a backward-compatible fashion to allow for new flags and response codes, and to provide support for longer responses. Since 1999, extension mechanisms for DNS (EDNS) has been the approach taken to address this challenge.

As shown in Figure 1, the OPT resource record (OPT RR) is key to extending DNS. It is structured to permit various options, including the EDNS client subnet option (ECS), which allows authoritative DNS providers to use the extra information to make more informed traffic routing decisions, for example, * To provide more accurate client location information for use by the geo algorithm * To convey the client's source-IP address for use by the consistent-hash algorithm * When serving clients coming from a mixture of private and public networks

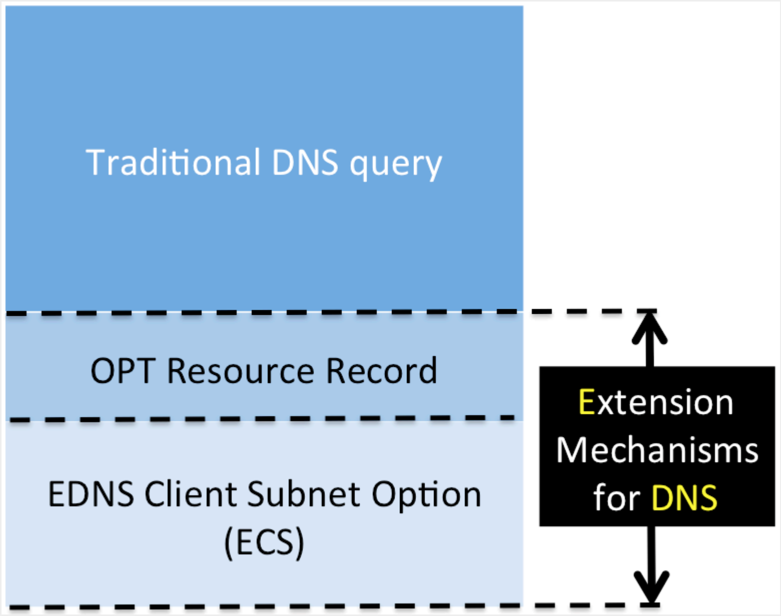

Figure 1. How the traditional DNS query can be extended

How EDNS and the ECS Option Work with Avi DNS

Enabling EDNS on the Avi DNS Virtual Service

As mentioned, the Avi DNS virtual service can directly profit from the information in OPT RR and the ECS option while acting as an authoritative DNS. To have it parse that information and append EDNS extension information into the client logs, check the Process EDNS Extensions box in the Avi UI as shown in Figure 2, or set the corresponding `edns` parameter to True in the Avi CLI.

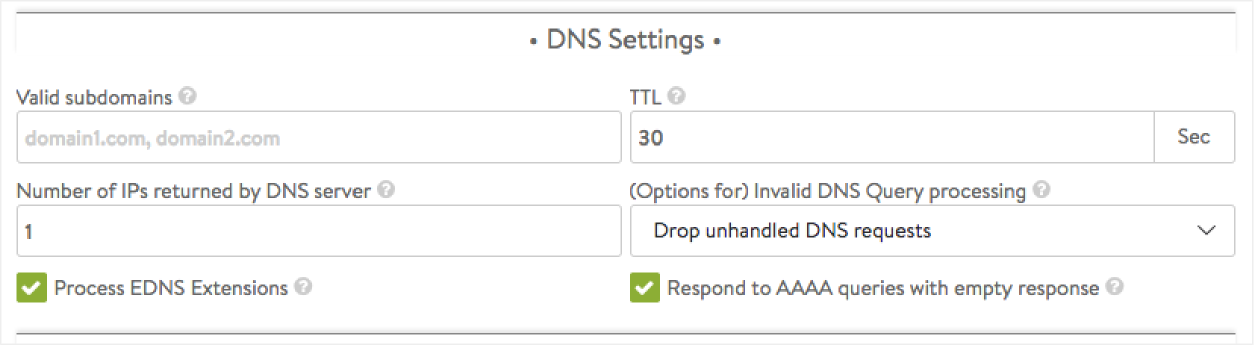

• DNS Settings •

Valid subdomains: domain1.com, domain2.com

TTL: 30 Sec

Number of IPs returned by DNS server: 1

(Options for) Invalid DNS Query processing: Drop unhandled DNS requests

☑ Process EDNS Extensions

☑ Respond to AAAA queries with empty response

Figure 2. DNS settings, as revealed by the application profile editor

Case 1: Avi DNS Virtual Service is Authoritative, OPT RR + ECS Option Received

In addition to checking the Process EDNS Extensions box having been checked, ensure that a list of authoritative domain names has been provided, as for example in the below Avi CLI sequence.

```
{%cli%} [admin:10-10-25-20]: > configure applicationprofile System-DNS [admin:10-10-25-20]: applicationprofile>
dns_service_profile [admin:10-10-25-20]: applicationprofile:dns_service_profile> authoritative_domain_names avi.com [admin:
10-10-25-20]: applicationprofile:dns_service_profile> authoritative_domain_names foo.com [admin:10-10-25-20]:
applicationprofile:dns_service_profile> save [admin:10-10-25-20]: applicationprofile> save {%endcli%}
```

Now suppose that the in-bound DNS request shown in Figure 3 is for one of these domains.

- The client system sends a traditional DNS query to its DNS resolver. Note that the request it sends contains neither an OPT RR nor an ECS option.
- Based on the source address of the client, the DNS resolver may amend the DNS query it receives. It does this to enable the authoritative DNS to respond in a more informed way, i.e., based on the address of the client, as opposed to the source IP of the DNS resolver itself.
- The Avi DNS forms the response based on address information it finds within the ECS option.

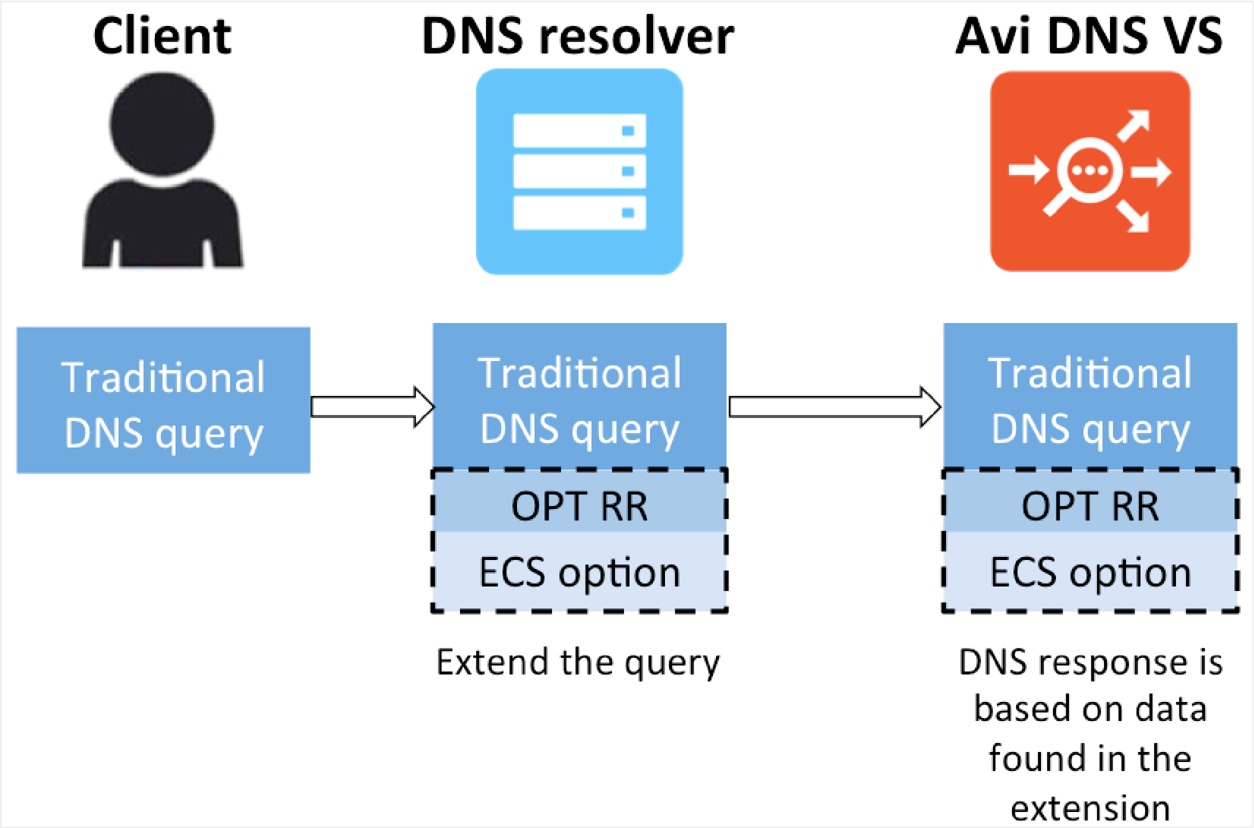

Figure 3. Authoritative Avi DNS responds to query based on ECS option

Case 2: Avi DNS Virtual Service is not Authoritative, OPT RR + ECS Option Received

In contrast to Case 1, in Figure 4 we see a query for which the Avi DNS VS is not authoritative. If a DNS server pool has been defined for the Avi DNS VS, the request will be passed through to it. The client subnet address information incorporated in the ECS option of the forwarded request depends on two values:

1. The value of the subnet prefix length parameter contained within the ECS option attached and sent by the DNS resolver, and
2. `edns_client_subnet_prefix_len`, an Avi DNS VS application profile parameter set by the administrator via the CLI. Its value ranges between 1 and 32.

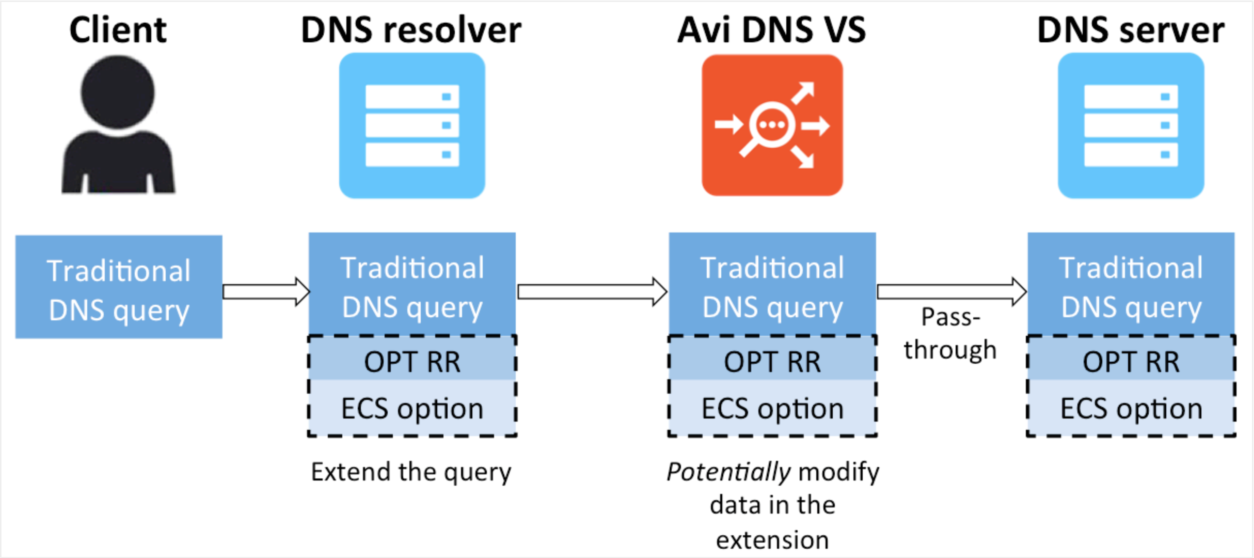

Figure 4. Avi DNS passes DNS request on to DNS server

Either prefix length is interpreted two ways:

1. It indicates the leading number of address bits after which all address bits are zero, and
2. When rounded up to an integer multiple of eight bits, it specifies the number of octets needing to be passed.

As an example, a prefix length of 19 implies the following about the subnet: * Bits 20 through 32 in the subnet address are zero, and * 24 bits, i.e., three octets need only be passed to identify the subnet. A fourth octet would be superfluous.

When passing the ECS option through to the DNS server, the Avi DNS VS will ensure the client subnet address is governed by the *lesser* of the two prefix lengths. That is to say:

- If the incoming subnet prefix length is *less* than the value of the Avi DNS's `edns_client_subnet_prefix_len` parameter, the ECS option will be untouched as it passes through.
- If the incoming subnet prefix length (e.g., 26) is *greater* than the value of the Avi DNS's `edns_client_subnet_prefix_len` parameter (e.g., 16), Avi will zero out some incoming bits (e.g., 10 in this case), and, if the lengths are sufficiently far apart, forward fewer octets (e.g., 2 not 4) to the DNS server shown on the right of Figure 3.

Case 3: Avi DNS Virtual Service is not Authoritative, Neither OPT RR nor ECS Option Received

In Figure 5, we see the DNS request arriving from the DNS resolver with no EDNS information whatsover. In addition, the DNS request is for a domain for which the Avi DNS is *not* authoritative. Thus, a pass-through is required. In this case, the Avi DNS VS will create an OPT RR, and, for the ECS option, insert a client subnet address with 1 to 4 octets and an appropriate number of trailing zeroes, in a fashion as described above.

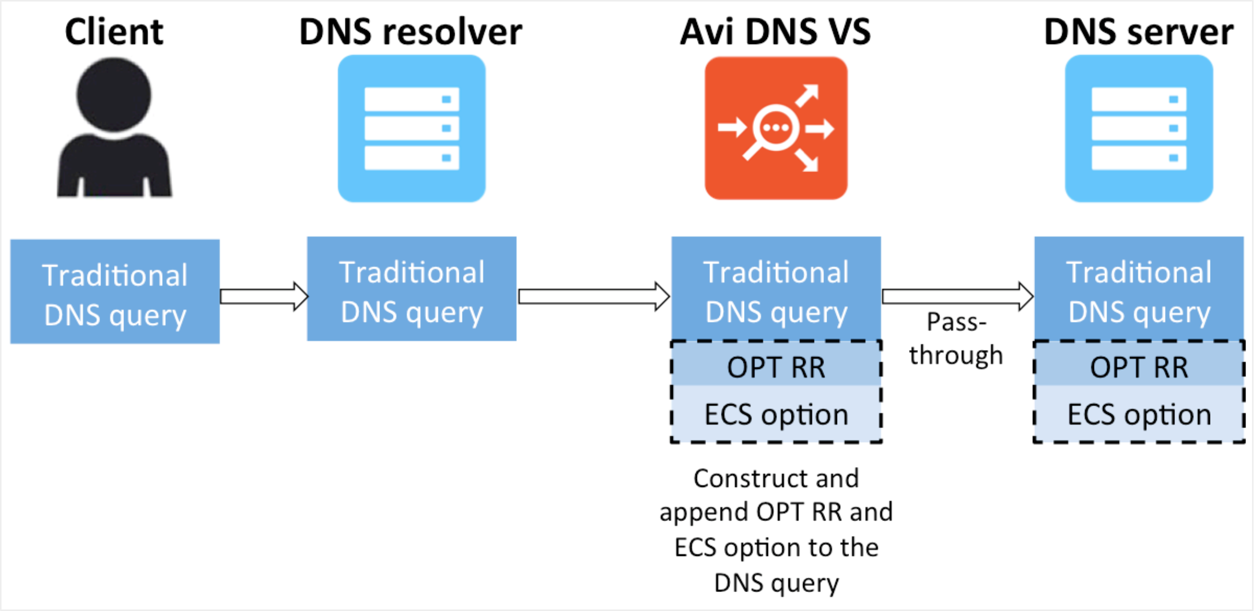

Figure 5. Avi DNS appends OPT RR and ECS option to a forwarded DNS query

Related RFCs

- RFC6891: Extension Mechanisms for DNS (EDNS(0))
- RFC7871: Client Subnet in DNS Queries