White Paper

Best Practices: Upgrading to DHCPv6 Failover

Using Cisco Prime Network Registrar

November 2013

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Prior to the release of Cisco Prime™ Network Registrar Version 8.2, in which Dynamic Host Configuration Protocol Version 6 (DHCPv6) failover functionality was introduced, any one of several alternatives may have been used to provide increased availability in a DHCPv6 deployment. Now that DHCPv6 failover is available, this document provides information on how to upgrade from the most frequently used alternative approaches to DHCPv6 failover using Cisco Prime Network Registrar.

This document will first explore the alternatives that have been recommended prior to the availability of Cisco Prime Network Registrar 8.2 DHCPv6 failover capabilities. If you haven't deployed one of these alternatives in the past, then you don't need to read this document - you can just implement DHCPv6 failover by using the latest version of Cisco Prime Network Registrar. However, if you have implemented one of these alternatives, or something similar, and you want to upgrade your DHCPv6 deployment to use DHCPv6 failover, then you should identify the alternative approach that has been used and find the upgrade instructions appropriate to that approach.

The Alternatives to DHCPv6 Failover

One of the main reasons for implementing DHCPv4 failover was to avoid having to split the address space between two DHCPv4 servers. This was problematic because IPv4 address space was (and now is even more so as the demand for new IP addresses has exploded) a scarce resource and few network operators could obtain twice the address space they required. DHCPv4 failover eliminated the need to split the address space, though it does require a small increase in the address space necessary to allow each failover partner some pool of addresses to use when unable to communicate with its partner. IPv6 generally does not have these limitations.

There are also some other differences between DHCPv4 and DHCPv6, or more precisely between IPv4 and IPv6, that are important to point out. In particular, IPv6 provides for graceful address (or prefix) transition between addresses while IPv4 has a general concept of one address per interface. With IPv6, an interface is expected to have multiple addresses and each address has its own lifetime. Addresses transition between states (see RFC 4862), including:

- Tentative to preferred (this occurs after testing for address conflicts). An address in the preferred state can be used for communication.
- Preferred to deprecated (this occurs when the preferred lifetime has expired). An address in the deprecated state may continue to be used for existing communications but should not be used for new communications.
- Deprecated to invalid (when the valid lifetime has expired, the address is removed and is no longer usable).

BasicDHCPV6 Operation Overview

A DHCPv6 client sends a Solicit message and waits for one or more Advertise messages. The client selects one of the Advertise messages, and sends a Request message. The DHCPv6server selected by the client then responds with a Reply message with the lease(s) that the client has been assigned. This is very similar to the DHCPDISCOVER/DHCPOFFER and DHCPREQUEST/DHCPACK exchange used with DHCPv4.

For a more thorough description of DHCPv6 operations see the Cisco®DHCPv6 Based IPv6 Access Service white paper.

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Two Servers - Split Address Space

One technique for redundancy in DHCPv6 is for the two servers to split the address space. For example, if devices are being assigned addresses from a /64 prefix, one server could be given the 64-bitprefix0:0:0:0/65 from which to allocate addresses and the other server the 64-bitprefix8000:0:0:0/65 from which to allocate addresses.

When a client gets a lease, it will get it from one of the servers and will usually be able to renew the address from that same server. Only if that server is down during the renewal period (50 percent to 85 percent of the lease time; see the “Longer Lease Times” section below), will the client likely end up having to obtain a new address.

But, if the client gets a lease from a different server, it will transition to that new address gracefully per the IPv6 rules. (**Note**: Some client types, such as cable modems, do not necessarily follow the IPv6 address transition rules and may switch to the new address immediately.)

There are two possibilities for configuring Cisco Prime Network Registrar DHCPv6 servers for split address spaces:

- Configure /64 prefixes with a /65 range (recommended).
- Configure /65 prefixes (not recommended).

Configuring the /64 prefixes with a /65 range is the recommended approach. The reason for this is that if a DHCPv6 client issues a Confirm message, and if the /64 prefixes are configured, either server will respond with success to the Confirm if the client has an address in that /64 prefix. If instead the /65 prefixes are configured, one of the servers will respond with "NOT ON LINK", which is not desirable (as this will cause the client to solicit a new address).

An example of prefixes configured for address assignment is shown below (note that many unset/default attributes have been excluded from this discussion for clarity):

On Server 1:

```
nrcmd> prefix server1
100 Ok
server1:
    address = 2001:db8::/64
    dhcp-type = [default=dhcp]
    range = 2001:db8::/65
```

On Server 2:

```
nrcmd> prefix server2
100 Ok
server2:
    address = 2001:db8::/64
    dhcp-type = [default=dhcp]
    range = 2001:db8:0:0:8000::/65
```

Note: The above applies to prefix delegation as well. For example, if you wanted to delegate prefixes from a /48 prefix, you would configure a prefix delegation prefix on both servers with the /48 prefix but one server with the range of 48-bitprefix0:0:0:0:0/49 and the other the range of 48-bitprefix8000:0:0:0:0/49.

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

An example of prefixes configured for prefix delegation is shown below (note that many unset/default attributes have been excluded for clarity):

On Server 1:

```
nrcmd> prefix server1-pd
100 Ok
server1-pd:
    address = 2001:db8::/32
    dhcp-type = prefix-delegation
    range = 2001:db8::/33
```

On Server 2:

```
nrcmd> prefix server2-pd
100 Ok
server2-pd:
    address = 2001:db8::/32
    dhcp-type = prefix-delegation
    range = 2001:db8:8000::/33
```

Two Servers - Interface Identifier Addresses

If a DHCP server is only performing address assignment (and not prefix delegation), it may be possible to use the client's IPv6 link-local address interface identifier to generate the address (see RFC 4291, Appendix A on EUI-64 formatted interface identifiers). Both servers would then always generate the same address, and thus the client could use either server just as easily.

To enable this, a prefix's allocation-algorithm attribute must be configured to only use interface identifiers.

An example of a prefix configured for interface-identifier-based address assignment is shown below (note that many unset/default attributes have been excluded for clarity):

```
nrcmd> prefix sample1
100 Ok
sample1:
    address = 2001:db8::/64
    allocation-algorithms = interface-identifier
    dhcp-type = [default=dhcp]
```

Two Servers - External Address Assignments

Another technique that might be particularly interesting to network operators with DHCPv6 servers performing prefix delegation to clients is to use external address or prefix delegation assignments.

The basic principle is that the DHCPv6 servers would query an external service that would provide the delegated prefix (or address) to the DHCPv6 server - which would then lease the prefix or address to the client. The "generate-lease" extension point (which is only available for DHCPv6) would be used to do this.

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Using LDAP

One option would be to use Lightweight Directory Access Protocol (LDAP) as the external source. Cisco Prime Network Registrar is configured to query the LDAP server for the client's entry. An LDAP attribute would contain the delegated prefix or address. This LDAP attribute is mapped to a Cisco Prime Network Registrar environment dictionary attribute. When the DHCPv6 server needs to assign a new prefix or address, it calls the "generate-lease" extension point. That code then uses the environment dictionary attribute initialized from LDAP to set the prefix or address that the server is to use for the client. The allocation-algorithm attribute (for the prefixes for which this is to be used) should be modified to specify that only extensions may be used to generate a new lease. This does require that each client be registered in LDAP.

Using Other External Sources

This technique can also be used with other external sources (databases or other servers), but requires a more complex extension. In general, the concepts are the same as those for LDAP but the query must be done by the extension. And this is where things get a bit tricky, as it is an extremely bad idea to "block" the DHCPv6 server's processing while waiting for a reply from the external source. The best course of action is to keep a cache and look up the client in that cache to see if the client's information is present and not stale; if so, the cached information can be used to provide the prefix or address. If not available or the information is stale, a query is initiated and the DHCPv6 server is told to drop the packet. The client will retransmit the request (assuming the other server hasn't responded), and hopefully by then the information will be in the cache.

There are several important notes regarding the above:

- Dropping the packet (setting drop to "true") is not allowed from the generate-lease extension. Rather, that extension would set "skip-lease" to "true" n the environment dictionary and also add some other flag (such as drop-request) to the environment dictionary. A later extension (such as atprepacketencode) would check this flag and request the packet to be dropped.
- The extension must read replies from the external server and place that information in the cache. This can be done either in a separate thread or whenever a new packet arrives before checking the cache.
- As occurs with LDAP, it is best to do the lookup for the client much earlier in processing (such as at postpacketdecode, preclientlookup, postclientlookup, and/or postclasslookup), store the delegated prefix(es) or address(es) in the environment dictionary, and use them later when generate-lease is called.

Multiple Techniques

Cisco Prime Network Prime Registrar's DHCPv6 server can actually use different techniques for different prefixes at the same time. This may be useful to optimize the assignment technique based on the client types or prefix type (stateful address versus prefix delegation). The techniques used are controlled by the prefix configuration, especially the allocation-algorithms attribute.

Longer Lease Times

One reason DHCPv4 lease times are configured to be relatively short is that a service provider has a limited number of addresses. However, DHCPv6 addresses are not in short supply. Thus, preferred and valid lifetimes for DHCPv6 leases can be set much longer.

Clients have between 50 percent and 80 percent of the preferred DHCPv6 lease lifetime in which to renew the address from the original server. Thus, for a 30-day preferred lifetime, the client will periodically try to RENEW its lease between 15 days and 24 days; and only after 24 days will it attempt to contact any available server using a REBIND message - which may then trigger the client to solicit a new lease.

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

The important point here, though, is that with longer lease times, an operator has more time in which to restore a downed server before existing clients will attempt to contact another DHCPv6 server (24 days with 30-day preferred lifetimes). However, new clients will need to get leases from the other server(s) during this time.

Preferring One Server over Another

The DHCPv6 protocol was design with a means to control which server a client chooses to use. This is done using the DHCPv6 Preference Option (see RFC 3315, sections 22.8 and 17.1.2).

One server can be configured with a high preference or the highest preference (255). Using a preference of 255 has the advantage that a client need not wait for additional server responses after receiving an Advertise message with a preference of 255 (as normally a client waits for a short period to receive responses from all servers). The other server can be configured with a lower preference and thus will be less preferred.

When both servers respond to a client's Solicit with an Advertise, the client will prefer the higher preference server or immediately use the 255 preference server.

Note that this preference only applies to the Solicit/Advertise phase. For Cisco Prime Network Registrar, the preference is specified by adding an instance of the preference option on an appropriate policy. For example, to set the default server wide preference:

```
nrcmd> policy system_default_policy setv6option preference 255
```

Issues with Independent Servers

Common issues with using multiple independent servers may include:

- Domain Name System (DNS) updates will occur correctly and multiple addresses can be entered under a device's name. This generally isn't a problem, but it may slow down connections to that device if one of those addresses is not in use by the device. This can happen if Server A provides the original lease, the device is shut down, Server A goes down, and the device is powered up and obtains a new address from Server B. The DNS will now contain both addresses, but only Server B's address is in use by the client.
- DNS updates may be prematurely removed if an external source assigns the address or delegated prefix or if interface-identifier-based address generation is used. This can happen in the previously mentioned situation, as when Server A returns, it may find that the lease expired and therefore remove the DNS mapping for the client.
- Lease query for cable source address verify or lease recovery by a relay agent (CMTS) after reboot can be more complex and troublesome. When one of the DHCPv6 servers is down, it cannot respond to lease-query requests, and thus the relay agent is unable to obtain information. And when both servers are up, one server may respond that the address is leased to the client but the other may indicate it is not.
- Lease information can be lost if a server's disk is lost or if the lease database is irretrievably corrupted. DHCPv4 failover has been used to provide a "backup" of the lease data. For DHCPv6, this can result in the risk of duplicate address assignment if the disk or database is lost on one server. However, as Cisco Prime Network Registrar generates addresses using a random number generator (that uses up to 62 bits of randomness), the probability is fairly small that a duplicate address will be generated.

These are key reasons why a DHCPv6 failover solution will have significant value.

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Upgrading to DHCPv6 Failover

The remainder of this document describes how to upgrade to DHCPv6 failover from two different existing deployment approaches, each of which uses two servers and each of which uses the split address approach discussed first above:

- Two servers that are currently DHCPv4 failover partners, but each server's DHCPv6 configuration uses the split address space approach. In this situation, failover will be extended to cover not only DHCPv4 addresses, but also DHCPv6 addresses.
- Two servers that have no existing connection - no DHCPv4 failover or DHCPv4 configuration. These two servers are also using the split address approach. In this situation, DHCPv6 failover will be configured since there is no DHCPv4 failover configuration to extend.

Note: DHCPv4 failover as implemented in Cisco Prime Network Registrar Version 8.2 and later is incompatible with prior versions, and therefore both local clusters must be upgraded before the two servers can connect as failover partners. There are two distinct differences in DHCPv4 failover in Cisco Prime Network Registrar Version 8.2:

- All failover (DHCPv4 and DHCPv6) operates over TCP, not User Datagram Protocol (UDP).
- Failover covers all scopes and prefixes (only "simple" failover is supported).

Thus, the DHCPv4 failover capability in Cisco Prime Network Registrar Version 8.2 will not interoperate with the DHCPv4 failover capability available in previous versions of the product. Once both servers are upgraded to Version 8.2 or later, they can communicate as failover partners.

Please refer to the Cisco Prime Network Registrar 8.2 Release Notes and Installation Guide for additional details.

Overview of the Procedures

The procedures discussed below make use of failover to "merge" the lease databases of the two DHCPv6 servers.

The procedures need to be followed carefully, as skipping steps or reordering them may cause problems in merging the lease state databases.

The procedures will merge the lease data. This has some implications when each existing server has lease information for the same client bindings, as the merged data will result in the client having multiple leases per binding. While such a situation is perfectly reasonable with respect to the DHCPv6 protocol, it may not be desired in a particular installation. This can be corrected over time by using the new policy max-leases-per-binding attribute (which can be applied server wide by changing the value in the system_default_policy setting). However, the leases will not begin to be corrected until the client communicates with the server, and even then it can take some time to assure that the client has received the updated information and is no longer using the lease.

The procedures also assume that all prefixes used the split prefix approach discussed above and, further, that all of the prefixes will be covered by failover (particularly since no other approach is possible in Version 8.2 and future versions of Cisco Prime Network Registrar).

It is also important that firewalls between the DHCP clusters be configured to allow the failover communication over TCP. The default failover port is 647 (unless explicitly configured by the customer). These connections will always be initiated from the main server to the backup server's failover port (647).

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Note that after the upgrade is complete, if firewall rules had permitted UDP traffic on the failover port, they can be removed as UDP is no longer used for failover in Cisco Prime Network Registrar Version 8.2.

Should you wish to do so, you can adjust the failover configuration to use TCP over IPv6 (with proper firewall adjustments) instead of IPv4, although either transport works correctly. However, it is recommended that you wait until after the upgrade is complete to perform this adjustment. In order to make this change, configure the fail over pair's main IPv6address and backup IPv6address attributes (note that both must be set to use IPv6) and assure that both partners are updated and reloaded.

Upgrading Existing DHCPv4 Failover Partners

For a failover pair, the steps are as follows.

1. On both the main and backup failover clusters, set the DHCP server's start-on-reboot setting to disable restarts (that is, **nrcmd> dhcp disable start-on-reboot**).
2. Upgrade the main and backup failover clusters to Cisco Prime Network Registrar Version 8.2. Follow the normal procedures. Note: The DHCP server will not be running after the upgrade because start-on-reboot was disabled. Do not start the server until instructed to do so!
3. On the **main** cluster:
    a. Modify all of the prefixes to remove the split range. (See below.)
    b. Set the DHCP server to only process DHCPv4 packets - this requires expert mode (that is, **nrcmd**> **session set visibility=3**) and is done by setting the DHCP server's dhcp-support attribute to v4-only (that is, **nrcmd> dhcp set dhcp-support=v4-only**).

       This step is critical!
    c. You can now start the main DHCP server if you want to restore DHCPv4 service. If you do, be sure to set it into partner-down mode (that is, **nrcmd> failover-pair pair-name setpartnerdown**). The remaining steps before the backup can come online may take a bit of time, but you will still have DHCPv4 service during this period. Note that **no** DHCPv6 packets will be received/processed, as the server was set to only start DHCPv4 interfaces, above.
    d. Set the max-leases-per-binding to 1. This is best done on the system_default_policy as then it applies to all DHCPv6 leasing activity; but it may also be set on more specific policies. (that is, **nrcmd> policy system_default_policy set max-leases-per-binding=1**).
    e. Unset the server preference option; likely this was set in the system_default_policy (that is, **nrcmd> policy system_default_policy unsetV6Option server-preference**). As an alternative, you may set it to a specific value (but this will apply to the partner as well).
    f. Start or reload the main server.
    g. Review the logs to assure that the DHCPv6 configuration is correct. If there are issues, correct them and reload the server and re-review the logs. No DHCPv6 packets are being processed, but the configuration is being loaded and any issues identified by the server will appear as warning messages. But you shouldn't just be looking for warning messages - you **must** ensure that the prefix configuration is correct at this stage.
    h. Perform failover configuration synchronization on the fail over pair from main to backup. Note that ideally this should be done in **exact** mode to make sure that the configurations on the main and backup are completely the same. If you had any intentional difference between the configurations, **exact** mode will remove them and you will have to reconfigure them.

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

4. Next, on the **backup** cluster:
    a. Start the DHCP server. DHCPv4 and DHCPv6 failover should operate now and the servers will exchange the leases with each other.
    b. Let the servers operate for several minutes to assure that all binding updates are exchanged. Review the DHCP server failover statistics (that is, **nrcmd> dhcp getstats failover**) and assure that the "v6-binding-" received and sent counts have stabilized (when viewed several seconds apart). Note that while no DHCPv6 incoming packets are being processed, timeouts may still occur and result in a small number of binding updates.
    c. Review the backup server's log (grep or search for "Error" or "Warning" messages).
5. Next, on the **main** cluster:
    a. Unset the expert-mode DHCP server's dhcp-support attribute (that is, **nrcmd> dhcp unset dhcp-support**).
    b. Enable start-on-reboot for the DHCP server (that is, **nrcmd> dhcp enable start-on-reboot**).
    c. Optional - You can do a failover configuration synchronization here (main to backup) as that avoids having to make these changes on the backup as well (below).
    d. Reload the **main** server. At this point the main server is live and processing DHCPv6 requests from DHCPv6 clients.
6. Next, on the **backup** cluster:
    a. If no failover configuration synchronization was done from the main to backup (in step 5 c above):
        i. Unset the expert-mode DHCP server's dhcp-support attribute (that is, **nrcmd> dhcp unset dhcp-support**).You will have to go into expert mode to perform this operation (see step 3b above).
        ii. Enable start-on-reboot for the DHCP server (that is, **nrcmd> dhcp enable start-on-reboot**).
    b. Reload the server.

The servers are now upgraded and operating in DHCPv6 (and DHCPv4) failover mode.

Note: If you do NOT want to leave the **max-leases-per-binding** in place permanently, you can remove this setting (and reload) **after** waiting at least the longest lease lifetime after upgrading.

However, be aware that failover can end up assigning multiple leases for a single client binding (from the same allocation group/prefix). This can happen in situations where the client obtains leases from one server, that server goes down before updating the partner, and then the DHCPv6 client undergoes a "reset" and sends Solicit messages to the partner; the partner will assign new leases. When the partners reconnect, the client would end up with multiple leases/binding. In practice this situation is rare, because one server rarely goes down prior to updating its failover partner with failover information, but it can happen.

Upgrading Two Independent Partners

When the DHCPv6 clients are being serviced by two independent DHCPv6 servers using the split address space approach and with no DHCPv4 support is involved, these are the steps to create a DHCPv6 failover pair:

1. On both the clusters, set the DHCP start-on-reboot setting to disable restarts (that is, **nrcmd> dhcp disable start-on-reboot**).
2. Upgrade both clusters to Cisco Prime Network Registrar Version 8.2. Follow the normal procedures. Note: The DHCP server will not be running after the upgrade because **start-on-reboot** was disabled.

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

3. Select one server to be main and other to be backup.
4. On the selected **main** cluster:
    a. Create a cluster object for the backup partner.
    b. Create a failover-pair object using the local cluster and the backup cluster.
    c. Modify all of the prefixes to remove the split range. (See below.)
    d. Set the **max-leases-per-binding** to 1. This is best done on the system_default_policy as then it applies to all DHCPv6 leasing activity; but it may also be set on more specific policies (that is, **nrcmd> policy system_default_policy set max-leases-per-binding=1**).
    e. Unset the server preference option, likely set in the system_default_policy (that is, **nrcmd> policy system_default_policy unsetV6Option server-preference**). You may alternatively set it to a specific value (but this will apply to the partner as well).
    f. Start the **main** server. At this point, the main server is live and processing DHCPv6 requests.
    g. Move the main server to partner-down mode (that is, **nrcmd> failover-pair *pair-name* setpartnerdown**).
    h. Review the logs to assure that the DHCPv6 configuration is correct. If there are issues, correct them and reload the server and re-review the logs.
    i. Perform a failover configuration synchronization on the failover pair from main to backup. Note that ideally this should be done in **exact** mode to make sure that the configurations on the main and backup are completely the same. If you had any intentional difference between the configurations, **exact** mode will remove them and you will have to reconfigure them.
5. Next, on the **backup** cluster:
    a. Start the DHCP server. DHCPv6 failover should operate now and the servers will exchange the leases with each other. Both servers will be live to DHCPv6 client requests.
    b. Let the servers operate for several minutes to assure that all binding updates are exchanged.
    c. Review the server logs on both main and backup servers (grep or search for "Error" or "Warning" messages).
6. Next, on the main cluster enable start-on-reboot for the DHCP server (that is, **nrcmd> dhcp enable start-on-reboot**).
7. Next, on the backup cluster enable start-on-reboot for the DHCP server (that is, **nrcmd> dhcp enable start-on-reboot**).

The servers are now upgraded and operating in DHCPv6 (and DHCPv4) failover mode.

Note: If you do NOT want to leave the **max-leases-per-binding** in place permanently, you can remove this setting (and reload) **after** waiting at least the longest lease lifetime after upgrading.

Note, however, that failover can end up assigning multiple leases for a single client binding (from the same allocation group/prefix). This can happen in situations where the client obtains leases from one server and that server goes down before updating the partner and then the DHCPv6 client undergoes a "reset" and sends Solicit messages to the partner; the partner will assign new leases. Then when the partners reconnect, the client would end up with multiple leases/binding. In practice this situation is rare, because one server rarely goes down prior to updating its failover partner with failover information, but it can happen.

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Removing the Split Ranges

The following steps can be used to generate an nrcmd script that can be used to unset the range attribute on all of the prefixes. This should typically be run on the main server.

WARNING: This unsets the range on all prefixes regardless of what the value is! If the range is something other than the address prefix length + 1, this should not be used.

```
Use nrcmd to generate this script as follows:
nrcmd> session listbrief set prefix prefix <name> unset range
nrcmd> session log remove-range-script.txt
nrcmd> prefix listbrief
nrcmd> session log
nrcmd> session listbrief set prefix
```

Here's a version that can be copied and pasted into a file or into an nrcmd window:

```
session listbrief set prefix prefix <name> unset range
session log remove-range-script.txt
prefix listbrief
session log
session listbrief set prefix
```

Then, you can use the following command to execute this script:

```
nrcmd -b <remove-range-script.txt
```

Note that the above can be done with Cisco Prime Network Registrar8.1 before upgrading to Version8.2 to prepare the operations in advance (though make sure no additional prefixes are provisioned). You should not execute the script until after upgrading to Version8.2 at the appropriate place the procedure.

Note: For more information on the list brief formatting syntax, see the /opt/nwreg2/local/conf/nrcmd-listbrief-defaults.conf file.

Other Options

In some situations it may be more desirable to discard the leases from the less preferred server. This however has some complications:

1. The Cisco Prime Network Registrar 8.2 lease admin tool must be used (with Cisco Prime Network Registrar stopped) to remove all IPv6 lease records (that is, **leaseadmin -d 0::0/0**). This must be done **before** the DHCP server on the less preferred server is started. Note: If you have multiple VPNs, you must issue this command once for each VPN and specify the vpn-id (-n <vpn-id>).
2. The prefix split ranges on the main **must not** be unset until the longest lifetime for leases from the less preferred server have expired. If you have 7-day leases, this means you must wait 7 days before adjusting the prefix configuration to remove the split range. It is critical that the split range not be removed until all leases have expired. **Tip**: You may want to reduce the valid lifetimes on the less preferred server of at least the longest lease time before attempting the upgrade as that will shorten the time you will need to wait.
3. This avoids the need for the max-leases-per-binding setting (though you may still want to set that for other reasons, since when using failover this situation can occur in rare cases) - see the discussion at the end of either section on Upgrading, above, for more details.

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

DHCPv6 Failover Protocol

As of this writing, the Internet Engineering Task Force Dynamic Host Configuration Working Group (DHC WG) has begun work on DHCPv6 failover.RFC 7031 details the requirements for DHCPv6 failover. There is an active IETF draft on a design for a DHCPv6 failover protocol (though this does not specify an over-the-wire protocol). The implementation of DHCPv6 failover in Cisco Prime Network Registrar Version 8.2 follows the work described in that draft.

References

- RFC 3315 - Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
- RFC 3633 -IPv6 Prefix Options for Dynamic Host Configuration Protocol(DHCP) Version 6
- RFC 4291 - IP Version 6 Addressing Architecture
- RFC 4862 - IPv6 Stateless Address Autoconfiguration
- RFC 7031 - DHCPv6 Failover Requirements

For More Information

For more information about Cisco Prime Network Registrar, please visit http://www.cisco.com/go/networkregistrar.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Printed in USA C11-730228-00 11/13

© 2013 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.