



Network Registrar Components

Cisco CNS Network Registrar provides the tools to configure and control the servers necessary to manage your IP address space. This chapter provides an overview of the related management and network concepts and protocols.

Management Components

Network Registrar contains two management components:

- Regional component—Includes a Web-based user interface (Web UI) and Central Configuration Management (CCM) server to provide to local cluster, address space, and router management.
- Local component—Includes a Web UI, command line interface (CLI), and CCM server, and manages the Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and Trivial File Transport Protocol (TFTP) servers, local address space, zones, scopes, and user administration.

The CCM server, Web UIs, and CLI are described in [Chapter 3, “Network Registrar User Interfaces.”](#) The remainder of this chapter describes the network protocols—DNS, DHCP, TFTP, and the Simple Network Management Protocol (SNMP).

Domain Name System and Zone Administration

The Domain Name System (DNS) was designed to handle the growing number of Internet users. DNS translates names, such as `www.cisco.com`, into Internet Protocol (IP) addresses, such as `192.168.40.0`, so that computers can communicate with each other. DNS makes using Internet applications, such as the World Wide Web, easy. The process is as if, when phoning your friends and relatives, you could autodial them based on their names instead of having to remember their phone numbers.


How DNS Works

To understand how DNS works, imagine a typical user, John, logging on to his computer. He launches his Web browser so that he can view the website at a company, ExampleCo (see [Figure 2-1](#)). He enters the name of their website—`http://www.example.com`. Then:

1. John’s workstation sends a request to the DNS server about the IP address of `www.example.com`.
2. The DNS server checks its database to find that `www.example.com` corresponds to `192.168.1.4`.

3. The server returns this address to John's browser.
4. The browser uses the address to locate the website.
5. The browser displays the website on John's monitor.


Figure 2-1 Domain Names and Addresses



Domains

John can access ExampleCo's website because his DNS server knows the `www.example.com` IP address. The server learned the address by searching through the domain namespace. DNS was designed as a tree structure, where each named domain is a node in the tree. The top-most node of the tree is the DNS root domain (`.`), under which there are subdomains, such as `.com`, `.edu`, `.gov`, and `.mil` (see [Figure 2-2](#)).

Figure 2-2 The Domain Name System Hierarchy



The fully qualified domain name (FQDN) is a dot-separated string of all the network domains leading back to the root. This name is unique for each host on the Internet. The FQDN for the sample domain is `example.com.`, with its domain `example`, parent domain `.com`, and root domain `"."` (`dot`).


Learning ExampleCo's Address

When John's workstation requests the IP address of the website `www.example.com` (see [Figure 2-3](#)):

1. The local DNS server looks for the `www.example.com` domain in its database, but cannot find it, indicating that the server is not authoritative for this domain.

2. The server asks the root nameserver that is authoritative for the top-level (root) domain “.” (dot).
3. The root nameserver directs the query to a nameserver for the .com domain that knows about its subdomains.
4. The .com nameserver responds that example.com is one of its subdomains and responds with its server’s address.
5. The local server asks the example.com nameserver for www.example.com’s location.
6. The example.com nameserver replies that its address is 192.168.1.4.
7. The local server sends this address to John’s Web browser.

Figure 2-3 DNS Hierarchical Name Search



Establishing a Domain

ExampleCo has a website that John could reach because it registered its domain with an accredited domain registry. ExampleCo also entered its domain name in the .com server’s database, and requested a network number, which defines a range of IP addresses. In this case, the network number is 192.168.1.0, which includes all addresses in the range 192.168.1.1 through 192.168.1.255. You can only have the numbers 0 through 256 (2^8) in each of the address fields, known as octets. However, the numbers 0 and 256 are reserved for network and broadcast addresses, respectively, and are not used for hosts.

Difference Between Domains and Zones


The domain namespace is divided into areas called zones that are points of delegation in the DNS tree. A zone contains all domains from a certain point downward, except those for which other zones are authoritative.

A zone usually has an authoritative nameserver, often more than one. In an organization, you can have many nameservers, but Internet clients can query only those that the root nameservers know. The other nameservers answer internal queries only.

The ExampleCo company registered its domain, example.com. It established three zones—example.com, marketing.example.com, and finance.example.com. ExampleCo delegated authority for marketing.example.com and finance.example.com to the DNS servers in the Marketing and Finance groups in the company. If someone queries example.com about hosts in marketing.example.com, example.com directs the query to the marketing.example.com nameserver.

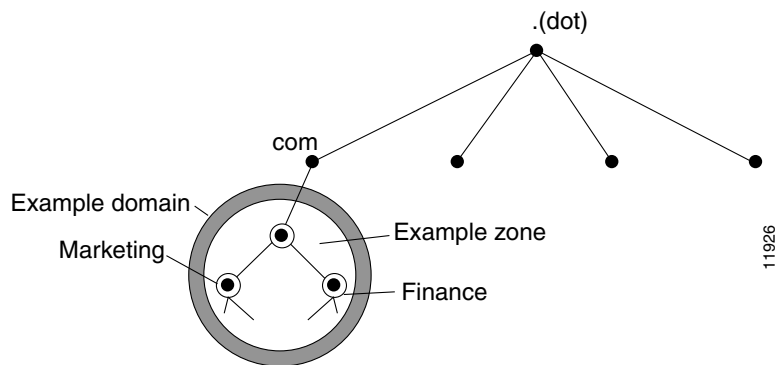
In [Figure 2-4](#), the domain `example.com` includes three zones, with the `example.com` zone being authoritative only for itself.

Figure 2-4 Example.com With Delegated Subdomains

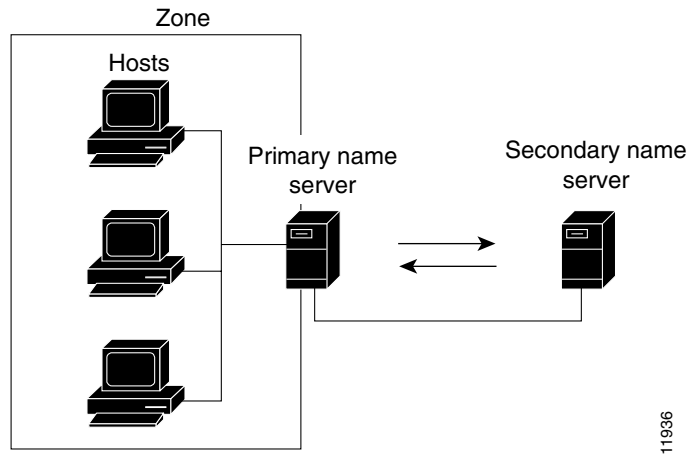


ExampleCo could choose not to delegate authority to its subdomains. In that situation, the `example.com` domain is a zone that is authoritative for the subdomains for marketing and finance (see [Figure 2-5](#)). The `example.com` server answers all outside queries about marketing and finance.

Figure 2-5 Example.com Without Delegation




As you begin to configure zones using Network Registrar, you must configure a nameserver for each zone. Each zone has one primary server, which loads the zone's contents from a local configuration database. Each zone can also have any number of secondary servers, which load the zone contents by fetching the data from the primary server. [Figure 2-6](#) shows a configuration with one secondary server.

Figure 2-6 Primary and Secondary Servers for Zones

Nameservers


DNS is based on a client/server model. In this model, nameservers store data about a portion of the DNS database and provide it to clients that query the nameserver across the network. Nameservers are programs that run on a physical host and store zone data. As administrator for a domain, you set up a nameserver with the database of all the resource records describing the hosts in your zone or zones (see [Figure 2-7](#)). For details about DNS resource records, see [Appendix A, “Resource Records.”](#)

Figure 2-7 Client/Server Name Resolution

The DNS servers provide name-to-address translation, or name resolution. They interpret the information in a fully qualified domain name (FQDN) to find its address. If a local nameserver does not have the data requested in a query, it asks other nameservers until it finds it. For commonly requested names, this process can go quickly, because nameservers continuously cache the information they learn from queries about the domain namespace.

Each zone must have one primary nameserver that loads the zone contents from a local database, and a number of secondary servers, which load a copy of the data from the primary server (see [Figure 2-8](#)). This process of updating the secondary server from the primary server is called a zone transfer.

Figure 2-8 DNS Zone Transfer



Even though a secondary nameserver acts as a kind of backup to a primary server, both types of servers can be authoritative for the zone. They both learn about host names in the zone from the zone's authoritative database, not from information learned while answering queries. Clients can query both servers for name resolution.

As you configure Network Registrar's DNS nameserver, you specify what role you want the server to perform for a zone—primary, secondary, or caching-only. The type of server is meaningful only in context to its role. A server can be a primary for some zones and a secondary for others. It can be a primary or secondary only, or it can serve no zones and just answer queries by means of its cache.

Although all servers are caching servers, because they save the information until it expires, a caching-only server is one that is not authoritative for any zone. This server answers internal queries and asks other authoritative servers for the information. Sites create caching-only servers to unburden the authoritative servers so that they do not need to have every query directed to the authoritative servers.

To configure the:

- Primary nameserver, see the [“Managing Primary DNS Servers”](#) section on page 8-3.
- Secondary server, see the [“Managing Secondary Servers”](#) section on page 8-10.
- Caching-only server, see the [“Configuring Caching-Only Servers”](#) section on page 10-2.

Reverse Nameservers


The DNS servers described so far perform name-to-address resolution. They can do this easily by searching through their database for the correct address, because they index all the data by name. However, there are times when you need address-to-name resolution so that you can interpret certain output, such as computer log files.

Finding a domain name when you only know the address, however, would require searching the entire namespace. DNS solves this problem by supporting a domain namespace that uses addresses as names, known as the `in-addr.arpa` domain. This reverse zone contains subdomains for each network based on the network number. For consistency and natural grouping, the four octets of a host number are reversed.

When you read the IP address as a domain name, it appears backwards, because the name is in leaf-to-root order. For example, ExampleCo's example domain's network number is 192.168.1.0. Its reverse zone is 1.168.192.in-addr.arpa. If you only know the DNS server address (192.168.1.1), the query to the reverse domain would find the host entry 1.1.168.192.in-addr.arpa that maps back to example.com (see [Figure 2-9](#)).

Reverse domains are handled through Pointer (PTR) resource records, as indicated in [Figure 2-9](#).

Figure 2-9 Reverse Domains



Dynamic Host Configuration and Leases

All hosts seeking Internet access must have an IP address. As Internet administrator, you must do 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 workstation.
3. Define workstation configuration parameters.
4. Update the DNS database, mapping the workstation 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.

How DHCP Works

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

Sample DHCP User


After Beth's workstation (bethpc) is configured to use DHCP, these actions occur when she first starts her workstation (see [Figure 2-10](#)):

1. Her workstation 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 the configuration data necessary to use the Internet. Nobody else uses the leased address, and it is valid only for her workstation.
3. Before the address's lease expires, bethpc renews it, thereby extending the expiration time. It continues to use the lease right up to its expiration or if it cannot reach the server.
4. If Beth relocates to another department and her workstation moves to a different subnet, her current address expires and becomes available for others. When Beth starts her workstation at its new location, it leases an address from an appropriate DHCP server on the subnet.

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 workstations and servers that are difficult to trace.

Figure 2-10 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:

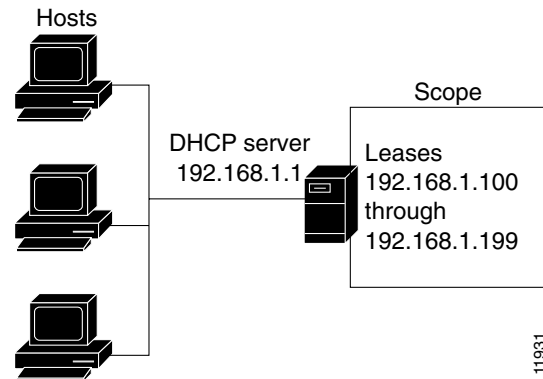
- Define a scope of IP addresses that the DHCP server can offer to DHCP clients. You no longer need to keep track of which addresses are in use and which are available.
- Configure a secondary server to share the distribution or handle leases if the first DHCP server goes down. This is known as DHCP failover, and is described further in the [“DHCP Failover” section on page 2-14](#).

Leases

One of the most significant benefits of DHCP is that it can dynamically configure workstations 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 ExampleCo DHCP server is configured to lease addresses 192.168.1.100 through 192.168.1.199 (see [Figure 2-11](#)).

Figure 2-11 DHCP Hosts Requesting Leases from a DHCP Server



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.

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 [Figure 2-12](#)).

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 Network Registrar policy hierarchy lets you 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 the [“Configuring DHCP Policies”](#) section on page 11-16). You can also write extensions to handle policy assignments (see the [“Using Extensions to Affect DHCP Server Behavior”](#) section on page 14-12).

Figure 2-12 Scopes and Policies



Network Registrar's DHCP Implementations

The 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 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.

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 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.

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. 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 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.

Figure 2-13 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 the “[DHCP Failover](#)”

section on page 2-14). The Relay Agent can identify the special, distinguished route of the client's address request to the DHCP server, as coordinated between the Relay Agent and Network Registrar administrators. The DHCP servers can now issue leases based on overlapping IP addresses to the clients on both VPNs.

Figure 2-13 Virtual Private Network DHCP Configuration



Subnet Allocation and DHCP Address Blocks


Network Registrar supports creating on-demand address pools as a network infrastructure for address provisioning and VPNs. Traditionally, the DHCP server is limited to interact with individual host devices. Through subnet allocation, the server can interact with VPN routers and other provisioning devices to provision entire IP subnets. This Network Registrar feature enhances the on-demand address pool capability currently supported by the Cisco IOS Relay Agent.

Network Registrar supports explicitly provisioned subnets. You must explicitly configure the DHCP server's address space and subnet allocation policies before the server can allocate pools or leases. You can thereby configure a server as a pool manager to manage subnets and delegate them to client devices.

You manage DHCP subnet allocation using DHCP server address block objects in Network Registrar. A DHCP address block is a range of contiguous IP addresses delegated to the DHCP server for assignment. The server expects to subdivide these addresses into pools so that it or other servers or devices can allocate them. DHCP address blocks are parents to subnets. These DHCP address blocks are distinct from the address blocks you can create using the Network Registrar Web UI, which are static. DHCP address blocks cannot include static address ranges or lease reservations.

Figure 2-14 shows a sample environment where a DHCP server allocates entire subnets to access concentrators or other provisioning devices, in addition to servicing individual clients. The traditional client/server relationship is shown on the left of the diagram, while the subnet allocation to access concentrators is shown on the right of the diagram. Dialup customers, for example, connect to the service provider's network at two ISP gateways (routers), which connect to the management network segment where the DHCP server resides. The gateways provision addresses to their connected clients based on the subnet requested from the DHCP server.

Figure 2-14 Sample DHCP Subnet Allocation Configuration



Dynamic 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. Dynamic DNS update automates the task of keeping the names and addresses current. With Network Registrar's dynamic 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, Network Registrar tells the DNS server to add the host data. When the lease expires or when the host gives it up, Network Registrar tells the DNS server to remove the association.

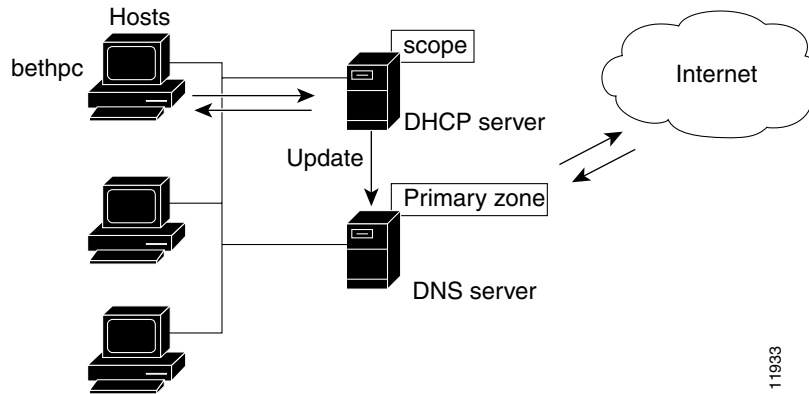
In normal operation, you do not have to manually reconfigure DNS, no matter how frequently clients' addresses change through DHCP. Network Registrar uses the host name that the client workstation provides. You also can have Network Registrar synthesize names for clients who do not provide them, or use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

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 workstation gets its owner's name. The administrator also configures the DHCP server to use dynamic 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 on to a website without having an address. When her host starts up, it broadcasts an address request (see [Figure 2-15](#)).

Figure 2-15 Dynamic DNS Update at ExampleCo Company



11933

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 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 Figure 2-16).
2. Updates the DNS server by removing the hostname and address. The DNS server no longer stores data about bethpc or its address.

Figure 2-16 Relinquishing a Lease



11934

Effect on DNS of Re-acquiring Leases

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

1. Her workstation 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.

DHCP Failover

Because DHCP, as described in RFC 2131, provides for multiple servers, you can configure these servers so that if one cannot provide leases to requesting clients, another one can take over. Network Registrar provides this capability in its DHCP failover feature, where two servers operate as redundant partners. Existing DHCP clients can continue to keep and renew their leases without needing to know or care which server is responding to their requests.

How Failover Works

Failover is based on a partner server relationship. The partners must have identical scopes, leases, policies, and client-classes. After the servers start up, each contacts the other. The main server provides its partner with a private pool of addresses and updates its partner with every client operation. If the main server fails, then the partner takes over offering and renewing leases, using its private pool. When the main server becomes operational again, it re-integrates with its partner without administrative intervention. These servers are in a relationship known as a failover pair.

The failover protocol keeps DHCP operational if:

- The main server fails—The partner takes over services during the time the main server is down. The servers cannot generate duplicate addresses, even if the main server fails before updating its partner.
- Communication fails—A partner can operate correctly even though it cannot tell whether it was the other server or the communication with it that failed. The servers cannot issue duplicate addresses, even if they are both running and each can communicate with only a subset of clients.

Failover configurations are usually in a simple, back office, or symmetrical fashion. Once configured:

1. The partners connect.
2. The main server supplies data about all existing leases to its partner.
3. The backup server requests a pool of backup addresses from the main server.
4. The main server replies with a percentage of available addresses from each scope to its partner.
5. The backup server ignores all DHCPDISCOVER requests, unless it senses that the main server is down. In normal operations, it handles only DHCPRENEW and DHCPREBINDING requests. A DHCPDISCOVER request is a broadcast to locate available servers.
6. The main server updates its partner with the results of all client operations.

Failover States and Transitions

During normal operation, the failover partners transition between states. They stay in their current state until all the actions for the state transition are completed and, if communication fails, until the conditions for the next state are fulfilled. The states and their transitions are described in [Table 2-1](#).

Table 2-1 Failover States and Transitions

State	Server Action
STARTUP	Tries to contact its partner to learn its state, then transitions to another state after a short time, typically seconds.
NORMAL	<p>Can communicate with its partner. The main and backup servers act differently in this state:</p> <ul style="list-style-type: none"> • The main server responds to all client requests using its pool. If its partner requests a backup pool, the main server provides it. • The backup server only responds to renewal and rebinding requests. It requests a backup pool from the main server.
COMMUNICATIONS-INTERRUPTED	<p>Cannot communicate with its partner, whether it or the communication with it is down. The servers cycle between this state and NORMAL state as the connection fails and recovers, or as they cycle between operational and nonoperational. During this time, the servers cannot give duplicate addresses.</p> <p>During this state, you usually do not need to intervene and move a server into the PARTNER-DOWN state. However, this is not practical in some cases. A server running in this state is not using the available pool efficiently. This can restrict the time a server can effectively service clients.</p> <p>A server is restricted in COMMUNICATIONS-INTERRUPTED state:</p> <ul style="list-style-type: none"> • It cannot re-allocate an expired address to another client. • It cannot offer a lease or renewal beyond the maximum client lead time (MCLT) longer than the current lease time. The MCLT is a small additional time added that controls how much ahead of the backup server's lease expiration the client's is. • A backup server can run out of addresses to give new clients, because it normally has only a small pool, while the main server has most of them. <p>The server is limited only by the number of addresses allocated to it and the arrival rate of DHCPDISCOVER or INIT-REBOOT packets for new clients. With a high new client arrival or turnover rate, you may need to move the server into PARTNER-DOWN state more quickly.</p>
PARTNER-DOWN	<p>Acts as if it were the only operating server, based on one of these facts:</p> <ul style="list-style-type: none"> • The partner notified it during its shutdown. • The administrator put the server into PARTNER-DOWN state. • The safe period expired and the partner automatically went into this state. <p>In this state, the server ignores that the other server might still operate and could service a different set of clients. It can control all its addresses, offer leases and extensions, and re-allocate addresses. The same restrictions to servers in COMMUNICATIONS-INTERRUPTED state do not apply.</p> <p>Either server can be in this state, but only one should be in it at a time so that the servers do not issue duplicate addresses and can properly resynchronize later on. Until then, an address is in a pending-available state.</p>

Table 2-1 Failover States and Transitions (continued)

State	Server Action
POTENTIAL-CONFLICT	Might be in a situation that does not guarantee automatic re-integration, and is trying to re-integrate with its partner. The server might determine that two clients (who might not be operating) were offered and accepted the same address, and tries to resolve this conflict.
RECOVER	Has no data in its stable storage, or is trying to re-integrate after recovering from PARTNER-DOWN state, from which it tries to refresh its stable storage. A main server in this state does not immediately start serving leases again. Because of this, do not reload a server in this state.
RECOVER-DONE	Can transition from RECOVER or PARTNER-DOWN state, or from COMMUNICATIONS-INTERRUPTED into NORMAL state.
PAUSED	Can inform its partner that it will be out of service for a short time. The partner then transitions to COMMUNICATIONS-INTERRUPTED state and begins servicing clients.
SHUTDOWN	Can inform its partner that it will be out of service for a long time. The partner then transitions to PARTNER-DOWN state to take over completely.

Allocating Addresses Through Failover

To keep your failover pair operating in spite of a network partition, in which both can communicate with clients but not with each other, you must allocate more addresses than are needed to run a single server. Configure the main server to allocate a percentage of the currently available (unassigned) addresses in each scope's 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.

How many additional addresses are needed? There is no single percentage for all environments. It depends on the arrival rate of new DHCP clients and the reaction time of your network administration staff. The backup server needs enough addresses from each scope 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.

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 re-allocating any leases. When these times expire, the backup server offers:

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

During the day, 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 nonpeak times, expressed as a percentage of the currently available (unassigned) addresses in each scope.

Client-Class Quality of Service

Assigning classes to clients is an important adjunct to DHCP addressing. You can use the 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 Network Registrar's 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, 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's access, he can register his laptop with the 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.

Here is what the DHCP server does:

1. Assigns an address to the client from a defined scope—To choose an address for the client, the DHCP server determines the client's 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, 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 selection by using DHCP server's address allocation priority feature (see the [“Configuring Multiple Scopes Using Allocation Priority”](#) section on page 11-6). After the server chooses a scope, it chooses an available (unassigned) address from that scope.

- a. Assigns DHCP option values from a defined policy—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:
 - a. If the scope-specific policy contains the option, the server returns its value to the client and stops searching.
 - b. If not found, the server looks in the system default policy, returns its value, and stops searching.

- c. If neither policy contains the option, the server returns no value to the client and logs an error.
 - d. The server repeats this process for each requested option.
2. With dynamic DNS update in effect, the server assigns an FQDN to the client—If you enabled dynamic DNS update, Network Registrar enters the client’s name and address in the DNS host table. See the “[Dynamic DNS Update](#)” section on page 2-12. The client’s name can be:
- Its name as specified in the client’s lease request (the default)
 - 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 the “[DHCP Processing Without Client-Classes](#)” section on page 2-17, 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’s 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, the client is made a member of this client-class.
 - b. No MAC address, it uses the default client specification. For example, if the client is assigned guest access, its client specification is Guest.
 - c. No MAC address and no default client, the server handles the client through regular DHCP processing.
 - d. A MAC address, the server checks if the client is a member of a client-class, determines its subnet based on the request packet, and applies the appropriate access properties based on its scope assignment.

The scopes must have addresses on client-accessible subnets—they must have a scope-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 scope-selection tag of Employee or Guest, but not both. In this case, there are two scopes for each subnet—one with the scope-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 host name the client requests, the server can:
 - Provide an explicit host name that overrides it.
 - Drop the client-requested host name and not replace it.
 - Synthesize a host name from the client’s 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 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 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's network.

Choosing Networks and Scopes

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

- Gateway address (*giaddr*), if there was one, for packets sent through a BOOTP relay.
- 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.

The Network Registrar scope-selection tag and the client-class features 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 scope-selection tags, determines the scope on the network to use.

Trivial File Transfer

The Trivial File Transfer Protocol (TFTP) is a way of transferring files across the network using the User Datagram Protocol (UDP), a connectionless TCP/IP transport layer protocol. Network Registrar maintains a TFTP server so that systems can provide device provisioning files to cable modems that comply with the Data Over Cable Service Interface Specification (DOCSIS) standard. The TFTP server buffers the DOCSIS file in its local memory as it sends the file to the modem. When the TFTP transfer is completed, the server flushes the file from local memory. TFTP also supports other, non-DOCSIS configuration files.

Here are some of the features of the Network Registrar TFTP server:

- Complies with RFCs 1350 and 1123.
- Includes a high performance multithreaded architecture.
- Caches data for performance enhancements.
- Is configurable and controllable using the **ftfp** command in the CLI.
- Includes flexible path and file access controls.
- Includes audit logging of TFTP connections and file transfers.
- Has a default root directory in the Network Registrar *install-path/data/tftp*.

Simple Network Management

The Network Registrar Simple Network Management Protocol (SNMP) notification support allows you to be warned of error conditions and possible problems with the DNS and DHCP servers, and to monitor threshold conditions that may indicate failure or impending failure conditions.

How Notification Works

Network Registrar SNMP notification support allows a standard SNMP management station to receive notification messages from the two servers. These messages contain the details of the event that triggered the SNMP trap.

Network Registrar generates notifications in response to predetermined events that are detected and signaled by the application code. In addition to the knowledge that a particular event occurred, each event can also carry with it a particular set of parameters or current values. For example, the *free-address-low-threshold* event may occur in the FDDI-Devices scope with a value of 10 percent free. Other scopes and values are also possible for such an event and each type of event can have different parameters associated with it. The scope level threshold settings override those set globally.

[Table 2-2](#) describes the events that can generate notifications.

Table 2-2 Notification Events

Event	Result
Address conflict with another DHCP server detected	Notification when an address conflict with another DHCP server is detected.
Configuration mismatch	Notification when a configuration mismatch between DHCP failover partners occurs.
DNS queue becomes full	Notification when the DHCP server's DNS queue fills and the DHCP server stops processing requests. This is a rare internal condition.
Duplicate IP address detected	Notification whenever a duplicate IP address is detected.
Change in free address count	The <i>free-address-low</i> trap when the number of free IP addresses becomes less than or equal to the low threshold; or a <i>free-address-high</i> trap when the number of free IP addresses exceeds the high threshold after having previously triggered the <i>free-address-low</i> trap.
Other server not responding	Notification when another server (DHCP, DNS, or LDAP) stops responding to the DHCP server.

Table 2-2 Notification Events (continued)

Event	Result
Other server responding	Notification when another server (DHCP, DNS, or LDAP) responds after having been unresponsive.
Server start	Notification whenever the DHCP or DNS server is started or re-initialized.
Server stop	Notification whenever the DHCP or DNS server is stopped.

Handling Notification Events

When Network Registrar generates a notification, a single copy of the notification is transmitted as an SNMP Trap Protocol Data Unit (PDU) to each recipient. The list of recipients and other notification configuration data are shared by all events (and scopes) and are read when the server is initialized.

You configure notifications using the **trap** command in the CLI. The notification configuration information is persistent and is re-initialized when you run the **reload** command on the respective server. (See the **trap** command section in the *Network Registrar CLI Reference Guide*.)

To use SNMP notifications on your system, you must specify trap recipients. These recipients indicate where Network Registrar notifications are directed. By default, all notifications are enabled, but no trap recipients are defined. Until you define the recipients, no notifications are sent. For details about adding recipients, see the **trap addRecipient**, **listRecipients**, and **removeRecipient** command sections in the *Network Registrar CLI Reference Guide*.

Network Registrar implements SNMP Trap PDUs according to the SNMP v1 standard. Each trap PDU contains:

- Generic-notification code, if enterprise-specific
- Specific-notification field that contains a code indicating the event or threshold crossing that has occurred
- Variable-bindings field that contains additional information about certain events

Refer to the Management Information Base (MIB) for the details. You can find the MIB in the locations based on the operating system in the following subsections. The MIB requires these MIB files to compile:

- SNMPv2-SMI.my
- SNMPv2-CONF.my
- SNMPv2-TC.my
- CISCO-TC.my
- CISCO-SMI.my

These individual MIBs are available at this public website:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

The location of the CISCO-NETWORK-REGISTRAR-MIB.my file in the default Network Registrar installations is on:

- Windows—*install-path*\misc
- Solaris and Linux—*install-path*/misc

Low Disk Space SNMP Trap

The Network Registrar **mcshadow** utility generates a new SNMP trap if the utility cannot create a shadow backup of the data due to inadequate disk space.

SNMP Troubleshooting

To diagnose Network Registrar conditions, set the following SNMP traps using the CLI (you can also override the free address threshold levels at the scope level for DHCP). These traps ensure that a central monitor is informed when unexpected events occur, so that you can respond more quickly before things become critical:

```
nrcmd> trap show
nrcmd> trap enable address-conflict
nrcmd> trap enable other-server-not-responding
nrcmd> trap enable free-address-low
nrcmd> trap set free-address-low-threshold=15%
nrcmd> trap set free-address-high-threshold=30%
```

You can also set these traps:

```
nrcmd> trap enable server-start
nrcmd> trap enable server-stop
nrcmd> trap enable free-address-high
nrcmd> trap enable dns-queue-too-big
nrcmd> trap enable other-server-responding
nrcmd> trap enable duplicate-address
```

The free-address traps catch when the number of free IP addresses on a server falls below a certain threshold, and when to notify that they again move out of this area. You arm the traps using the **trap enable free-address-low** and **trap enable free-address-high** commands. You set the thresholds for each using the **trap set free-address-low-threshold** and **trap set free-address-high-threshold** commands, respectively. You set the low and high threshold values either by an absolute number, or by a percentage (followed by the percent sign). You must use the same unit of measure for both thresholds; for example, if the low threshold value is a percentage, the high threshold value must be as well. The *free-address-low* trap catches when the free addresses fall below the low threshold. The *free-address-high* trap catches when they are no longer too low. The high value must be equal to or greater than the low one. Both values default to 20 percent. These traps, like all others, apply on a server and not a scope by scope basis.

You generally set the low and high thresholds at a certain offset. For example, you can set the low value to 20%, in which case the DHCP server catches when the number of free addresses fall below 20%. You can then set the high threshold to 25% so that you get a notification at a slightly higher point that the addresses have again become free. As soon as the DHCP server issues a trap for one threshold condition, it arms the trap for the opposite condition. Because of this, creating a safety zone between the two thresholds eliminates issuing traps each time the free addresses hover close to and keep crossing the low threshold point.

Even if you disable one trap through the **trap disable** command, Network Registrar still sends its opposite as needed.