Implementing DHCPv6 on an IPv6 network

Benjamin Long
benlong@iol.unh.edu

8-11-2009

Table of Contents

DHCPv6 Overview

While Internet Protocol Version Six (IPv6) is capable of handling address auto-configuration, it is often an ineffective replacement for Dynamic Host Configuration Protocol (DHCP) in more complex networks for the distribution of other information. DHCP is used in current Internet Protocol Version Four (IPv4) networks to distribute addresses and other information. The most commonly distributed information is Domain Name Service (DNS) configurations. IPv6 auto-configuration only handles addressing configuration.

Dynamic Host Configuration Protocol Version Six (DHCPv6) was designed to handle the the automatic configuration of addressing, DNS information and is further capable of distributing other information to be covered further in this document. A minimum of two types of devices need to be configured on an IPv6 network to successfully implement DHCPv6: the server and the client devices. Additionally, depending on your network configuration and needs, you may need to configure a relay agent.

Terms used by DHCPv6

DHCPv6 Unique Identifier (DUID): All DHCPv6 devices identify themselves using a DUID. There are several methods for generating a DUID, but it is most often generated using the Media Access Control (MAC) address.

Global (address scope): Global addresses are used for communication outside the local network.

Lease: Addresses assigned by DHCPv6 are also assigned with a lease time as two values (T1 and T2). The lease time governs the duration the client device is able to use the address for and how often it should attempt to renew the lease. T1 defines when the client should begin attempting to renew the lease from the server that assigned the lease. After time T2 the the client attempts to find any available servers to renew its lease from.

Link (referring to a network): Functionally equivalent to a subnet in IPv4. Devices that are on link with each other can communicate without a router to route the traffic.

Link-Local (address scope): Link-local addresses are used only for communication with other devices on the same link. These addresses are never routed. They are also generated automatically by IPv6 devices when they are brought online. The link-local unicast prefix is fe80::/64 and the link-local multicast prefix is ff02::/64.

Multicast: Multicast addresses are destined for multiple nodes on a network, specifically the nodes that have joined the multicast group that corresponds to a multicast address. The All_DHCP_Relay_Agents_and_Servers address is ff02::1:2 (a link-local multicast address) and the All_DHCP_Servers address is ff05::1:3.

Prefix: IPv6 addresses are composed of a prefix and a link identifier which are equivalent to the network and host bits of an IPv4 address respectively. Prefixes are commonly recorded as <prefix>::/<prefix length>. For example: 3000::/64.

Session Initiation Protocol (SIP): SIP is used primarily for Voice over Internet Protocol (VoIP) to make calls to other clients.

Unicast: Unicast addresses are destined for a single node on a network.

Valid Lifetime: The maximum amount of time that an address can be used for. The lease time should always be set lower than the valid lifetime or client will not be able to communicate with other off link devices periodically.

DHCPv6 Message Exchange Process

The DHCPv6 message exchange process consists of four steps during initialization. DHCPv6 and IPv6 auto-configuration interact in beneficial way as well. When a client is first brought online, it performs IPv6 auto-configuration to generate its link-local address. The client then sends a Router Solicitation message to the all IPv6 routers multicast address. If DHCPv6 is in use on the network, the router then replies with a Router Advertisement message with the Managed Address Configuration Flag set.

At this point the client sends a DHCPv6 Solicit message to the All_DHCP_Relay_Agents_and_Servers address to probe for any available DHCPv6 servers. The servers that receive the Solicit message reply with an Advertise message containing configuration information available from that server. The client selects a server and transmits a Request message to the server, to confirm the configuration information. The server responds with a Reply message to the client confirming the information the client should bind to its interface.

Implementing the DHCPv6 Server

The DHCPv6 server is the most important and configuration intensive device on the network for DHCPv6 implementation. The server controls the configuration options that are available, as well the the lease time information.

Network Implementations and Requirements

For a network to implement DHCPv6 a network designer should be aware of a few requirements and restrictions. Because the DHCPv6 message exchange process occurs on a link-local scope, the server must have an interface located on the same link as the client devices, or a relay agent must be used to bridge between networks. For the purpose of this section, it is assumed that the clients and server are on the same link; for situations where the server is off-link, refer to section IV covering DHCPv6 relay agents.

For DHCPv6 to occur automatically when client devices are brought online, the router for the DHCPv6 network should be configured to set the Managed Address Configuration bit in outgoing Router Advertisements. If non-address information is also available over DHCPv6, set the Other Configuration Flag to indicate this to the clients.

DHCPv6 is capable of distributing additional information (such as DNS and SIP) while still using IPv6 address auto-configuration. To use this method, configure the router to set only the Other Config Flag (and not the Managed Config flag). When a client receives this option, it still configures its global address using the prefix provided in the Router Advertisement message, but the client also sends a DHCPv6 Information-Request message to all the DHCPv6 servers. Servers reply with a Reply message providing the configuration information from the server configuration file.

Server Configuration

To function effectively, the server will need to be configured with an address pool, timing information, and any other information that needs to be distributed across the network. Refer to the documentation available for your implementation for how to configure each option.

Address Pools

The server needs to be configured with addresses to be distributed to DHCPv6 clients. This is done through address pools which are used to define address ranges that can be distributed, as well as lease times (T1 and T2), as well as valid and preferred lifetimes for that address. You may also be able to use address pools to control what clients are allowed addresses by pool which can be useful for dividing your network into trusted and untrusted networks.

Access Control(optional)

You can set access control on your DHCPv6 server to only allow access to address pools to specifically listed. Access control is typically done using client DUID, but may vary from implementation to implementation. Common options include banning a specific DUID or range of DUIDs from an address pool as well as restricting access to only to a specific DUID or range of DUIDs.

Preference

Preference is a value between zero and 255 that clients use to determine which server to use in the case where they receive multiple DHCPv6 Advertise messages. A client will pick the server that has the highest preference value. If a client receives an Advertise message with a preference of 255 it will immediately continue the DHCPv6 message exchange process with the sending server without waiting the normal amount of time.

Rapid Commit (optional)

Rapid commit shortens the DHCPv6 message exchange to two messages instead of the normal four. Both the server and the client must be configured to allow rapid commit. Rapid commit is useful in networks under heavy load.

Unicast (optional)

A server can be configured to send a unicast option in its Advertise messages. Clients that support this option can then complete the message exchange using unicast instead of multicast. The unicast option is useful for reducing traffic in networks as multicast packets must be processed by multiple nodes.

DNS options (optional)

The DNS server option (also known as a DNS Recursive Name Server option) is used to inform clients of DNS server addresses. The address or addresses of the DNS server must be statically configured in the DHCPv6 server configuration. The domain option (also known as a domain search list option) is used to inform clients of the domain name the client is a member of. The domain option allows clients to refer to other devices with just their device name; without this option, clients must refer to other devices by their full domain name.

SIP domain options (optional)

DHCPv6 can be configured to distribute SIP information in the same manner as DNS information. SIP has server and domain information that can be provided.

Information only exchange (optional)

Servers can be configured to only provide non-address configuration information (such as DNS and SIP) information to clients. In this case the server will reply only to Information-Request messages.

Implementing the Client

Setting up the clients is relatively easy. There are several implementations available, and you typically won't have to do much configuration. Many systems that support IPv6 will come with DHCPv6 client software already installed.

Client Configuration

The client has very little configuration involved. You may have to configure which interfaces the client should run DHCPv6 on (usually the only interface the device has) as well as what type of address to request. If you want to use rapid commit or unicast DHCPv6, you should set the flags in the configuration accordingly. You can also configure the client to request a specific set of options (DNS and SIP for example). You can also force the client to use the information only message exchange process instead of the full DHCPv6 message exchange process. Forcing an information only exchange can be useful if the client uses a statically configured IPv6 address in a DHCPv6 managed network.

Relay Agent implementation

Relay agents are used in networks that have multiple links because the DHCPv6 message exchange occurs using link-local multicast addresses. To have a single server cover multiple links, relay agents are placed in each link. The relay agents accept the clients messages and forward them to the servers by encapsulating them in a Relay-Forward messages sent to the DHCPv6 server. The server then replies to the relay agent with a Relay-Reply message containing the reply to the encapsulated packet in the Relay-Forward. Both the relay agent and server must be specifically configured to handle this network configuration.

In some very large networks you may need to implement cascading relay agents to cover the entire network with a single server. A cascade relay agent accepts a Relay-Forward or Relay-Reply messages, encapsulates them in new Relay-Forward messages or decapsulating them into smaller Relay-Reply messages respectively before forwarding them to the next node. The client side Relay-Agent must be configured to decapsulate the packet and forward the servers response.

Relay Agent Server Configuration

The server must be configured to accept and decapsulate incoming Relay-Forward messages and to encapsulate outgoing messages in Relay-Reply messages. The specific configuration requirements for performing this task vary from implementation to implementation, refer to your implementation's documentation for details on configuring the server for this role. You must also configure the server to distribute information as covered under the server section of this document.

Relay Agent Configuration

The relay agent must be configured to know which interface it is receiving client requests from and which interface it is sending Relay-Forward messages out. These interfaces may be the same interface but this configuration only makes sense if you instruct the relay agent to also use the an address larger than link-local to get to the DHCPv6 server (such as the server's global address). The relay agent may also need be configured with an interface identifier to uniquely identify the receiving interface to the server. Relay agents may be configured to send and receive messages on specific unicast addresses which is primarily used for cascading relays. Relay agent's listening on client networks should be listening for client multicast messages.

Cascading Relay Agent Configuration

When using cascading relay agents you must still configure the same items as a normal relay agent, however there are additional steps that you must take. To minimize traffic load you should consider configuring the relay agents to transmit traffic to each other using unicast addresses where available. You have to ensure that each cascade node is configured to accept incoming Relay-Forward messages and forward them toward the server.

References

RFC 3315, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)"
RFC 3646, "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)"
RFC 3736, "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6"