Introduction to Dynamic Host Configuration

All hosts seeking Internet access must have an IP address. As Internet administrator, you must perform the following for every new user and for every user whose computer was moved to another subnet:

1. Choose a legal IP address.
2. Assign the address to the individual device.
3. Define device configuration parameters.
4. Update the DNS database, mapping the device name to the IP address.

These activities are time consuming and error prone, hence the Dynamic Host Configuration Protocol (DHCP). DHCP frees you from the burden of individually assigning IP addresses. It was designed by the Internet Engineering Task Force (IETF) to reduce the amount of configuration required when using TCP/IP. DHCP allocates IP addresses to hosts. It also provides all the parameters that hosts require to operate and exchange information on the Internet network to which they are attached.

DHCP localizes TCP/IP configuration information. It also manages allocating TCP/IP configuration data by automatically assigning IP addresses to systems configured to use DHCP. Thus, you can ensure that hosts have Internet access without having to configure each host individually.

This chapter contains the following sections:

How DHCP Works

DHCP makes dynamic address allocation possible by shifting device configuration to global address pools at the server level. DHCP is based on a client/server model. The client software runs on the device and the server software runs on the DHCP server.

Sample DHCP User

After Beth's workstation (bethpc) is configured with DHCP, these actions occur when she first starts up:

1. Her PC automatically requests an IP address from a DHCP server on the network.

2. The DHCP server offers her a lease that is an IP address, with assigned lease time, and with other configuration data necessary to use the Internet. Nobody else uses the leased address and it is valid only for her PC.

3. Before the address lease expires, bethpc renews it by requesting a lease extension from the server that provided the lease (this process usually begins about half way through the original assigned lease time), thereby extending the expiration time. If unable to renew the lease by about 85% of the lease time, bethpc will initiate sending a slightly different request attempting to renew the lease from any available server. Bethpc continues to use the lease right up to its expiration if it cannot reach the server.

    To summarize, there are three important client times:

    • **Lease Expiration Time** (**Valid Lifetime**)—The time at which the lease expires. This is always explicitly communicated to the client.

    • **Renewal Time (T1)**—The time at which the client can start the renew process with the server that granted or last extended the lease. The renewals for DHCPv4 are unicast. For DHCPv6, the client specifies the server from which the lease was granted or last renewed.

    Renewal Time (T1) is either explicitly communicated by the server or left for the client to generate. It is by default at 50% of the lease time.

    • **Rebinding Time (T2)**—The time at which the client can start the rebind process, which is similar to the renewal process but no longer restricted to a single server. For DHCPv4, these requests are broadcast (hence picked up by relay and forwarded to both failover partners). For DHCPv6, the client does not specify a server and hence any server can respond.

    Rebinding Time (T2) is either explicitly communicated by the server or left for the client to generate. It is usually at about 87.5% (for DHCPv4) and 85% (for DHCPv6) of the lease time.

4. If Beth relocates to another department and her PC moves to a different subnet, her current address expires and becomes available for others. When Beth starts her PC at its new location, it leases an address from an appropriate DHCP server on the subnet (see the image below).

As long as the DHCP server has the correct configuration data, none of the workstations or servers using DHCP will ever be configured incorrectly. Therefore, there is less chance of incurring network problems from incorrectly configured devices and servers that are difficult to trace.

Figure 1: Hosts Request an IP Address

The example shows the DHCP protocol with a set of DHCP servers that provide addresses on different subnets. To further simplify the administration of address pools, network routers are often configured as DHCP relay agents to forward client messages to a central DHCP server. This server is configured with address pools for a group of subnets.

Typical DHCP Administration

To use DHCP, you must have at least one DHCP server on the network. After you install the server:

- Define a scope of IP addresses that the DHCP server can offer to DHCP clients. You no longer need to keep track of which addresses are in use and which are available.
- Configure a secondary server to share the distribution or handle leases if the first DHCP server goes down. This is known as DHCP failover. For information on Managing DHCP Failover, see Managing DHCP Failover.
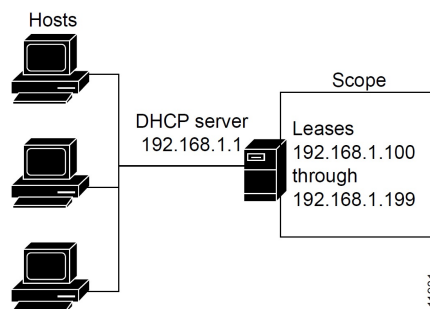
Leases

One of the most significant benefits of DHCP is that it can dynamically configure devices with IP addresses and associate leases with the assigned addresses. DHCP uses a lease mechanism that offers an automated, reliable, and safe method for distributing and reusing addresses in networks, with little need for administrative intervention. As system administrator, you can tailor the lease policy to meet the specific needs of your network.

Leases are grouped together in an address pool, called a scope, which defines the set of IP addresses available for requesting hosts. A lease can be reserved (the host always receives the same IP address) or dynamic (the host receives the next available, unassigned lease in the scope). The DHCP server of the site is configured to lease addresses 192.168.1.100 through 192.168.1.199 (see the image below).

If you plan not to have more network devices than configured addresses for the scope, you can define long lease times, such as one to two weeks, to reduce network traffic and DHCP server load.

Figure 2: DHCP Hosts Requesting Leases from a DHCP Server

Scopes and Policies

A scope contains a set of addresses for a subnet, along with the necessary configuration parameters. You must define at least one scope for each subnet for which you want dynamic addressing.

A policy includes lease times and other configuration parameters that a DHCP server communicates to clients. Use policies to configure DHCP options that the DHCP server supplies to a client upon request. Policies

ensure that the DHCP server supplies all the correct options for scopes without having to do so separately for each scope (see the image below).

The difference between scopes and policies is that scopes contain server information about addresses, such as which address is leasable and whether to ping clients before offering a lease. Policies contain client configuration data, such as the lease duration and address of the local DNS server.

Policies are especially useful if you have multiple scopes on a server. You can create policies that apply to all or selected scopes. The Cisco Prime Network Registrar policy hierarchy is a way to define policies from least to most specific. For example, you usually specify a router option for each policy, which means that you would need a policy for each scope. Scope-specific policies like this can be defined in a scope-embedded policy. More general policies, such as those referring to lease times, can be applied in a system-wide policy (see Configuring DHCP Policies). You can also write extensions to handle policy assignments (see Using Extensions to Affect DHCP Server Behavior).

Figure 3: Scopes and Policies

Links and Prefixes

The explicit DHCPv6 configuration objects are links and prefixes:

- **Link**—Network segment that can have one or more prefixes, and adds an additional layer at which policies can be applied for DHCPv6 clients.
- **Prefix**—Equates to a scope in IPv4. The link associated with a prefix is similar to a primary scope, except that it names a link and not another prefix.

Just as with scopes, you can create multiple prefix objects for the same IPv6 prefix. However, rather than supporting multiple ranges with explicit start and end addresses, prefixes support only a single range that must be an IPv6 prefix with a length the same as, or longer than, the prefix object. For example, if you define a 2001::/64 prefix with a 2001::/96 range, the server can assign addresses from 2001:0:0:0:0:0:0:0 through 2001:0:0:0:0:0:ffff:ffff only. The range:

- Is limited to powers of 2.
- Must be unique (cannot be duplicated by any other range, except in a different VPN).
- Cannot be contained in, or contain, another range, except for prefix delegation prefixes, as explained below.
- Is the full IPv6 prefix if not specified, except for prefix delegation prefixes, as explained below.

If a prefix delegation prefix object is defined with an unspecified range, it may contain non prefix-delegation prefixes, and the effective range is either:

- The full IPv6 prefix if no other prefixes exist with the same IPv6 prefix, or
- The prefixes that remain when all other ranges for prefix objects with the same IPv6 prefix are removed from the IPv6 prefix.

You create a link only if more than one prefix object with a different IPv6 prefix exists on a link. When the server loads the configuration, if a prefix has no explicit link, the server searches for or creates an implicit link with the name Link-[*vpn.name*/]*prefix*. All prefix objects with the same IPv6 prefix must either not specify a link or explicitly specify the same link.

The DHCPv6-enabled server supports VPN address spaces for DHCPv6. Both the link and prefix objects may be assigned to a VPN. But all prefixes on a link must use the same VPN ID. Because there is presently no DHCPv6 VPN option, clients can only be assigned addresses from a VPN by using the client or client-class *override-vpn* attribute.

Related Topics

Determining Links and Prefixes

Generating Addresses

Generating Delegated Prefixes

Prefix Stability

Cisco Prime Network Registrar DHCP Implementations

The Cisco Prime Network Registrar DHCP server provides a reliable method for automatically assigning IP addresses to hosts on your network. You can define DHCP client configurations, and use the Cisco Prime Network Registrar database to manage assigning client IP addresses and other optional TCP/IP and system configuration parameters. The TCP/IP assignable parameters include:

- IP addresses for each network adapter card in a host.
- Subnet masks for the part of an IP address that is the physical (subnet) network identifier.
- Default gateway (router) that connects the subnet to other network segments.
- Additional configuration parameters you can assign to DHCP clients, such as a domain name.

Cisco Prime Network Registrar automatically creates the databases when you install the DHCP server software. You add data through the web UI or CLI as you define DHCP scopes and policies.

The Cisco Prime Network Registrar DHCP server also supports allocating addresses in virtual private networks (VPNs) and subnets to pool manager devices for on-demand address pools. These features are described in the following sections.

Related Topics

Virtual Private Networks, on page 5

Subnet Allocation and DHCP Address Blocks

Virtual Private Networks

Virtual private networks (VPNs) allow the possibility that two pools in separate networks can have the same address space, with these two pools having overlapping private network addresses. This can save address resources without having to use valuable public addresses. These VPN addresses, however, require a special designator to distinguish them from other overlapping IP addresses. Cisco Prime Network Registrar DHCP

servers that are not on the same VPN as their clients can now allocate leases and addresses to these clients, and can distinguish the addresses from one VPN to another.

Through changes made to the Cisco Prime Network Registrar DHCP server and Cisco IOS DHCP Relay Agent, the DHCP server can service clients on multiple VPNs. A VPN distinguishes a set of DHCP server objects, making them independent of otherwise identical objects in other address spaces. You can define multiple VPNs containing the same addresses. You create a VPN based on the VPN identifier configured in the Cisco IOS Relay Agent.

The illustration below shows a typical VPN-aware DHCP environment. The DHCP Relay Agent services two distinct VPNs, blue and red, with overlapping address spaces. The Relay Agent has the interface address 192.168.1.1 on VPN blue and is known to DHCP Server 1 as 172.27.180.232. The server, which services address requests from DHCP Client 1 in VPN blue, can be on a different network or network segment than the client, and can be in a failover configuration with DHCP Server 2 (see Managing DHCP Failover). The Relay Agent can identify the special, distinguished route of the client address request to the DHCP server, as coordinated between the Relay Agent and Cisco Prime Network Registrar administrators (see RFC 6607). The DHCP servers can now issue leases based on overlapping IP addresses to the clients on both VPNs.

*Figure 4: Virtual Private Network DHCP Configuration*
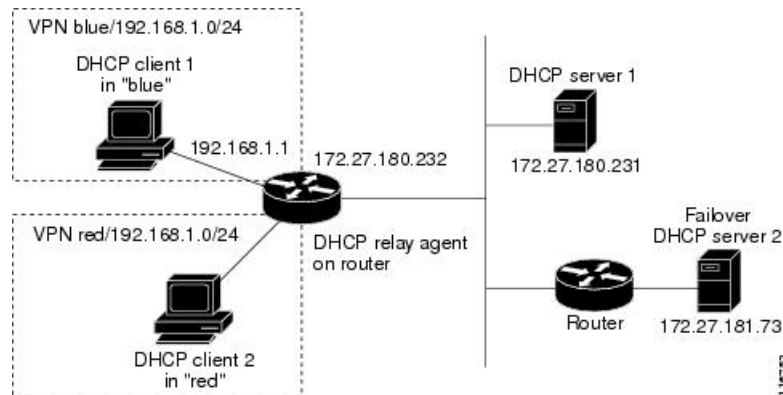
Prefix Delegation

Prefix delegation enables delegation of prefixes from a DHCPv6 server to a requesting device. Prefix Delegation is used by service providers to assign a prefix to a Customer Premise Equipment (CPE) device. It is also used by an ISP to delegate a prefix to a subscriber.

During operation, a DHCPv6 server is provided IPv6 prefixes to be delegated to the requesting device. The requesting device requests prefix(es) from the DHCPv6 server. The DHCPv6 server chooses prefix(es) for delegation, and responds with prefix(es) to the requesting device. The requesting device is then responsible for the delegated prefix(es). For example, the requesting device might assign a subnet from a delegated prefix to one of its interfaces, and begin sending advertisements for the prefix on that link. Each prefix has an associated valid and preferred lifetime, which constitutes an agreement about the length of time over which the requesting device is allowed to use the prefix. A requesting device can request an extension of the lifetimes on a delegated prefix and is required to terminate the use of a delegated prefix if the valid lifetime of the prefix expires.

Figure 5: Model Topology for the end-user network
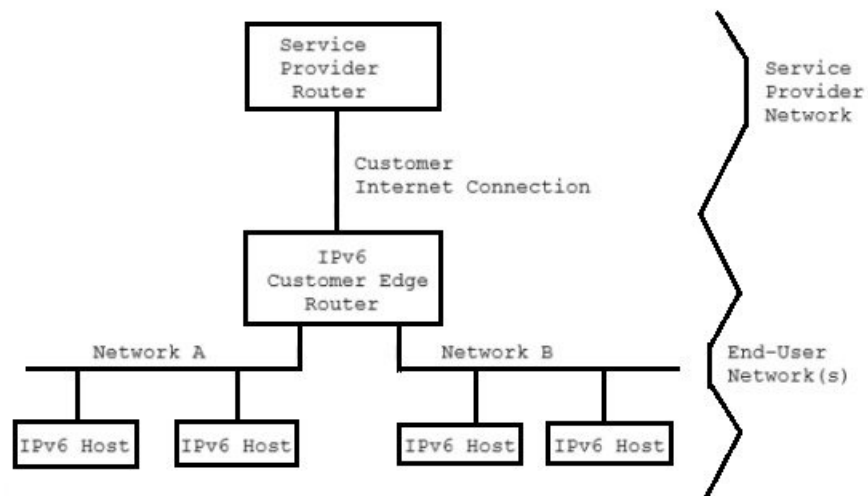
DNS Update

Although DHCP frees you from the burden of distributing IP addresses, it still requires updating the DNS server with DHCP client names and addresses. DNS update automates the task of keeping the names and addresses current. With the Cisco Prime Network Registrar DNS update feature, the DHCP server can tell the corresponding DNS server when a name-to-address association occurs or changes. When a client gets a lease, Cisco Prime Network Registrar tells the DNS server to add the host data. When the lease expires or when the host gives it up, Cisco Prime Network Registrar tells the DNS server to remove the association.

In normal operation, you do not have to manually reconfigure DNS, no matter how frequently clients' addresses change through DHCP. Cisco Prime Network Registrar uses the hostname that the client device provides. You also can have Cisco Prime Network Registrar synthesize names for clients who do not provide them, or use the client lookup feature to use a preconfigured hostname for the client.

Different use-cases for DHCPv4 and DHCPv6 DNS update made server design different to handle hostname updates. So, the difference of behavior in DHCPv4 and DHCPv6 DNS updates for hostname is expected.

Effect on DNS of Obtaining Leases

For ExampleCo, the administrator creates a scope on the DHCP server and allocates 100 leases (192.168.1.100 through 192.168.1.199). Each device gets its owner name. The administrator also configures the DHCP server to use DNS update and associates it with the correspondingly configured DNS server. The administrator does not need to enter the names in the DNS server database.

Monday morning, Beth (user of bethpc) tries to log in to a website without having an address. When her host starts up, it broadcasts an address request (see the image below).

Figure 6: DNS Update at ExampleCo Company

The DHCP server then:

1. Gives bethpc the next available (unassigned) IP address (192.168.1.125).
2. Updates her DNS server with the hostname and address (bethpc 192.168.1.125).

Beth can now access the website. In addition, programs that need to translate the name of Beth's machine to her IP address, or the other way around, can query the DNS server.

Effect on DNS of Reacquiring Leases

When Beth returns from her trip to start up her host again:

1. Her PC broadcasts for an IP address.
2. The DHCP server checks if the host is on the correct network. If so, the server issues an address. If not, the server on the correct network issues the address.
3. The DHCP server updates the DNS server again with the host and address data.

Effect on DNS of Releasing Leases

Later that day, Beth learns that she needs to travel out of town. She turns off her host, which still has a leased address that is supposed to expire after three days. When the lease is released, the DHCP server:

1. Acknowledges that the IP address is now available for other users (see the figure below).
2. Updates the DNS server by removing the hostname and address. The DNS server no longer stores data about bethpc or its address.

Figure 7: Relinquishing a Lease

DHCP Failover

Cisco Prime Network Registrar failover protocol is designed to allow a backup DHCP server to take over for a main server if the main server is taken offline for any reason. Prior to 8.2, this protocol was UDP based, only operated over IPv4, and only supported DHCPv4. Starting with 8.2, this protocol is TCP based, can be configured to use either IPv4 or IPv6, and supports both DHCPv4 and DHCPv6 over a single connection. The DHCP server will try both IPv4 and IPv6 transports if configured to use both, and will use whichever connection comes up first. The existing DHCP clients can keep and renew their leases without the need to know which server is responding to their requests.

You can create and synchronize failover pairs at the local and regional clusters in Cisco Prime Network Registrar. For details, see Managing DHCP Failover.

Allocating Addresses Through Failover

In order to keep the failover pair operating in spite of a network partition, in which both can communicate with clients but not with each other, you must make available more addresses than the addresses needed to run a single server. Configure the main server to allocate a percentage of the currently available (unassigned) addresses in each scope or prefix delegation address pool to its partner. These addresses become unavailable to the main server. The partner uses them when it cannot talk to the main server and does not know if it is down. However, when the failover partners are in communication, they periodically rebalance these pools.

The backup server needs enough addresses from each scope or prefix to satisfy the requests of all new DHCP clients that arrive during the period in which the backup does not know if the main server is down. The default backup percentage for a failover pair is 50%. This ensures that during the failover the other partner has equal number of addresses.

Even during PARTNER-DOWN state, the backup server waits for the lease expiration and the maximum client lead time (MCLT), a small additional time buffer, before reallocating any leases. When these times expire, the backup server offers:

- Leases from its private pool of addresses.
- Leases from the main server pool of addresses.
- Expired leases to new clients.

During the working hours, if the administrative staff can respond within two hours to a COMMUNICATIONS INTERRUPTED state to determine if the main server is working, the backup server needs enough addresses to support a reasonable upper bound on the number of new DHCP clients that might arrive during those two hours.

During off-hours, if the administrative staff can respond within 12 hours to the same situation, and considering that the arrival rate of previously unheard from DHCP clients is also less, the backup server then needs enough addresses to support a reasonable upper bound on the number of DHCP clients that might arrive during those 12 hours.

Consequently, the number of addresses over which the backup server requires sole control would be the greater of the numbers of addresses given out during peak and non-peak times, expressed as a percentage of the currently available (unassigned) addresses in each scope or prefix.

Note The default use-safe-period is enabled for the DHCP failover pair and the default safe period is 4 hours. This ensures that if the failover partner is in COMMUNICATIONS-INTERRUPTED state for 4 hours, it will enter PARTNER-DOWN state automatically after the safe period elapses.

Client-Classes

You can use the Cisco Prime Network Registrar client and client-class facility to provide differentiated services to users that are connected to a common network. You can group your user community based on administrative criteria, and then ensure that each user receives the appropriate class of service.

Although you can use the Cisco Prime Network Registrar client-class facility to control any configuration parameter, the most common uses are for:

- **Lease periods**—How long a set of clients should keep their addresses.
- **IP address ranges**—From which lease pool to assign clients addresses.
- **DNS server addresses**—Where clients should direct their DNS queries.
- **DNS hostnames**—What name to assign clients.
- **Denial of service**—Whether unauthorized clients should be offered leases.

One way to use the client-class facility is to allow visitors access to some, but not all, of your network. For example, when Joe, a visitor to ExampleCo, tries to attach his laptop to the example.com network, Cisco Prime Network Registrar recognizes the laptop as being foreign. ExampleCo creates one class of clients known as having access to the entire network, and creates another visitor class with access to a subnet only. If Joe needs more than the standard visitor access, he can register his laptop with the Cisco Prime Network Registrar system administrator, who adds him to a different class with the appropriate service.

The following sections describe how DHCP normally processes an address assignment, and then how it would handle it with the client-class facility in effect.

DHCP Processing Without Client-Classes

To understand how you can apply client-class processing, it is helpful to know how the DHCP server handles client requests. The server can perform three tasks:

- Assign an IP address.
- Assign the appropriate DHCP options (configuration parameters).

• Optionally assign a fully qualified domain name (FQDN) and update the DNS server with that name.

The DHCP server:

1. Assigns an address to the client from a defined scope—To choose an address for the client, the DHCP server determines the client subnet, based on the request packet contents, and finds an appropriate scope for that subnet.

   If you have multiple scopes on one subnet or several network segments, which is known as multinetting, the DHCP server may choose among these scopes in a round-robin fashion, or you can change the priority of the scope choice by using the DHCP server address allocation priority feature (see Configuring Multiple Scopes Using Allocation Priority). After the server chooses a scope, it chooses an available (unassigned) address from that scope:

   a. It assigns DHCP option values from a defined policy. Cisco Prime Network Registrar uses policies to group options. There are two types of policies: scope-specific and system default. For each DHCP option the client requests, the DHCP server searches for its value in a defined sequence.

   b. If the scope-specific policy contains the option, the server returns its value to the client and stops searching.

   c. If not found, the server looks in the system default policy, returns its value, and stops searching.

   d. If neither policy contains the option, the server returns no value to the client and logs an error.

   e. The server repeats this process for each requested option.

2. With DNS update in effect, the server assigns an FQDN to the client. If you enabled DNS update, Cisco Prime Network Registrar enters the client name and address in the DNS host table. See DNS Update, on page 7. The client name can be:

   • Its name as specified in the client lease request (the default value).
   • Its MAC address (hardware address; for example, 00:d0:ba:d3:bd:3b).
   • A unique name using the default prefix *dhcp* or a specified prefix.

DHCP Processing with Client-Classes

When you enable the client-class facility for your DHCP server, the request processing performs the same three tasks of assigning IP addresses, options, and domain names as described in DHCP Processing Without Client-Classes, on page 10, but with added capability. The DHCP server:

1. **Considers the client properties and client-class inclusion before assigning an address**—As in regular DHCP processing, the DHCP server determines the client subnet. The server then checks if there is a client-class defined or a MAC address for this client in its database. If there is:

   a. A client-class defined by a client-class lookup ID expression, the client is made a member of this client-class.

   b. No MAC address, it uses the default client. For example, the default client could have its client-class name set to Guest, and that client-class could limit (using options and address selection) what network operations such clients are permitted.

   c. No MAC address and no default client, the server handles the client through regular DHCP processing.

   d. No client-specifier, but a MAC address, the MAC address is converted into a client-specifier. An unknown client is mapped to the default client, if the default client is defined.

   The scopes must have addresses on client-accessible subnets. That is, they must have a selection tag that associates them with a client-class. To assign the same clients to different address pools, you must use separate scopes.

For example, a scope would either have a selection tag of Employee or Guest, but not both. In this case, there are two scopes for each subnet; one with the selection tag Employee, and the other with Guest. Each scope has a different associated policy and address range that provides the appropriate access rights for the user group.

2. **Checks for client-class DHCP options**—In regular DHCP processing, the server checks the scope-specific and system default DHCP options. With client-class, it also first checks the client-specific and client-class-specific options.

3. **Provides additional FQDN assignment options**—Beyond the usual name assignment process of using the hostname the client requests, the server can:

- Provide an explicit hostname that overrides it.
- Drop the client-requested hostname and not replace it.
- Synthesize a hostname from the client MAC address.

Defining Scopes for Client-Classes

The motivating factor for using client-classes is often to offer an address from one or another address pool to a client. Another motivating factor might be to provide clients with different option values or lease times. Offering clients addresses from separate pools requires defining more than one scope.

To get more than one scope on a subnet, they must come from the same network segment. Networks are not configured directly in Cisco Prime Network Registrar, but are inferred from scope configurations. Scopes become related (end up in the same network):

- **Implicitly**—Two scopes have the same network number and subnet mask. These scopes naturally end up on the same network without explicit configuration.
- **Explicitly**—One scope is marked as a secondary to another. This is required when the scope marked as a secondary has a network and subnet mask unrelated to the primary. An example is putting a set of 10.0.0.0 network addresses on a normal, routable network segment.

When the Cisco Prime Network Registrar DHCP server reads the scope configuration from its database, it places every scope in a network, and logs this information. Scopes with the same network number and subnet mask end up on the same network, while a secondary scope ends up on the primary scope network.

Choosing Networks and Scopes

When a DHCP packet arrives, the server determines the address from which it came by:

- When a DHCPv4 packet arrives the server determines the gateway address (*giaddr*), if there was one, for packets sent through a BOOTP relay.
- For information on DHCPv6, see Determining Links and Prefixes.
- Interface address of the interface on which the broadcast packet arrived, if the DHCP client is on a network segment to which the DHCP server is also directly connected.

In all cases, the DHCP server determines a network from the gateway or interface address. Then, if the network has multiple scopes, the server determines from which scope to allocate an address to the DHCP client. It always looks for a scope that can allocate addresses to this type of client. For example, a DHCP client needs a scope that supports DHCP, and a BOOTP client needs one that supports BOOTP. If the client is a DHCP client and there are multiple scopes that support DHCP, each with available (unassigned) addresses, the DHCP server allocates an IP address from any of those scopes, in a round-robin manner, or by allocation priority.

Selection tags and client-classes let you configure the DHCP server to allocate IP addresses from:

- One or more scopes on a network to one class of clients.
- A different set of scopes to a different class of clients.

In the latter case, the gateway or interface address determines the network. The client-class capability, through the mechanism of the selection tags, determines the scope on the network to use.