

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 7, 2016

T. Lemon
Nominum, Inc.
T. Mrugalski
ISC
July 6, 2015

Customizing DHCP Configuration on the Basis of Network Topology
draft-ietf-dhc-topo-conf-05

Abstract

DHCP servers have evolved over the years to provide significant functionality beyond that which is described in the DHCP base specifications. One aspect of this functionality is support for context-specific configuration information. This memo describes some such features and makes recommendations as to how they can be used.

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 January 7, 2016.

Copyright Notice

Copyright (c) 2015 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 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.	Terminology	3
3.	Identifying Client's Location by DHCP Servers	3
3.1.	DHCPv4 Specific Behavior	7
3.2.	DHCPv6 Specific Behavior	7
4.	Simple Subnetted Network	9
5.	Relay agent running on a host	11
6.	Cascade relays	11
7.	Regional Configuration Example	12
8.	Dynamic Lookup	14
9.	Multiple subnets on the same link	15
10.	Acknowledgments	15
11.	Security Considerations	16
12.	IANA Considerations	16
13.	References	16
13.1.	Normative References	16
13.2.	Informative References	16
	Authors' Addresses	17

1. Introduction

The DHCPv4 [[RFC2131](#)] and DHCPv6 [[RFC3315](#)] protocol specifications describe how addresses can be allocated to clients based on network topology information provided by the DHCP relay infrastructure. Address allocation decisions are integral to the allocation of addresses and prefixes in DHCP.

The DHCP protocol also describes mechanisms for provisioning devices with additional configuration information; for example, DNS [[RFC1034](#)] server addresses, default DNS search domains, and similar information.

Although it was the intent of the authors of these specifications that DHCP servers would provision devices with configuration information appropriate to each device's location on the network, this practice was never documented, much less described in detail.

Existing DHCP server implementations do in fact provide such capabilities; the goal of this document is to describe those capabilities for the benefit both of operators and of protocol designers who may wish to use DHCP as a means for configuring their own services, but may not be aware of the capabilities provided by most modern DHCP servers.

2. Terminology

- o Routable IP address: an IP address with a scope of use wider than the local link.
- o PE router: provider edge router. The provider router closest to the customer.
- o CPE device: customer premise equipment device. Typically a router belonging to the customer that connects directly to the provider link.
- o Shared subnet: a case where two or more subnets of the same protocol family are available on the same link. 'Shared subnet' terminology is typically used in Unix environments. It is typically called 'multinet' in Windows environment. The administrative configuration inside a Microsoft DHCP server is called 'DHCP Superscope'.

3. Identifying Client's Location by DHCP Servers

Figure 1 illustrates a small hierarchy of network links with Link D serving as a backbone to which the DHCP server is attached.

Figure 2 illustrates a more complex case. Although some of its aspects are unlikely to be seen in an actual production networks, they are beneficial for explaining finer aspects of the DHCP protocols. Note that some nodes act as routers (which forward all IPv6 traffic) and some are relay agents (i.e. run DHCPv6 specific software that forwards only DHCPv6 traffic).




Figure 1: A simple network with a small hierarchy of links




Figure 2: Complex network

Those diagrams allow us to represent a variety of different network configurations and illustrate how existing DHCP servers can provide configuration information customized to the particular location from which a client is making its request.

It is important to understand the background of how DHCP works when considering those diagrams. It is assumed that the DHCP clients may not have routable IP addresses when they are attempting to obtain configuration information.

The reason for making this assumption is that one of the functions of DHCP is to bootstrap the DHCP client's IP address configuration; if the client does not yet have an IP address configured, it cannot route packets to an off-link DHCP server, therefore some kind of relay mechanism is required.

The details of how packet delivery between clients and servers works are different between DHCPv4 and DHCPv6, but the essence is the same: whether or not the client actually has an IP configuration, it generally communicates with the DHCP server by sending its requests to a DHCP relay agent on the local link; this relay agent, which has a routable IP address, then forwards the DHCP requests to the DHCP server (directly or via other relays). In later stages of the configuration when the client has acquired an address and certain conditions are met, it is possible for the client to send packets directly to the server, thus bypassing the relays. The conditions for such behavior are different for DHCPv4 and DHCPv6 and are discussed in sections [Section 3.1](#) and [Section 3.2](#).

The DHCP server uses an IP address from the client's message which is on the same link as the client to perform address assignment decisions or to select subnet-specific configuration for the client. The address that the server uses is the DHCP client's routable IP address or the client facing address of the relay agent. The server is then able to determine the client's point of attachment and select appropriate subnet- or link-specific configuration.

Sometimes it is useful for the relay agents to provide additional information about the topology. A number of extensions have been defined for this purpose. The specifics are different, but the core principle remains the same: the relay agent knows exactly where the original request came from, so it provides an identifier that will help the server to choose appropriate address pool and configuration parameters. Examples of such options are mentioned in the following sections.

Finally, clients may be connected to the same link as the server, so no relay agents are required. In such cases, the DHCPv4 server typically uses the IPv4 address assigned to the network interface over which the transmission was received to select appropriate subnet. This is more complicated for DHCPv6, as the DHCPv6 server is not required to have any globally unique addresses. In such cases, an additional configuration information may be required. Some

servers allow indicating that a given subnet is directly reachable over specific local network interface.

3.1. DHCPv4 Specific Behavior

In some cases in DHCPv4, when a DHCPv4 client has a routable IPv4 address, the message is unicast to the DHCPv4 server rather than going through a relay agent. Examples of such transmissions are renewal (DHCPREQUEST) and address release (DHCPRELEASE).

The relay agent that receives client's message sets GIADDR field to the address of the network interface the message was received on. The relay agent may insert a relay agent option [[RFC3046](#)].

There are several options defined that are useful for subnet selection in DHCPv4. [[RFC3527](#)] defines Link Selection sub-option that is inserted by a relay agent. This option is particularly useful when the relay agent needs to specify the subnet/link on which a DHCP client resides, which is different from an IP address that can be used to communicate with the relay agent. Virtual Subnet Selection Option, specified in [[RFC6607](#)] is used for the same purpose (i.e. relay agents insert that information), but it also covers additional use cases in VPN environment. In certain cases it is useful for the client itself to specify this option, e.g. when there are no relay agents involved during VPN set up process.

Another option that may influence the subnet selection is IPv4 Subnet Selection Option, defined in [[RFC3011](#)], which allows the client to explicitly request allocation from a given subnet.

3.2. DHCPv6 Specific Behavior

In DHCPv6 unicast communication is possible in case where the server is configured with a Server Unicast option (see [Section 22.12 in \[RFC3315\]](#)) and clients are able to take advantage of it. In such cases, once a client is assigned a, presumably global, address, it is able to contact the server directly, bypassing any relays. It should be noted that such a mode is completely controllable by administrators in DHCPv6. (They may simply choose to not configure server unicast option, thus forcing clients to send their messages always via relay agents in every case).

In the DHCPv6 protocol, there are two core mechanisms defined in [[RFC3315](#)] that allow server to distinguish which link the relay agent is connected to. The first mechanism is a link-address field in the Relay-forward and Relay-reply messages. Somewhat contrary to its name, relay agents insert in the link-address field an address that is typically global and can be used to uniquely identify the link on

which the client is located. In normal circumstances this is the solution that is easiest to maintain, as existing address assignments can be used and no additional administrative actions (like assigning dedicated identifiers for each relay agent, making sure they are unique and maintaining a list of such identifiers) are needed. It requires, however, for the relay agent to have an address with a scope larger than link-local configured on its client-facing interface.

If for whatever reason that is not feasible (e.g. because the relay agent does not have a global address or ULA [[RFC4193](#)] configured on the client-facing interface), the relay agent includes an Interface-Id option (see [Section 22.18 of \[RFC3315\]](#)) that identifies the link clients are connected to. If the interface-id is unique within an administrative domain, the interface-id value may be used to select the appropriate subnet. As there is no guarantee for the uniqueness ([\[RFC3315\]](#) only mandates the interface-id to be unique within a single relay agent context), it is up to the administrator to check whether the relay agents deployed use unique interface-id values. If they aren't, Interface-id cannot be used to determine client's point of attachment.

It should be noted that Relay-forward and Relay-reply messages are exchanged between relays and servers only. Clients are never exposed to those messages. Also, servers never receive Relay-reply messages. Relay agents must be able to process both Relay-forward (sending already relayed message further towards the server, when there is more than one relay agent in a chain) and Relay-reply (when sending back the response towards the client, when there is more than one relay agent in a chain).

For completeness, we also mention an uncommon, but valid case, where relay agents set link-local address in the link-address field in relayed Relay-forward messages. This may happen if the relay agent doesn't have any address with a larger scope. Even though link local addresses cannot be automatically used to associate relay agent with a given link, with sufficient information provided the server is still able to correctly select the proper link. That requires the DHCP server software to be able to specify relay agent link-address or a feature similar to 'shared subnets' (see [Section 9](#)). Network administrator has to manually configure additional information that a given subnet uses a relay agent with link-address X. Alternatively, if the relay agent uses link address X and relays messages from a subnet A, an administrator can configure that subnet A is a shared subnet with a very small X/128 subnet. That is not a recommended configuration, but in cases where it is impossible for relay agents to get an address from the subnet they are relaying from, it may be a viable solution.

DHCPv6 also has support for more finely grained link identification, using Lightweight DHCPv6 Relay Agents [[RFC6221](#)] (LDRA). In this case, the link-address field is set to Unspecified_address (::), but the DHCPv6 server also receives an Interface-Id option from the relay agent that can be used to more precisely identify the client's location on the network.

What this means in practice is that the DHCP server in all cases has sufficient information to pinpoint, at the very least, the layer 3 link to which the client is connected, and in some cases which layer 2 link the client is connected to, when the layer 3 link is aggregated out of multiple layer 2 links.

In all cases, then, the DHCP server will have a link-identifying IP address, and in some cases it may also have a link-specific identifier (e.g. Interface-Id Option or Link Address Option defined in [Section 5 of \[RFC6977\]](#)). It should be noted that there the link-specific identifier is unique only within the scope of the link-identifying IP address. For example, link-specific identifier of "eth0" for a relay agent with IPv4 address 192.0.2.1 means something different than "eth0" for a relay agent with address 192.0.2.123.

It is also possible for link-specific identifiers to be nested, so that the actual identifier that identifies the link is an aggregate of two or more link-specific identifiers sent by a set of LDRAs in a chain; in general this functions exactly as if a single identifier were received from a single LDRA, so we do not treat it specially in the discussion below, but sites that use chained LDRA configurations will need to be aware of this when configuring their DHCP servers.

The Virtual Subnet Selection Options, present in DHCPv4, are also defined for DHCPv6. The use case is the same as in DHCPv4: the relay agent inserts VSS options that can help the server to select the appropriate subnet with its address pool and associated configuration options. See [[RFC6607](#)] for details.

4. Simple Subnetted Network

Consider Figure 1 in the context of a simple subnetted network. In this network, there are four leaf subnets: links A, B, F and G, on which DHCP clients will be configured. Relays A, B, C and D in this example are represented in the diagram as IP routers with an embedded relay function, because this is a very typical configuration, but the relay function can also be provided in a separate node on each link.

In a simple network like this, there may be no need for link-specific configuration in DHCPv6, since local routing information is delivered through router advertisements. However, in IPv4, it is very typical

to configure the default route using DHCP; in this case, the default route will be different on each link. In order to accomplish this, the DHCP server will need link-specific configuration for the default route.

To illustrate, we will use an example from a hypothetical DHCP server that uses a simple JSON notation [[RFC7159](#)] for configuration. Although we know of no DHCP server that uses this specific syntax, most modern DHCP server provides similar functionality.

```
{
  "prefixes": {
    "192.0.2.0/26": {
      "options": {
        "routers": ["192.0.2.1"]
      },
      "on-link": ["A"]
    },
    "192.0.2.64/26": {
      "options": {
        "routers": ["192.0.2.65"]
      },
      "on-link": ["B"]
    },
    "192.0.2.128/26": {
      "options": {
        "routers": ["192.0.2.129"]
      },
      "on-link": ["F"]
    },
    "192.0.2.192/26": {
      "options": {
        "routers": ["192.0.2.193"]
      },
      "on-link": ["G"]
    }
  }
}
```

Figure 3: Configuration example

In Figure 3, we see a configuration example for this scenario: a set of prefixes, each of which has a set of options and a list of links for which it is on-link. We have defined one option for each prefix: a routers option. This option contains a list of values; each list only has one value, and that value is the IP address of the router specific to the prefix.

When the DHCP server receives a request, it searches the list of prefixes for one that encloses the link-identifying IP address provided by the client or relay agent. The DHCP server then examines the options list associated with that prefix and returns those options to the client.

So for example a client connected to link A in the example would have a link-identifying IP address within the 192.0.2.0/26 prefix, so the DHCP server would match it to that prefix. Based on the configuration, the DHCP server would then return a routers option containing a single IP address: 192.0.2.1. A client on link F would have a link-identifying address in the 192.0.2.128/26 prefix, and would receive a routers option containing the IP address 192.0.2.129.

5. Relay agent running on a host

A relay agent is a DHCP software that may be run on any IP node. Although it is typically run on a router, this is by no means required by the DHCP protocol. The relay agent is simply a service that operates on a link, receiving link-local multicasts (IPv6) or broadcasts (IPv4) and relaying them, using IP routing, to a DHCP server. As long as the relay has an IP address on the link, and a default route or more specific route through which it can reach a DHCP server, it need not be a router, or even have multiple interfaces.

A relay agent can be run on a host connected to two links. That case is presented in Figure 2. There is router B that is connected to links D and E. At the same time there is also a host that is connected to the same links. The relay agent software is running on that host. That is uncommon, but a valid configuration.

6. Cascade relays

Let's observe another case, shown in Figure 2. Note that in this configuration, the clients connected to link G will send their requests to relay D which will forward its packets directly to the DHCP server. That is typical, but not the only possible configuration. It is possible to configure relay agent D to forward client messages to relay E which in turn will send it to the DHCP server. This configuration is sometimes referred to as cascade relay agents.

Note that the relaying mechanism works differently in DHCPv4 and in DHCPv6. In DHCPv4 only the first relay is able to set the GIADDR field in the DHCPv4 packet. Any following relays that receive that packet will not change it as the server needs GIADDR information from the first relay (i.e. the closest to the client). The server will

send the response back to the GIADDR address, which is the address of the first relay agent that saw the client's message. That means that the client messages travel on a different path than the server's responses. A message from client connected to link G will travel via relay D, relay E and to the server. A response message will be sent from the server to relay D via router B, and relay D will send it to the client on link G.

Relaying in DHCPv6 is more structured. Each relay agent encapsulates a packet that is destined to the server and sends it towards the server. Depending on the configuration, that can be a server's unicast address, a multicast address or next relay agent address. The next relay repeats the encapsulation process. Although the resulting packet is more complex (may have up to 32 levels of encapsulation if the packet traveled through 32 relays), every relay may insert its own options and it is clear which relay agent inserted which option.

7. Regional Configuration Example

In the Figure 2 example, link C is a regional backbone for an ISP. Link E is also a regional backbone for that ISP. Relays A, B, C and D are PE routers, and Links A, B, F and G are actually link aggregators with individual layer 2 circuits to each customer--for example, the relays might be DSLAMs or cable head-end systems. At each customer site we assume there is a single CPE device attached to the link.

We further assume that links A, B, F and G are each addressed by a single prefix, although it would be equally valid for each CPE device to be numbered on a separate prefix.

In a real-world deployment, there would likely be many more than two PE routers connected to each regional backbone; we have kept the number small for simplicity.

In the example presented in Figure 4, the goal is to configure all the devices within a region with server addresses local to that region, so that service traffic does not have to be routed between regions unnecessarily.


```
{
  "prefixes": {
    "2001:db8:0:0::/40": {
      "on-link": ["A"]
    },
    "2001:db8:100:0::/40": {
      "on-link": ["B"]
    },
    "2001:db8:200:0::/40": {
      "on-link": ["F"]
    },
    "2001:db8:300:0::/40": {
      "on-link": ["G"]
    }
  },
  "links": {
    "A": {"region": "omashu"},
    "B": {"region": "omashu"},
    "F": {"region": "gaoling"},
    "G": {"region": "gaoling"}
  },
  "regions": {
    "omashu": {
      "options": {
        "sip-servers": ["sip.omashu.example.org"],
        "dns-servers": ["dns1.omashu.example.org",
                        "dns2.omashu.example.org"]
      }
    },
    "gaoling": {
      "options": {
        "sip-servers": ["sip.gaoling.example.org"],
        "dns-servers": ["dns1.gaoling.example.org",
                        "dns2.gaoling.example.org"]
      }
    }
  }
}
```

Figure 4: An example regions configuration

In this example, when a request comes in to the DHCP server with a link-identifying IP address in the 2001:DB8:0:0::/40 prefix, it is identified as being on link A. The DHCP server then looks on the list of links to see what region the client is in. Link A is identified as being in omashu. The DHCP server then looks up omashu in the set of regions, and discovers a list of region-specific options.

The DHCP server then resolves the domain names listed in the options and sends a sip-server option containing the IP addresses that the resolver returned for sip.omashu.example.org, and a dns-server option containing the IP addresses returned by the resolver for dns1.omashu.example.org and dns2.omashu.example.org. Depending on the server capability and configuration, it may cache resolved responses for specific period of time, repeat queries every time or even keep the response until reconfiguration or shutdown.

Similarly, if the DHCP server receives a request from a DHCP client where the link-identifying IP address is contained by the prefix 2001:DB8:300:0::/40, then the DHCP server identifies the client as being connected to link G. The DHCP server then identifies link G as being in the gaoling region, and returns the sip-servers and dns-servers options specific to that region.

As with the previous example, the exact configuration syntax and structure shown above does not precisely match what existing DHCP servers do, but the behavior illustrated in this example can be accomplished with most existing modern DHCP servers.

8. Dynamic Lookup

In the Regional example, the configuration listed several domain names as values for the sip-servers and dns-servers options. The wire format of both of these options contains one or more IPv6 addresses--there is no way to return a domain name to the client.

This was understood to be an issue when the original DHCP protocol was defined, and historical implementations even from the very early days would accept domain names and resolve them. Some early DHCP implementations, particularly those based on earlier BOOTP implementations, had very limited capacity for reconfiguration.

However, most modern DHCP servers handle name resolution by querying the resolver each time a DHCP packet comes in. This means that if DHCP servers and DNS servers are managed by different administrative entities, there is no need for the administrators of the DHCP servers and DNS servers to communicate when changes are made. When changes are made to the DNS server, these changes are promptly and automatically adopted by the DHCP server, as long as the DNS server is managed appropriately (see the next paragraph). Similarly, when DHCP server configurations change, DNS server administrators need not be aware of this.

It should be noted that even though the DHCP server may be configured to query the DNS resolver every time it uses configured names, the changes made in the DNS zone may not be visible to the server until

the DNS cache expires. In general, it is the responsibility of the DNS zone's administrator to ensure that existing cache does not cause a trouble when a change is made to the zone; it should be usually reasonable for the DHCP server to rely on it. However, if this is not desired or if the management of the DNS zone is not very reliable, the DHCP server can be configured to query the authoritative DNS server directly, bypassing any caching DNS servers.

It is worth noting that DNS is not the only way to resolve names, and not all DHCP servers support other techniques (e.g., NIS+ or WINS). However, since these protocols have all but vanished from common use, this won't be an issue in new deployments.

9. Multiple subnets on the same link

There are scenarios where there is more than one subnet from the same protocol family (i.e. two or more IPv4 subnets or two or more IPv6 subnets) configured on the same layer 3 link. One example is a slow network renumbering where some services are migrated to the new addressing scheme, but some aren't yet. Second example is a cable network, where cable modems and the devices connected behind them are connected to the same layer 2 link. However, operators want the cable modems and user devices to get addresses from distinct address spaces, so users couldn't easily access their modems management interfaces. Such a configuration is often referred to as 'shared subnets' in Unix environments or 'multinet' in Microsoft terminology.

To support such a configuration, additional differentiating information is required. Many DHCP server implementations offer a feature that is typically called client classification. The server segregates incoming packets into one or more classes based on certain packet characteristics, e.g. presence or value of certain options or even a match between existing options. Servers require additional information to handle such configuration, as they cannot use the topographical property of the relay addresses alone to properly choose a subnet. Exact details of such operation is not part of the DHCPv4 or DHCPv6 protocols and is implementation dependent.

10. Acknowledgments

Thanks to Dave Thaler for suggesting that even though "everybody knows" how DHCP servers are deployed in the real world, it might be worthwhile to have an IETF document that explains what everybody knows, because in reality not everybody is an expert in how DHCP servers are administered. Thanks to Andre Kostur, Carsten Strotmann, Simon Perreault, Jinmei Tatuya, Suresh Krishnan, Qi Sun, Jean-Francois Tremblay, Marcin Siodelski and Bernie Volz for their reviews, comments and feedback.

11. Security Considerations

This document explains existing practice with respect to the use of Dynamic Host Configuration Protocol [[RFC2131](#)] and Dynamic Host Configuration Protocol Version 6 [[RFC3315](#)]. The security considerations for these protocols are described in their specifications and in related documents that extend these protocols. This document introduces no new functionality, and hence no new security considerations.

12. IANA Considerations

The IANA is hereby absolved of any requirement to take any action in relation to this document.

13. References

13.1. Normative References

- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.

13.2. Informative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC3011] Waters, G., "The IPv4 Subnet Selection Option for DHCP", [RFC 3011](#), November 2000.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", [RFC 3046](#), January 2001.
- [RFC3527] Kinnear, K., Stapp, M., Johnson, R., and J. Kumarasamy, "Link Selection sub-option for the Relay Agent Information Option for DHCPv4", [RFC 3527](#), April 2003.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.
- [RFC6221] Miles, D., Ooghe, S., Dec, W., Krishnan, S., and A. Kavanagh, "Lightweight DHCPv6 Relay Agent", [RFC 6221](#), May 2011.

- [RFC6607] Kinnear, K., Johnson, R., and M. Stapp, "Virtual Subnet Selection Options for DHCPv4 and DHCPv6", [RFC 6607](#), April 2012.
- [RFC6977] Boucadair, M. and X. Pougnaud, "Triggering DHCPv6 Reconfiguration from Relay Agents", [RFC 6977](#), July 2013.
- [RFC7159] Bray, T., "The JavaScript Object Notation (JSON) Data Interchange Format", [RFC 7159](#), March 2014.

Authors' Addresses

Ted Lemon
Nominum, Inc.
2000 Seaport Blvd
Redwood City, CA 94063
USA

Phone: +1-650-381-6000
Email: Ted.Lemon@nominum.com

Tomek Mrugalski
Internet Systems Consortium, Inc.
950 Charter Street
Redwood City, CA 94063
USA

Phone: +1 650 423 1345
Email: tomasz.mrugalski@gmail.com

