



Managing the Central Configuration

This chapter explains how to manage the central configuration at the Cisco CNS Network Registrar regional cluster, which requires use of the regional cluster Web UI.

Central Configuration Tasks

Central configuration management at the regional cluster can involve:

- Setting up server clusters, and polling subnet utilization and lease history data from them.
- Setting up routers.
- Creating virtual private networks (VPNs), or pulling them from or pushing them to the local clusters.
- Creating DHCP scope templates, or pushing them to or pulling them from the local clusters.
- Creating DHCP policies, or pushing them to or pulling them from the local clusters.
- Creating DHCP client-classes, or pushing them to or pulling them from the local clusters.
- Managing DHCP failover server pairs.
- Managing zone distributions.

These functions are available only to administrators assigned the central-cfg-admin role. Note that central configuration management does not involve setting up administrators and checking the status of the regional servers. These functions are performed by the regional administrator, as described in the [“Licensing” section on page 4-6](#) and [“Controlling Servers” section on page 6-1](#).

Setting Up Server Clusters

Server clusters are groupings of CCM, DNS, DHCP, and TFTP servers at local cluster locations. For example, an organization might have Boston and Chicago clusters of DNS and DHCP servers. A central administrator might want to affect how addresses are allocated on these clusters, or poll subnet utilization or lease history data from them. The central administrator might even want to connect to those local clusters, if the required permissions exist, to view changes there or restart the servers.



You view the created clusters on the View Tree of Cluster Servers page. To get there, click **Clusters** on the Primary Navigation bar. Once the page is populated with clusters, it shows some rich information and provides some useful functions. [Figure 5-1](#) shows an example of two clusters with their server types, listed alphabetically. The Go Local icon () allows single sign-on to a local cluster’s Web UI, if an equivalent administrator account exists at the local cluster.

Figure 5-1 View Tree of Server Cluster Page

Name	Connect	Type	IP Address	SCP Port	Resynchronize	Replicate Data	Poll Subnet Utilization	Poll Lease History
1 All								
Boston-cluster		Cluster	10.86.145.82	1234				
DHCP		Server						
DNS		Server						
TFTP		Server						
CCM		Server						
CNRAGENT		Server						

This page might have been populated by manually adding clusters on the List Server Clusters page, or it might have occurred automatically when adding and synchronizing with routers, which also creates server clusters. The cluster names are links that you can click to edit the cluster information. The resynchronization, replication, and polling functions are described further on in this chapter.

The DHCP server may have the Related Servers icon () next to the DHCP server for the cluster. Click this icon to open the List Related Servers for DHCP Server page (see the “[Listing Related Servers for DHCP Servers](#)” section on page 5-3). These servers can be DNS, TFTP, or DHCP failover servers.

Adding Local Clusters

You add clusters manually on the List Server Clusters page. To get there, click **Cluster List** on the Secondary Navigation bar. This is core functionality of the central-cfg-admin role. The List Server Cluster page is similar to the View Tree of Server Clusters page (see [Figure 5-1](#)), except that you cannot expand the clusters to show their servers. However, you can add server clusters on the List Server Clusters page, which you cannot do on the View Tree of Server Clusters page. Both pages provide the following functions:

- Connect to a local cluster Web UI for local administration.
- Resynchronize with a local cluster to reconcile updates there.
- Pull data over to a regional cluster’s replica database.
- Query subnet utilization data from a local cluster—This function appears only if you have the address space license entered and are assigned the regional-addr-admin role with at least the subnet-utilization subrole.
- Query lease history data from a local cluster—This function appears only if you have the address space license entered and are assigned the regional-addr-admin role with at least the lease-history subrole.

To enable subnet utilization and lease history data collection, see the “[Polling Subnet Utilization and Lease History Data](#)” section on page 5-7.

To add a cluster, click **Add Cluster**. This opens the Add Server Cluster page (see [Figure 4-18 on page 4-26](#)).

The minimum required values to add a cluster are its name, IP address of the machine, administrator username, and password. The administrator account must be a superuser at the local cluster. The cluster name must be unique and its IP address must match that of the host where the CCM database is located. Obtain the SCP and HTTP ports, username, and password from the local cluster administrator. The default at Network Registrar installation for the SCP port is 1234 and the HTTP port is 8080.

You can also set whether you want outbound connections to local servers to be secure by setting the *use-ssl* attribute to optional or required. It is set to optional by default, and it requires the Network Registrar Communications Security Option installed to be effective.

Click **Add Cluster** to add the cluster and return to the List Server Clusters page.

Editing Local Clusters

To edit a local cluster, click its name on the View Tree of Server Clusters page or List Server Clusters page. The Edit Server Cluster page is essentially the same as the Add Server Cluster page, except for an additional attribute unset function. Make your changes, then click **Modify Cluster**.

Listing Related Servers for DHCP Servers


If you related DNS, TFTP, or DHCP failover partner servers, you can access the attributes for these servers from the View Tree of Server Clusters page (see [Figure 5-1 on page 5-2](#)). On this page, click the Related Servers icon () next to the DHCP server for the cluster to open the List Related Servers for DHCP Server page. This page shows the communications and failover states the servers are in. [Table 5-1](#) describes the attributes on this page. For this page to appear, you must be assigned the central-cfg-admin role with the dhcp-management subrole.

Table 5-1 Attributes for Related Servers to DHCP Servers

Related Server Attribute	Description
Related Server Address	IP address of the related server. For DHCP failover partners, click this link to open the View Failover Related Server page (see Table 5-2).
Communications	State of the communication—None, OK, or Interrupted.
Requests	Applies to DNS or LDAP related servers only, the number of requests from these servers.
State	For DHCP failover only, the server's state—None, Startup, Normal, Communications-interrupted, Partner-down, Potential-conflict, Recover, Paused, Shutdown, or Recover-done.
Partner Role	For DHCP failover only, the failover role of the partner—Main or Backup.
Partner State	For DHCP failover only, the partner's state—None, Startup, Normal, Communications-interrupted, Partner-down, Potential-conflict, Recover, Paused, Shutdown, or Recover-done.
Update Response Complete	For DHCP failover only, the percentage of completed update responses, valid only if there are outstanding update responses.

Table 5-2 Attributes for DHCP Related Failover Servers

Failover Partner Attribute	Description
General attributes:	
current-time	Current time on the server returning this object.
comm-state	None, OK, or Interrupted.

Table 5-2 Attributes for DHCP Related Failover Servers (continued)

Failover Partner Attribute	Description
maximum-client-lead-time	Current maximum client lead time (MCLT) on this system.
sequence-number	Sequence number unique across failover objects, if different from the sequence in the lease, the lease is considered “not up to date” independent of the sf-up-to-date lease flag.
Local server information:	
our-ipaddr	IP address of the interface to this server.
role	Failover role of the server returning this object—None, Main, or Backup.
state	State of the local server—None, Startup, Normal, Communications-interrupted, Partner-down, Potential-conflict, Recover, Paused, Shutdown, or Recover-done.
start-time-of-state	Time at which the current failover state began.
start-of-comm-interrupted	Time at which this partner most recently went into communications-interrupted state. This is valid across reloads, while the start-time-of-state never has a time earlier than the most recent server reload.
est-end-recover-time	Valid if <i>update-request-in-progress</i> is not set to None. If it appears, the time at which the server enters the recover-done state if the update request outstanding is complete. If it does not appear, then the server enters recover-done whenever update-request is completed.
use-other-available	If false or unset, then this server cannot use other-available leases. If true, then the server can use other-available leases. Valid at all times, but should only be true if in partner-down state.
use-other-available-time	If, in partner-down state, the <i>use-other-available</i> is false or unset, the time when <i>use-other-available</i> will go to true.
safe-period-remaining	Duration in seconds remaining in safe-period. If not set to 0, then this server is currently running down a safe period with respect to its partner.
Partner server information:	
ipaddr	IP address of the partner server.
partner-role	Failover role of the partner of the server returning this object—None, Main, or Backup.
partner-state	Last known state which the partner’s end of the failover relationship is in—None, Startup, Normal, Communications-interrupted, Partner-down, Potential-conflict, Recover, Paused, Shutdown, or Recover-done.
start-time-of-partner-state	Time at which the partner’s current failover state began.
est-partner-end-recover-time	If the <i>partner-state</i> is Recover, an estimated prediction of when the partner will time out its MCLT and finish being in recover state.
last-comm-ok-time	Time at which this server last found communications to be OK.
Update requests sent to partner:	
update-request-outstanding	If None or unset, then the server does not have an update request queued for its partner. If not set to None, then it does have an update request queued for its failover partner. Valid values are None, Update, and Update-all.
update-request-start-time	Time at which any <i>update-request-outstanding</i> request was started.



Table 5-2 Attributes for DHCP Related Failover Servers (continued)

Failover Partner Attribute	Description
update-request-done-time	Time at which the last of any update request completed.
Update requests processed for partner:	
update-response-in-progress	If this server is processing an update response, gives information about the type and origin of the response.
update-response-percent-complete	If <i>update-response-outstanding</i> appears, the percent complete of the current update response.
update-response-start-time	Time that the update response mentioned in <i>update-response-in-progress</i> was started.
update-response-done-time	Time that the most recent update response sent an update done to the partner server.

Other controls are available on these pages:

- To refresh the data on the View Failover Related Server page, click **Refresh Data**.
- On the View Failover Related Server page, if the partner is in the Communications-interrupted failover state, you can click **Set Partner Down** in association with an input field for the partner-down date setting. This setting is initialized to the value of the *start-of-communications-interrupted* attribute. (In Normal Web UI mode, you cannot set this date to be an earlier value than the initialized date. In Expert Web UI mode, you can set this value to any date.) After clicking **Set Partner Down**, you return to the List Related Servers for DHCP Server page to view the result of the partner-down action.
- To return from either of these pages, click **OK**.


Connecting to Local Clusters

If you have an equivalent administrator account on the local cluster, you can single sign-on to the local cluster's Manage Servers page by clicking the Go Local icon () next to the cluster name on the View Tree of Server Clusters page or List Server Clusters page. To return to the regional cluster Web UI, click the Go Regional icon () at the top right corner of the local cluster page. If you do not have an equivalent account on the local cluster, the Go Local icon opens the local cluster's login page.

Synchronizing with Local Clusters

Synchronization is configuring regional and local clusters so that they can work together in a unified fashion. When you synchronize:

1. The list of local servers are copied to the regional cluster.
2. A shared secret is established between the regional and local clusters for single sign-on.

Synchronization occurs once when you create a local cluster at the regional cluster. However, changes might occur on the local cluster periodically, requiring you to resynchronize with it. For example, you might change the username and password used to make local connections. Resynchronization does not happen automatically—you must click the Resynchronize icon () next to the cluster name on the List Server Clusters page. The result is a positive confirmation for success or an error message for a failure.


**Note**

For synchronization to be effective, the user account specified for the local cluster must be a superuser. If you get a synchronization error message, check the local cluster to ensure that it is running properly and includes the proper base license.

Replicating Local Cluster Data

Replication is copying the configuration data from a local server to the regional cluster's replica database. Replication occurs for DHCP scopes, address blocks, subnets, policies, scope templates, client-classes, VPNs, DNS zones, and zone templates. During replication:

1. The current data from the local database is copied to the regional cluster. This usually occurs once.
2. Any changes made in the master database since the last replication are copied over.

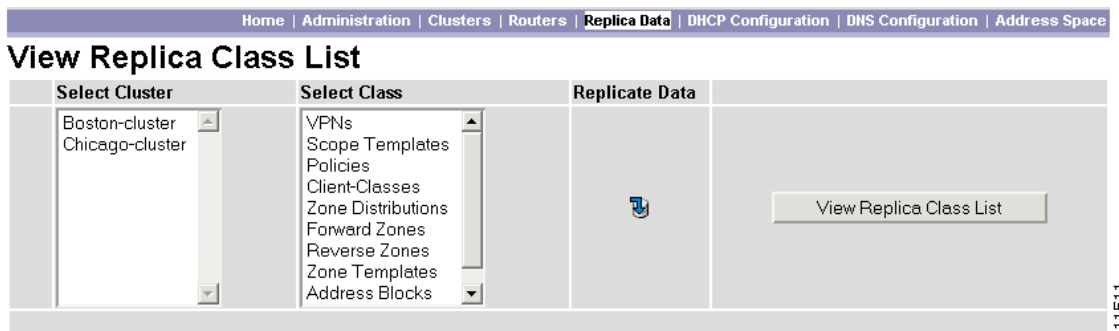
Replication happens when you first synchronize the clusters and also at a given time interval. You can also force an immediate replication by clicking the Replica icon () next to the cluster name on the View Tree of Server Clusters page or List Server Clusters page. Replication needs to occur before you can pull data into the regional server's database.

You can set the automatic replication interval on the Add Server Cluster page, or adjust it on the Edit Server Cluster page, using the *poll-replica-interval* attribute. This interval is set at four hours by default. You can also adjust the time-of-day offset between the regional and local clusters by using the *poll-replica-offset* attribute; its default is zero hours.


Viewing Replica Data

You can view the replica data cached in the replica database on the regional cluster by clicking **Replica Data** on the Primary Navigation bar. This opens the View Replica Class List page (see [Figure 5-2](#)).

Figure 5-2 View Replica Class List Page



On this page, you can:


1. Choose the cluster from the Select Cluster drop-down list.
2. Choose the object class from the Select Class drop-down list.
3. Replicate the data for the cluster and class chosen—Click the Replica icon ()

4. View the replica data—Click **View Replica Class List**, which opens the following pages based on the class you choose:
 - **VPNs**—List Replica DHCP VPNs for Cluster page
 - **Scope Templates**—List Replica DHCP Scope Templates for Cluster page
 - **Policies**—List Replica DHCP Policies for Cluster page
 - **Client-Classes**—List Replica DHCP Client-Classes for Cluster page
 - **Zone Distributions**—List Replica Zone Distributions for Cluster page
 - **Forward Zones**—List Replica Forward Zones for Cluster page
 - **Reverse Zones**—List Replica Reverse Zones for Cluster page
 - **Zone Templates**—List Replica Zone Templates for Cluster page
 - **Address Blocks**—List Replica Address Blocks for Cluster page
 - **Subnets**—List Replica Subnets for Cluster page
 - **Owners**—List Replica Subnets for Cluster page
 - **Regions**—List Replica Regions for Cluster page
 - **Administrators**—List Replica Admins for Cluster page
 - **Roles**—List Replica Roles for Cluster page
 - **Groups**—List Replica Groups for Cluster page

On each page, you can:

- Click the name of an object to open a View page on the regional cluster. Return to the List Replica page by clicking **Return to *object* List**.




Note The List Replica Address Blocks and List Replica Subnets pages do not provide this function. To view the address blocks or subnets for the local cluster, use the Go Local icon ()

- Click the Go Local icon () to go to the List page for the object on the local cluster. Return to the List Replica *object* page by clicking the Go Regional icon ()

Click **Return** on the View page, to return to the View Replica Class List page.

Polling Subnet Utilization and Lease History Data

Subnet utilization and lease history data are automatically collected at any regional cluster where these features are enabled for the DHCP server or failover pair. The default polling interval to update the regional databases is 4 hours. You can poll the servers immediately by clicking the Poll icon () for the cluster in the Poll Subnet Utilization column or Poll Lease History column on the View Tree of Server Clusters page or List Server Clusters page. For this manual polling, if the server is in a failover relationship, data is only retrieved for the subnets where the server is the main.

If you have address space privileges (you have an address space license entered in the system and you are assigned the regional-addr-admin role with at least the subnet-utilization and lease-history subroles), you can query the subnet utilization or lease history data by clicking **Address Space** on the Primary Navigation bar (see the “[Generating Subnet Utilization History Reports](#)” section on page 18-13 or “[Running IP Lease Histories](#)” section on page 12-13).

Polling Process

When the regional cluster polls the local cluster for subnet utilization or lease history, it first requests all available data up to the current time. This time is recorded in the history databases, and subsequent polls request only new data from this time forward. All times are stored relative to each local cluster's time, adjusted for that cluster's time zone.

If the times on each server are not synchronized, you might observe odd query results. For example, if the regional cluster's time lags behind a local cluster's, the collected history might be in the future relative to the time range queries on the regional cluster. If so, the result of the query would be an empty list. Data merged from the several clusters could also appear out of sequence, because of the different time skews between local clusters. This type of inconsistency would make it difficult to interpret trends. To avoid these issues, the use of a network time service for all Network Registrar clusters is strongly recommended.

Adjusting the Polling Intervals

You can adjust the automatic polling interval for subnet utilization and lease history, along with other attributes (see [Table 5-3 on page 5-8](#)). There are three places these attributes are set at the regional cluster, with the following priority:

- For the regional CCM server (the default polling interval is 4 hours)—This is set on the Edit CCM Server page, accessible from **Administration** on the Primary Navigation bar, **Servers** on the Secondary Navigation bar, then clicking the **Local CCM Server** link.
- For the cluster—These values override the server-wide settings, unless they are unset, in which case the server values are used. The cluster values are set when adding or editing the cluster.
- For each failover pair (see the [“Creating DHCP Failover Pairs” section on page 5-25](#))—These values override the cluster settings (only for subnets in the failover pair), and set additional attributes to control how polling to the backup server is done if the main is not available:
 - If the main failover server is unavailable, the subnets on the backup server are polled.
 - If there are no failover pair settings for these attributes, the main server values are used.



Note

If subnet utilization or lease history collection is not explicitly turned on at the local cluster DHCP server (see the [“Enabling Subnet Utilization Collection” section on page 5-9](#) and the [“Enabling Lease History Collection” section on page 5-10](#)), there is no data collected, even though polling is on by default. Subnet utilization collection at the DHCP server is distinct from polling at the regional cluster, and polling does not automatically trigger collection. Subnet utilization collection must occur before new polling picks up any new data. Because this collection is every 15 minutes by default, the polling interval should be set higher than this interval (the automatic polling interval is every 4 hours by default). This also means that subsequent explicit polling performed before the next *collect-addr-util-interval* will not return any new subnet utilization data.

Table 5-3 Subnet Utilization and Lease History Polling Regional Attributes

Attribute Type	Subnet Utilization	Lease History
Polling interval—How often to poll data	<i>poll-subnet-util-interval</i> 0 (no polling) to 1 year, defaults to 4 hours for the CCM server	<i>poll-lease-hist-interval</i> 0 (no polling) to 1 year, defaults to 4 hours for the CCM server




Table 5-3 Subnet Utilization and Lease History Polling Regional Attributes (continued)

Attribute Type	Subnet Utilization	Lease History
Retry interval—How often to retry after an unsuccessful polling	<i>poll-subnet-util-retry</i> 0 to 4 retries	<i>poll-lease-hist-retry</i> 0 to 4 retries
Polling priority for the regional failover pair—Pull data from the main or backup server first	<i>poll-subnet-util-server-first</i> choose mainserver (default) or backupserver	<i>poll-lease-history-server-first</i> choose mainserver (default) or backupserver
Offset—Hour of the day to guarantee polling	<i>poll-subnet-util-offset</i> 0 to 24h (0h= midnight)	<i>poll-lease-hist-offset</i> 0 to 24h (0h=midnight)

The polling offset attribute ensures that polling occurs at a specific hour of the day, set as 24-hour time, in relation to the polling interval. For example, if you set the interval to 4h and the offset to 6h (6 A.M.), the polling occurs at 2 A.M., 6 A.M., 10 A.M., 2 P.M., 6 P.M., and 10 P.M. each day.



Enabling Subnet Utilization Collection

To capture subnet utilization data:

-
- Step 1** Configure the local cluster DHCP server with scopes and address ranges so that clients have requested leases.
- Step 2** Explicitly enable subnet utilization collection. The DHCP server attributes to set are:
- *collect-addr-util-duration*—Maximum period the DHCP server maintains data. You must change this from the default of 0 (no collection) to some reasonable value (see the context sensitive help for this attribute for the impact on memory).
-  **Note** If you are configuring simple failover, disable individual polling of the main and backup DHCP servers. Instead, enable the failover pair polling by setting the failover pair attribute *poll-subnet-util-interval*, so as to collect one set of data from both servers.
-
- *collect-addr-util-interval*—Frequency the server collects snapshots of the data (set to 15 minutes by default). How you juggle this value with that of the *collect-addr-util-duration* attribute determines how much memory you use (see the context sensitive help for this attribute).
- Step 3** Reload the local cluster DHCP server.
- Step 4** On the regional cluster, create the cluster that includes this DHCP server.
- Step 5** Go to the Subnet Utilization Settings section of the Add Server Cluster or Edit Server Cluster page.
- Step 6** Set the attributes in [Table 5-3](#).
- Step 7** Click **Modify Cluster**.
- Step 8** On the List Server Clusters page, click the Replica icon () next to the cluster name.
- Step 9** Click the Poll Subnet Utilization icon () for the cluster to obtain the initial set of subnet utilization data. This data is refreshed automatically at each polling interval. Note that if you subsequently click the Poll Subnet Utilization icon, new subnet utilization data does not appear until after the next collection interval (*collect-addr-util-interval*) on the DHCP server (15 minutes by default).
-

Enabling Lease History Collection


To capture lease history data:

-
- Step 1** Configure the local cluster DHCP server with scopes and address ranges so that clients have requested leases.
- Step 2** Explicitly enable lease history data collection. The DHCP server attributes to set are:
- *ip-history*—Set this to enabled.
 - *ip-history-detail*—Set this to enabled if you want to collect detailed history data.
 - *ip-history-max-age*—Limit on the age of the history records (set to 4 weeks by default).
- Step 3** Reload the local cluster DHCP server.
- Step 4** On the regional cluster, create the cluster that includes this DHCP server.
- Step 5** Go to the Lease History Settings section of the Add Server Cluster or Edit Server Cluster page.
- Step 6** Set the attributes in [Table 5-3 on page 5-8](#).
- Step 7** Click **Modify Cluster**.
- Step 8** On the List Server Clusters page, click the Replica icon () next to the cluster name.
- Step 9** On the same page, click the Poll Lease History icon () for the cluster involved to obtain the initial set of lease history data. This data is refreshed automatically at each polling interval.
-

Setting Up Routers and Router Interfaces

The regional Router Interface Configuration (RIC) server is used to manage router interfaces on Cisco Systems Universal Broadband Routers (uBRs) that manage cable modem termination systems (CMTSs). This module interacts with the CMTS servers to push the required cable modem configuration to edge devices, such as routers and switches (see [Figure 5-3](#)). The RIC server module is accessible only if you entered the router license, and by administrators assigned the ric-management subrole of the central-cfg-admin or regional-addr-admin role.

Figure 5-3 Router Interface Configuration (RIC) Server Module



Note

Add routers before you add any other subnets. This avoids the subnets that the router creates from possibly overlapping with those explicitly added, and avoids router synchronization errors later on.

The view of the routers is available on the View Tree of Routers page (see Figure 5-4).

Figure 5-4 View Tree of Routers Page

Home Administration Clusters Routers Replica Data DHCP Configuration DNS Configuration Address Space					
Router Tree Router List					
View Tree of Routers					
Name	Type	Address	Primary Subnet	Description	
0 1 2 All					
router-1	Router	10.86.144.138			
Cable2/0	Router Interface		2.2.2.2/24		I'm the best Interface Ever made.
Cable2/0.1	Router Interface				
Cable3/0	Router Interface				I'm nothing compared to Cable2/0
Ethernet1/0	Router Interface		10.86.144.138/26		I just don't get it
Ethernet1/1	Router Interface				Free Thinking Interface
Ethernet1/2	Router Interface				
Ethernet1/3	Router Interface				Why do I have to be Ethernet1/3
FastEthernet0/0	Router Interface		3.3.0.1/16		
				10	Change Page Size

The tree levels are the routers, their router interfaces, and any child interfaces. Parent/child relationships can be either physical/virtual (as in Cable2/0 and Cable2/0.1) or primary/secondary (as with router interface bundling, where the bundle is identified by one of the interfaces—see the “[Bundling Interfaces](#)” section on page 5-14). This listing of router interfaces is available only once you create routers in the system and synchronize with them.

Adding Routers

The routers managed by the RIC server can be one of the Cisco Universal Broadband Routers in the family uBR 72xx or uBR 10xxx.

- Step 1** To add a router, click **Routers** on the Primary Navigation bar, then **Router List** on the Secondary Navigation bar. This opens the List Routers page (see [Figure 5-5](#)).

Figure 5-5 List Routers Page

Name	IP Address	Description	Resynchronize	Interfaces
router-1	10.86.144.138	null		

Add Router

111491

- Step 2** Click **Add Router**. This opens the Add Router page (see [Figure 4-19 on page 4-28](#)).
- Step 3** The Add Router page requires that you enter at least the router's IP address and type. The literal entries for the Type field are **Ubr72xx** or **Ubr10k**. You also need to check with the router administrator about the username, password, and enable password, and enter these values.
- Step 4** Click **Add Router**.

Secure Mode Connections with Routers

To enable secure communication between the RIC server and the routers, you must have the Cisco CNS Network Registrar Communications Security Option Release 1.1 installed. By default, secure connectivity is disabled and it is through Telnet. However, you can specify whether you want to require or desire an SSH connection. Use the *use-ssh* attribute in the (expandable) Reserved Attributes section of the page. This attribute has the values:

- disabled—The default. This uses simple Telnet for the connection.
- required—The router communicates with the edge device using SSH only, and not Telnet.
- desired—The router tries to communicate using SSH, but if it cannot, it uses Telnet.

Alternative Login Method to Routers

There are two types of login mechanisms provided in the RIC server that you can affect using the *login-template* attribute on the Add Router page:

- Discovery mode—The default mechanism, designed to understand login prompts on edge devices and respond to those dynamically. It does not force a particular login sequence, but supports the various login sequences and login prompts most customers use with these default prompts:


```
Username prompt - Username:
Password prompt - Password:
Login-prompt - >
Enable password prompt - Password:
Enable prompt - #
```

- Template mode—Use this in case the RIC server cannot log in using the discovery mechanism for some reason, such as with nonstandard prompts or a login sequence that the discovery mechanism does not understand. The *login-template* is the name of an optional login template to use to further customize the RIC server's login and enable interactive sessions. To create this template you must:
 - a. In the API, create an `ScpObj` of class `CCMRouterLoginTemplate`.
 - b. Add the object to the database using the `RICAdminSession.addRouterLoginTemplate` method.
 - c. Enter the name of the added template (`CCMRouterLoginTemplate.name`) as the value of the *login-template*.


Editing Routers

To edit a router, click its name on the View Tree of Routers page or List Routers page. The Edit Router page is essentially the same as the Add Router page, except for an additional attribute unset function. Make your changes, then click **Modify Router**.

Resynchronizing Routers

As soon as you add the router to the regional cluster, it is synchronized over the network. You can also explicitly resynchronize the router if you know that changes occurred. On the List Routers page, click the Resynchronize icon () next to the router name. If the synchronization could not occur or timed out, you get an error message to that effect.

Viewing and Editing the Router Interfaces

If you click the Interface icon () associated with the router on the List Routers page, the list of related cable or Ethernet interfaces appears on the List Router Interfaces page. Both from this page and the View Tree of Routers page, you can click the interface name to edit it. The only difference is an additional attribute unset function and that you can delete the interface on the List Router Interfaces page.

Changeable Router Interface Attributes

Editing the router interface opens the Edit Router Interface page. You cannot change the name, state, or MAC address of the interface on this page. However, you can change the following attributes:

- Description
- Address of the primary subnet address on the interface
- Addresses of the secondary subnets on the interface
- Address of any IP helper (DHCP relay agent) for the interface
- Address of any cable helper of the DHCP server to accept unicast packets for the interface
- Owner of the router
- Region of the router

Bundling Interfaces

An interface bundle provides load balancing among the router interfaces. When you define a bundle, all the participating interfaces in the bundle must have the same bundle identifier (ID), which is the name of the interface specified as the master.

If you want to use bundling, the attributes are in the Interface Bundling Settings section of the Edit Router Interface page:

- *bundle-id*—Interface bundle identifier, the name of the master interface. All participating interfaces in the bundle must have the same bundle ID.
- *is-master*—This interface is the master interface in the bundle.
- *is-virtual*—This interface is a virtual interface in the bundle.

Creating Virtual Private Networks

A virtual private network (VPN) is a specialized address space identified by a key (see also “[Configuring Virtual Private Networks Through DHCP](#)” section on page 19-1). A VPN allows address overlap in a network, because the addresses are distinguished by separate keys. Most IP addresses exist in the “global” address space outside of a VPN. You can create regional VPNs only if you are an administrator assigned the dhcp-management subrole of the central-cfg-admin role.

VPNs that you create provide a filtering mechanism for:

- Viewing the unified address space (see “[Viewing Unified Address Space](#)” section on page 18-2)
- Listing address blocks (see “[Adding Address Blocks](#)” section on page 18-5)
- Listing subnets (see “[Address Blocks and Subnets](#)” section on page 18-4)
- Querying subnet utilization (see “[Generating Subnet Utilization History Reports](#)” section on page 18-13)
- Querying lease history (see “[Running IP Lease Histories](#)” section on page 12-13)

To set up VPNs on the regional cluster, click **DHCP Configuration** on the Primary Navigation bar, then **VPNs** on the Secondary Navigation bar. This opens the List/Add VPNs page (see [Figure 5-6](#)).

Figure 5-6 List/Add VPNs Page

The screenshot shows the 'List/Add VPNs' page. At the top, there is a navigation bar with the following items: Home, Administration, Clusters, Routers, Replica Data, **DHCP Configuration**, DNS Configuration, Address Space, and a secondary navigation bar with VPNs, Scope Templates, Policies, Client Classes, and Failover. The main content area is titled 'List/Add VPNs' and contains a table with the following columns: Key*, Name*, VPN Id, VRF Name, Description, and Push Data. The table has one row with the following data: Key: 999, Name: vpn-1, VPN Id: a1:3f6c, VRF Name: null, Description: null, and a 'Push VPN' button. Below the table are three buttons: 'Add VPN', 'Pull Replica VPNs', and 'Push All VPNs'. The page number '111490' is visible in the bottom right corner.

Key*	Name*	VPN Id	VRF Name	Description	Push Data
999	vpn-1	a1:3f6c	null	null	Push VPN

To create a VPN on this page, enter the VPN key and its name, and you can also enter a description. The key must be a unique integer value. The name must also be unique across the regional cluster. You must also enter either of two values, depending on the configuration:

- **VPN ID**—VPN identifier, in the colon-separated hexadecimal format *oui:index*, per RFC 2685. The *oui* is the Organizationally Unique Identifier (OUI) that corresponds to the VPN owner or service provider. The *index* is the four-octet index number of the VPN. For example, **a1:3f6c**.
- **VRF Name**—VPN Routing and Forwarding instance (VRF) name. Cisco routers commonly use VRF names.

Editing VPNs

To edit a VPN, click its name on the List/Add VPNs page (see [Figure 5-6](#)). The Edit VPN page is essentially the same as the Add VPN page. Make your changes, then click **Modify VPN**.

Pushing VPNs to Local Clusters

You can also push the VPNs you create from the regional cluster to any of the local clusters. If you want to push a specific VPN to a cluster, click **Push VPN** on the List/Add VPNs page. If you want to push all of them, click **Push All VPNs**. Both open the Push VPN Data to Local Clusters page (see [Figure 5-7](#)).

Figure 5-7 Push VPN Data to Local Clusters Page

This page identifies the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure**—The default: Ensures that the local cluster has new data without affecting any existing data.
- **Replace**—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for “push all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field.



Tip

The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

After making these choices, click **Push Data to Clusters**. This opens the View Push VPN Data Report page.

Pulling VPNs from Replica Data

Instead of explicitly creating VPNs, you can pull them from the local clusters. (You may first want to update the VPN replica data by clicking the Replica icon [🔄] next to the cluster name.) To pull the replica data, click **Pull Replica VPNs** to open the Select Replica VPN Data to Pull page (see Figure 5-8).

Figure 5-8 Select Replica VPN Data to Pull Page

Name	Update Replica Data	Pull Data
Boston-cluster	[🔄]	Pull All VPNs Mode: <input type="radio"/> Ensure <input checked="" type="radio"/> Replace <input type="radio"/> Exact
Chicago-cluster	[🔄]	Pull All VPNs Mode: <input type="radio"/> Ensure <input checked="" type="radio"/> Replace <input type="radio"/> Exact

Cancel

10 Change Page Size

This page shows a tree view of the regional server’s replica data for the local clusters’ VPNs. The tree has two levels, one for the local clusters and one for the VPNs in each cluster. You can pull individual VPNs or you can pull all of them. To pull individual VPNs, expand the tree for the cluster, then click **Pull VPN** next to its name. To pull all the VPNs, click **Pull All VPNs from Cluster**. To pull the VPNs, you must choose a synchronization mode:

- Ensure—Ensures that the regional cluster has new data without affecting any existing data.
- Replace—The default: replaces data without affecting other objects unique to the regional cluster.
- Exact—Available for “pull all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the regional cluster.

Creating DHCP Scope Templates

Scope templates apply certain common attributes to multiple scopes. These common attributes include a scope name based on an expression, policies, address ranges, and an embedded policy options based on an expression. The scope templates you add or pull from the local clusters are visible on the List DHCP Scope Templates page.

You get to this page by clicking **DHCP Configuration** on the Primary Navigation bar, then **Scope Templates** on the Secondary Navigation bar. This functionality is available only to administrators assigned the dhcp-management subrole of the central-cfg-admin role.

To explicitly create a scope template, click **Add Scope Template** on this page. This opens the Add DHCP Scope Template page, which includes a number of fields and settings (see [Figure 4-21 on page 4-31](#)). You must give the template at least a name. The following sections describe other settings.

Using Expressions

There are three fields on the Add DHCP Scope Template page for which you must specify an expression:

- Scope name—This must return a string.
- Address range—This must return IP addresses.
- Embedded policy options.

The expressions you can use are documented in a link from the Web UI online help for this page.

Scope Name Expression Example

You might want to set an expression so that the template constructs scope names starting with “ISP–” and followed by the subnet of the scope and a derivative of its ping time-out value. You would use the following expression in the Scope Name Expression field:

```
(concat "ISP-" subnet "-" (+ template.ping-timeout 10))
```

The elements of the expression are:

- **(concat ...)**—Concatenation operation, which concatenates all the following values into one value.
- **“ISP–”**—String with which to start the scope name.
- **subnet**—Keyword variable that indicates to use the existing subnet defined for the scope.
- **“–”**—Indicates to include this hyphen to construct the value.
- **(+ template.ping-timeout 10)**—Indicates to add the *ping-timeout* property value for the scope to the number 10.

If the scope’s subnet happens to be 192.168.50.0/24 and its *ping time-out* value 100, the resulting constructed scope name would be:

```
ISP-192.168.50.0/24-110
```

Range Expression Example

You might want to set an expression so that the template constructs only certain addresses ranges for scopes. You can either be explicit about the actual starting and ending addresses, or you can make them relative to the subnet. Here are two ways of requesting relative ranges in the Range Expression field:

```
(create-range first-addr last-addr)  
(create-range 1 10)
```

The first **create-range** operation creates the address range based on the first through last usable address in the subnet. For the 192.168.50.0/24 subnet, for example, the address range would be 192.168.50.1 through 192.168.50.254. Because the second operation specifies integers instead of full IP addresses, it makes the range relative to the subnet based on its mask. If the template discovers the subnet to be 192.168.50.0/26, it takes the first through tenth address in this subnet, which would be 192.168.50.65 through 192.168.50.74.

Embedded Policy Option Expression Example

An embedded policy is important because the DHCP server looks at it before it looks at the scope's assigned, named policy. This is usually where you would set the DHCP options on a scope. You might want to set an expression so that the template constructs DHCP options for the scope's embedded policy. Here are some examples:

```
(create-option "domain-name" "example.com")
(create-option 3 "10.10.10.1")
(create-option "routers" (create-ipaddr subnet 10))
```

The first **create-option** operation associates the value `example.com` with the *domain-name* option for the scope. The second operation associates the address `10.10.10.1` with the *routers* option (number 3). The third operation creates an IP address for the *routers* option based on the tenth address in a subnet.

Additional Scope Template Attributes

The optional additional attributes appear in functional categories. For a description of each attribute, click the attribute name to open a help window. For example, you might want to enable dynamic DNS updates for the scope, or set the main and backup DHCP failover servers.

After you complete these fields, click **Add Scope Template**.

Editing Scope Templates

To edit a scope template, click its name on the List DHCP Scope Templates page. The Edit DHCP Scope Template page is essentially the same as the Add DHCP Scope Template page, except for an additional attribute unset function. Make your changes, then click **Modify Scope Template**.

Pushing Scope Templates to Local Clusters

You can also push the scope templates you create from the regional cluster to any of the local clusters. If you want to push a specific template to a cluster, click **Push Scope Template** on the List DHCP Scope Templates page. If you want to push all of them, click **Push All Scope Templates**. Both open the Push Scope Template Data to Local Clusters page (see [Figure 5-9](#)).

Figure 5-9 Push Scope Template Data to Local Clusters Page

This page identifies the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure**—The default: Ensures that the local cluster has new data without affecting any existing data.
- **Replace**—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for “push all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field.


Tip

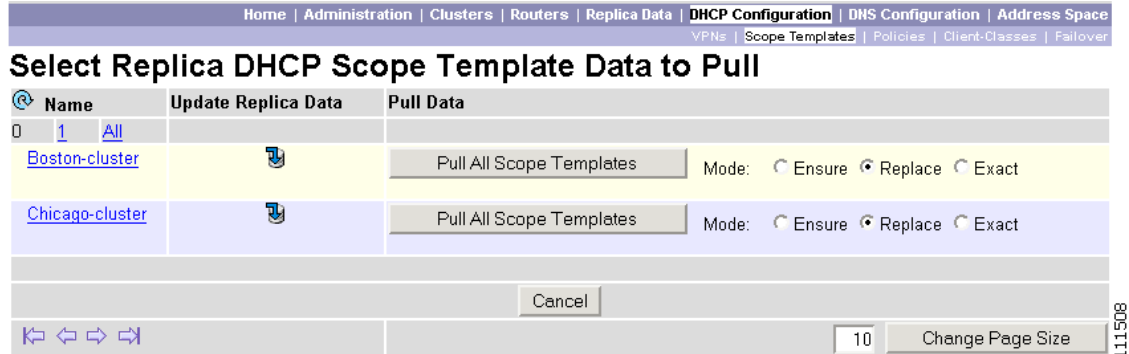
The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

After making these choices, click **Push Data to Clusters**. This opens the View Push Scope Template Data Report page.

Pulling Scope Templates from Replica Data

You may choose to pull scope templates from the replica data of the local clusters instead of explicitly creating them. (You may first want to update the policy replica data by clicking the Replica icon [🔄] next to the cluster name.) To pull the scope templates, click **Pull Replica Scope Templates** to open the Select Replica DHCP Scope Template Data to Pull page (see Figure 5-10).

Figure 5-10 Select Replica DHCP Scope Template Data to Pull Page



This page shows a tree view of the regional server's replica data for the local clusters' scope templates. The tree has two levels, one for the local clusters and one for the scope templates in each cluster. You can pull individual scope templates from the clusters, or you can pull all of their scope templates. To pull individual scope templates, expand the tree for the cluster, then click **Pull Policy** next to its name. To pull all the scope templates from a cluster, click **Pull All Scope Templates from Cluster**. To pull the scope templates, you must also choose a synchronization mode:

- Ensure—Ensures that the regional cluster has new data without affecting any existing data.
- Replace—The default: replaces data without affecting other objects unique to the regional cluster.
- Exact—Available for “pull all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the regional cluster.

Creating DHCP Policies

Every DHCP server must have one or more policies defined for it. Policies define lease duration, gateway routers, and other configuration parameters, in what are called DHCP options. Policies are especially useful if you have multiple scopes, because you need only define a policy once and apply it to the multiple scopes.

The policies you add or pull from the local clusters are visible on the List DHCP Policies page. The policies are listed alphabetically, with their offer time-out and grace period values displayed. The default and system_default_policy policies appear by default.

You get to this page by clicking **DHCP Configuration** on the Primary Navigation bar, then **Policies** on the Secondary Navigation bar to open the List DHCP Policies page. This functionality is available only to administrators who are assigned the dhcp-management subrole of the central-cfg-admin role.

To explicitly create a policy, click **Add Policy** on the List DHCP Policies page. This opens the Add DHCP Policy page, which includes a number of fields and settings. You must give the policy at least a name.

The other important properties of a policy are its offer time-out and grace period. By default, the DHCP server stops trying to offer leases to the scope after two minutes. Also, by default, the server gives a grace period of five minutes after a lease expires, during which time the lease becomes available for re-assignment (this must be set in a policy, because it is not available for an embedded policy). You can adjust these two values on this page (see Figure 5-11).

Figure 5-11 Add DHCP Policy Page

Attribute	Value
Name*	central-policy-1
Offer timeout	2m
Grace period	5m

Options	Number	Value
	[51] dhcp-lease-time (unsigned time)	2w

Attribute	Value	Data Type	Default
bootp-reply-options		list	
dhcp-reply-options		list	

111476

As with embedded policies, you can also set DHCP options for named policies. In this case, you can choose an option from the drop-down list, then associate a value with it. There are additional attributes you can set, such as allowing a lease time override.

When you are finished, click **Add Policy** to add the policy.

Editing Policies

To edit a policy, click its name on the List DHCP Policies page. The Edit DHCP Policy page is essentially the same as the Add DHCP Policy page, except for an additional attribute unset function. Make your changes, then click **Modify Policy**.

Pushing Policies to Local Clusters

You can also push the policies you create from the regional cluster to any of the local clusters. If you want to push a specific policy to a cluster, click **Push Policy** on the List DHCP Policies page. If you want to push all of them, click **Push All Policies**. Both actions open the Push DHCP Policy Data to Local Clusters page.

This page identifies the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure**—The default: Ensures that the local cluster has new data without affecting any existing data.
- **Replace**—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for “push all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field.




Tip

The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

After making these choices, click **Push Data to Clusters**. This opens the View Push Policy Data Report page.

Pulling Policies from Replica Data

You may choose to pull policies from the replica data of the local clusters instead of explicitly creating them. (You may first want to update the policy replica data by clicking the Replica icon [] next to the cluster name.) To pull the policies, click **Pull Replica Policies** to open the Select Replica DHCP Policy Data to Pull page.

This page shows a tree view of the regional server's replica data for the local clusters' policies. The tree has two levels, one for the local clusters and one for the policies in each cluster. You can pull individual policies from the clusters, or you can pull all of their policies. To pull individual policies, expand the tree for the cluster, then click **Pull Policy** next to its name. To pull all the policies from a cluster, click **Pull All Policies from Cluster**. To pull all the policies, you must also choose a synchronization mode:

- **Ensure**—Ensures that the regional cluster has new data without affecting any existing data.
- **Replace**—The default: replaces data without affecting other objects unique to the regional cluster.
- **Exact**—Available for “pull all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the regional cluster.

Creating DHCP Client-Classes

Client-classes 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:

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

You can enable or disable client-class processing for the DHCP server and apply a set of properties to groups of clients. With client-class processing enabled, the DHCP server assigns the client to an address from a matching scope. The server acts according to the client and client-class data in each packet. To configure client-class, enable client-class processing for the DHCP server, then assign clients to these classes based on selection criteria.

You can view any created client-classes on the List DHCP Client-Classes page. You get to this page by clicking **DHCP Configuration** on the Primary Navigation bar, then **Client-Classes** on the Secondary Navigation bar. This functionality is available only to administrators who are assigned the dhcp-management subrole of the central-cfg-admin role.

To explicitly create a policy, click **Add Client-Class** on this page. This opens the Add DHCP Client-Class page, which includes a number of fields and settings (see [Figure 5-12](#)). You must give the client-class at least a name.

Figure 5-12 Add DHCP Client-Class Page

Attribute	Value	Data Type	Default
Name*	example-client-class		
Host name	clienthost-1		
Domain name	example.com.		
Policy name	central-policy-1		
selection-criteria		list	
selection-criteria-excluded		list	

111473

The important properties of a policy are its host name and domain name. You can also associate a DHCP policy with the client-class, and you can set additional attributes. For example, you can set selection criteria for a scope to be acceptable to the client-class. You can also set a limitation ID to limit the number of IP addresses the DHCP server should give out to devices on customer premises. (For details on limitation IDs, see the “[Subscriber Limitation Using Option 82](#)” section on page 13-10.)

When you are finished, click **Add Client-Class** to add the client-class.

Editing Client-Classes

To edit a client-class, click its name on the List DHCP Client-Classes page. The Edit DHCP Client-Class page is essentially the same as the Add DHCP Client-Class page, except for an additional attribute unset function. Make your changes, then click **Modify Client-Class**.

Pushing Client-Classes to Local Clusters

You can also push the client-classes you create from the regional cluster to any of the local clusters. If you want to push a specific client-class to a cluster, click **Push Client-Class** on the List DHCP Client-Classes page. If you want to push all of them, click **Push All Client-Classes**. Both open the Push Client-Class Data to Local Clusters page (see [Figure 5-13](#)).

Figure 5-13 Push DHCP Client-Class Data to Local Clusters Page

This page identifies the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure**—The default: Ensures that the local cluster has new data without affecting any existing data.
- **Replace**—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for “push all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field.



Tip

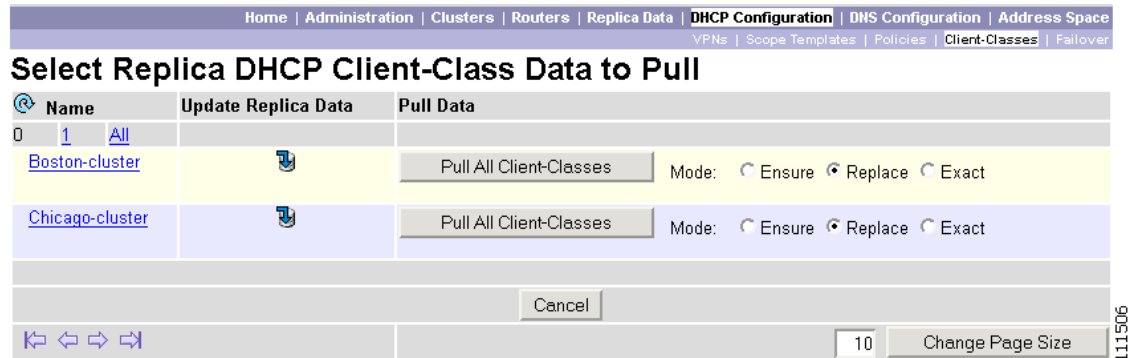
The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

After making these choices, click **Push Data to Clusters**. This opens the View Push Client-Class Data Report page.

Pulling Client-Classes from Replica Data

You may choose to pull client-classes from the replica data of the local clusters instead of explicitly creating them. (You may first want to update the client-class replica data by clicking the Replica icon [🔄] next to the cluster name.) To pull the client-classes, click **Pull Replica Client-Classes** to open the Select Replica DHCP Client-Class Data to Pull page (see [Figure 5-14](#)).

Figure 5-14 Select Replica DHCP Client-Class Data to Pull Page



This page shows a tree view of the regional server's replica data for the local clusters' client-classes. The tree has two levels, one for the local clusters and one for the client-classes in each cluster. You can pull individual client-classes from the clusters, or you can pull all of their client-classes. To pull individual client-classes, expand the tree for the cluster, then click **Pull Client-Class** next to its name. To pull all the client-classes from a cluster, click **Pull All Client-Classes from Cluster**. To pull the client-classes, you must also choose a synchronization mode:

- **Ensure**—Ensures that the regional cluster has new data without affecting any existing data.
- **Replace**—The default: replaces data without affecting other objects unique to the regional cluster.
- **Exact**—Available for “pull all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the regional cluster.

Creating DHCP Failover Pairs

DHCP failover allows a backup DHCP server to take over for a main server if the main server is taken off the network for any reason. You can use failover to configure two DHCP servers to operate as a redundant pair. If one server is down, the other server seamlessly takes over so that new DHCP clients can get, and existing clients can renew, their addresses. Clients requesting new leases need not know or care about which server is responding to their request for a lease. These clients can obtain leases even if the main server is down.

You can view any created failover pairs on the List Failover Pairs page. To access this page, click **DHCP Configuration** on the Primary Navigation bar, then **Failover** on the Secondary Navigation bar. This functionality is available only to administrators who are assigned the dhcp-management subrole of the central-cfg-admin role.

You can create a failover pair in two ways:

- Pull the address space from the replica data.
- Add the failover pair.


Adding Failover Pairs

To explicitly create a failover pair on the regional cluster, click **Add Failover Pair** on the List Failover Pairs page. This opens the Add Failover Pair page, which includes a number of fields and settings (see [Figure 4-22 on page 4-32](#)). You must give the failover pair a name, and choose a main and backup server

for it. Specifying a scope template and subnet is optional, unless you want to use the **Use Regional Subnets** option when synchronizing the failover pair (see the “[Synchronizing Failover Pairs with Local Clusters](#)” section).



The important properties of a failover pair are the local clusters where the main and backup DHCP servers reside. These local clusters must already be created on the regional cluster, and you choose them from a drop-down list. You can also choose a scope template for the failover pair, and you can choose the subnet or subnets for the ensuing leases.

You can also set additional attributes, such as for polling and replication, or to propagate scopes. When you are finished, click **Add Failover Pair** to add the failover pair.

After you add the failover pair, you can get the related server information for each server. On the List Failover Pairs page, click the Related Servers icon () next to the main or backup DHCP server. This opens the List Related Servers for DHCP Server page (see the “[Listing Related Servers for DHCP Servers](#)” section on page 5-3).

Synchronizing Failover Pairs with Local Clusters

You must synchronize the created failover pair with the local cluster DHCP servers. You can run a report of impending changes before you run the synchronization, or you can synchronize right away.

-
- Step 1** On the List Failover Pairs page, look at the icons in the Synchronize column:
- To set up the synchronization criteria and view a report first, click the Report icon (). This opens the Report Synchronize Failover Pair page.
 - To set up the synchronization criteria and run the synchronization right away, click the Run icon (). This opens the Run Synchronize Failover Pair page.
- Step 2** On the Report or Run page, choose what to use for synchronization, the regional subnets or the local scopes. In the Synchronize Subnets section, choose the appropriate radio button:
- Use Regional Subnets**—If there is a subnet and scope template selected for the failover pair in the regional cluster, use this subnet.
 - Use Local DHCP Scopes**—Use the scopes on the local cluster DHCP servers.
- Step 3** Click the radio button in the Synchronization Direction section corresponding to the direction you want to synchronize the servers:
- Main to Backup**
 - Backup to Main**
- Step 4** Click the radio button in the Operation section corresponding to the degree to which you want to synchronize the failover data:
- Update**—This is the default and least radical operation. It is appropriate for update synchronizations in that it has the least effect on the unique properties of the backup server.
 - Complete**—This operation is appropriate for all initial synchronizations. It is more complete than an update operation, while still preserving many of the backup server’s unique properties.
 - Exact**—This operation is appropriate for initial basic and symmetrical failover configurations. It makes the two servers as much as possible mirror images of each other, although it retains unique DHCP server attributes, LDAP event services, and extension points on the backup server.

Each operation performs a different mix of functions on the failover properties, as described in Table 5-4. There are four functions, with examples based on these property name-value pairs:

On the main server:	On the backup server:
Name1=A	Name2=B
Name2=C	Name3=D

- **no change**—Makes no change to the list of properties or their values on the backup server. For the example, the result would be Name2=B, Name3=D.
- **ensure**—Ensures that a copy of the main server property exists on the backup server, but does not replace its value. For the example, the result would be Name1=A, Name2=B, Name3=D.
- **replace**—Replaces the value of a property that the two servers have in common with that of the main server. For the example, the result would be Name1=A, Name2=C, Name3=D.
- **exact**—Puts an exact copy of the main server's list of properties and values on the backup server and removes the unique ones. For the example, the result would be Name1=A, Name2=C.

Table 5-4 Failover Synchronization Functions

Data Description	Update	Complete	Exact
DHCP Server (server level failover pair):	replace	replace	replace
Client Class Properties			
Failover Properties			
Failover Tuning Properties			
Dynamic DNS Security Properties			
(See the Web UI online help for the full list of properties affected.)			
All other Properties	no change	replace	replace
LDAP Event Service	no change	replace	replace
Policy:			
Option-list Property	ensure	replace	exact
All other Properties	replace	replace	exact
Client	replace	replace	exact
Client-Class	replace	replace	exact
Scopes (related to failover pair)	exact	exact	exact
VPN	replace	replace	exact
Key	replace	replace	exact
Extensions	ensure	replace	exact
Note You must manually copy over the extension files.			
Extension Point	no change	replace	replace
Option Information:	ensure	exact	exact
Custom options list			
Vendor options list			
Option-Data-types list			

- Step 5** If you are running a report, click **Report** to show a preview of the synchronized data, then click **Run Update** on the View Failover Pair Sync Report page. If you synchronizing right away, click **Run**. In both cases, you implement the changes and return to the View Failover Pair Sync Report page.







The parts of the report are:

- What changes were made on the main DHCP server
- What changes were made on the backup DHCP server
- What change set entries were made on the regional cluster

- Step 6** Click **Return to Failover Pair List**.
-

Restarting the Failover Servers

For any failover synchronization to take effect, you must first connect to, and restart, both the main and backup failover servers.

- Step 1** On the List Failover Pairs page, click the Go Local icon () in the Main DHCP Server column.
- Step 2** On the local Manage DHCP Server page for the main server, click the Reload icon () on the right hand side of the page.
- Step 3** Click the Go Regional icon () at the top right corner of the page.
- Step 4** On the regional List Failover Pairs page, click the Go Local icon () in the Backup DHCP Server column.
- Step 5** On the local Manage DHCP Server page for the backup server, click the Reload icon () on the right hand side of the page.
- Step 6** Click the Go Regional icon () at the top right corner of the page.
-

Editing Failover Pairs

The Edit Failover Pair page you access when you click the failover pair name on the List Failover Pairs page is essentially the same as the Add Failover Pair page (see [Figure 4-22 on page 4-32](#)). The difference is that the associated subnets may be different than what was originally entered, because the failover pair may have been synchronized with the Use Local DHCP Scopes mode selected (see the [“Synchronizing Failover Pairs with Local Clusters”](#) section on page 5-26). After you modify the failover pair on this page, click **Modify Failover Pair**.

Pulling Address Space from Replica Data

To pull the address space for a failover pair, you must have the address space license and regional-addr-admin privileges.

- Step 1** On the List Failover Pairs page or View Unified Address Space page, click **Pull Replica Address Space**.

- Step 2** Choose the data synchronization mode (**Update**, **Complete**, or **Exact**) on the Select Pull Replica Address Space page. The results of choosing these modes are described in the table on the page.
- Step 3** Click **Report** at the bottom of the page.
- Step 4** Click **Run** on the Report Pull Replica Address Space page.
- Step 5** Click **OK** on the Run Pull Replica Address Space page.

Creating DNS Zone Distributions

Creating a zone distribution map simplifies creating multiple zones that share the same secondary DNS server attributes. Like a scope template, the zone distribution map can have a unique name. The distribution map requires adding one or more predefined secondary servers. When you run a zone distribution synchronization, this creates secondary zones based on the primary zone.

In Network Registrar 6.0, you could manage only the default distribution. In Network Registrar 6.1, you can define additional ones at the regional cluster. The distribution must be in a star topology, that is, one authoritative server and multiple secondary servers. The authoritative server can only be the local primary DNS server where the zone distribution default is defined.

To create zone distributions, click **DNS Configuration** on the Primary Navigation bar, then **Zone Distributions** on the Secondary Navigation bar. This opens the List/Add Zone Distributions page (see [Figure 5-15](#)).

Figure 5-15 List/Add Zone Distributions Page

Name	DNS Primary Server	Synchronize
Default	local	Run Report
distr-1	Boston-cluster	Run Report

Buttons: Add Zone Distribution, Synchronize All Zone Distributions (Mode: Replace, Exact), Pull Replica Zone Data

116768

Adding Zone Distributions

To add a zone distribution, click **Add Zone Distribution** on the List/Add Zone Distributions page. This opens the Add Zone Distribution page (see [Figure 5-16](#)).

Figure 5-16 Add Zone Distribution Page

Home | Administration | Clusters | Routers | Replica Data | DHCP Configuration | **DNS Configuration** | Address Space

Zone Distributions | Forward Zones Tree | Forward Zones List | Reverse Zones Tree | Reverse Zones List | Zone Templates

Add Zone Distribution

Attribute	Value
Name	example-zone-distrib
Primary Server	Boston-cluster
DNS Primary Server IP Addresses	<input type="text"/> Add IP Key
Secondary Servers	Add Server
Forward Zones	
Selected [none]	Available testrest
<< >> Select All	
<input type="text"/> Search	<input type="text"/> Search
Reverse Zones	

111479

On this page, enter the name of the zone distribution, the cluster for its primary server, and the address or addresses of the master DNS server or servers. To add each master server, enter its IP address with an optional TSIG key, in the format *address-key* (with the TSIG key preceded by a hyphen) in the Master Servers field, then click **Add IP Key**.

To add a secondary DNS server, click **Add Secondary Server** in the Secondary Servers section of the page. This opens the Add Secondary Server page. On that page, choose the cluster for the secondary server, then enter the address or addresses of any master servers. When you click **Add Secondary Server** to return to the Add Zone Distribution page, you can sign on to the cluster to effect changes there, if necessary. You can also do this from the List/Add Zone Distributions page.

To add the forward and zones for the zone distribution, move them to the appropriate Selected field.

To finish adding the zone distribution, click **Add Zone Distribution**.

Editing Zone Distributions

To edit a zone distribution, click its name on the List/Add Zone Distributions page. The Edit Zone Distribution page is essentially the same as the Add Zone Distribution page. Make your changes, then click **Modify Zone Distribution**.

Synchronizing Zone Distributions with Local Clusters

You can synchronize the zone distribution with the local cluster DNS servers by clicking the Run icon (➕) next to the zone distribution name on the List/Add Zone Distributions page. This opens the Sync Zone Distribution page, which shows the actual data synchronized. If you click the Report icon (📄) in the Synchronize column, this opens the same page, except that it shows a preview of the data synchronized, and you have to click **Run** to activate it.

You can also synchronize all the created zone distributions by choosing a synchronization mode, then clicking the Run icon (➕) or Report icon (📄) in the Synchronize All Zone Distributions section of the List/Add Zone Distributions page. The synchronization modes are:

- **Replace**—Replaces the zone distribution data for those with the same name, but retains all others.
- **Exact**—Replaces the zone distribution data for those with the same name, and removes all others.

After synchronizing, click the Go Local icon (🏠➡) next to the zone distribution name to single sign-on to the local cluster. On the Manage DNS Server page, reload the server, then click the Go Regional icon (🏠←) at the top right corner of the page to return to the regional cluster pages. If you do not have single sign-on privileges, ask the local DNS administrators to reload the DNS server.

Listing Forward and Reverse Zones

You can list the forward and reverse zones for a zone distribution, once synchronized. You get there by clicking **DNS Management** on the Primary Navigation bar, then one of the following tabs on the Secondary Navigation bar:

- **Forward Zones Tree**—Opens the View Forward Zones Tree page. This page shows the hierarchy of forward zones. Note that the page control relates to the total number of nodes displayed, not just the top nodes in the tree. The tree always shows the related hierarchy in the yellow section of the table when you navigate to subsequent pages. You can also navigate using the small up or down arrows next to each zone level in these yellow areas.

This page lets you edit the zone distribution to which each zone belongs, and pull the replica zone data. To edit the zone attributes themselves, you must connect to the local cluster.

- **Forward Zones List**—Opens the List Forward Zones page.
- **Reverse Zones Tree**—Opens the View Reverse Zones Tree page. This page shows a hierarchy of reverse zones. The same navigation features exist as for the forward zone page.

This page lets you edit the zone distribution to which each reverse zone belongs, and pull the replica zone data. To edit the reverse zone attributes themselves, you must connect to the local cluster.

- **Reverse Zones List**—Opens the View Reverse Zones page.

See [Chapter 8, “Managing Zones,”](#) for details on managing zone data.

Pulling Zone Distributions from Replica Data

You may choose to pull zone distributions from the replica data of the local clusters instead of explicitly creating them. To pull the zone distributions:

-
- Step 1** On the List/Add Zone Distributions page, click **Pull Replica Zone Data**.
 - Step 2** Choose the data synchronization mode (**Update**, **Complete**, or **Exact**) on the Select Pull Replica Zone Data page. These modes are described in the table on the page.


- Step 3** Click **Report** at the bottom of the page.
- Step 4** Click **Run** on the Report Pull Replica Zone Data page.
- Step 5** Click **OK** on the Run Pull Replica Zone Data page.

Creating Zone Templates

A zone template is a convenient way to create a boilerplate for primary zones that share many of the same attributes. You can apply a zone template to any zone, and override the zone's attributes with those of the template.

To create zone templates, click **DNS Configuration** on the Primary Navigation bar, then **Zone Templates** on the Secondary Navigation bar. This opens the List Zone Templates page (see [Figure 5-17](#)).

Figure 5-17 List Zone Templates Page



Adding Zone Templates

To add a zone template, click **Add Zone Template** on the List Zone Templates page. This opens the Add Zone Template page. This page is almost identical to the Add Zone page for the local cluster (see [Figure 4-8](#) on page 4-18).

Enter the zone template name and at least the suggested serial number, nameserver, and contact E-mail address, because they are required for the zone itself. You might also want to specify any zone owners or zone distributions. The template name and zone default TTL are required. (For a description of the minimally required zone attributes, see the [“Creating Primary Zones”](#) section on page 8-4.)

Once you are done entering these value, click **Add Zone Template** at the bottom of the page.

Editing Zone Templates

Edit a zone template by clicking its name on the List Zone Templates page. The Edit Zone Template page is essentially the same as the Add Zone Template page. Make your changes, then click **Modify Zone Template**.

Pushing Zone Templates to Local Clusters

You can also push the zone templates you create from the regional cluster to any of the local clusters. If you want to push a specific zone template to a cluster, click **Push Zone Template** on the List DHCP Client-Classes page. If you want to push all of them, click **Push All Zone Templates**. Both open the Push Zone Template Data to Local Clusters page (see [Figure 5-18](#)).

Figure 5-18 Push Zone Template Data to Local Clusters Page

This page identifies the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure**—The default: Ensures that the local cluster has new data without affecting any existing data.
- **Replace**—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for “push all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field.



Tip

The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

After making these choices, click **Push Data to Clusters**. This opens the View Push Zone Template Data Report page.

Pulling Zone Templates from Replica Data

You may choose to pull zone templates from the replica data of the local clusters instead of explicitly creating them. To pull the zone templates:

- Step 1** On the List Zone Templates page, click **Pull Replica Zone Templates**.

- Step 2** Choose the data synchronization mode (**Update**, **Complete**, or **Exact**) on the Select Replica DNS Zone Template Data to Pull page (see [Figure 5-19](#)).

Figure 5-19 Select Replica DNS Zone Template Data to Pull Page

Name	Update Replica Data	Pull Data
Boston-cluster		Pull All Zone Templates Mode: <input type="radio"/> Ensure <input checked="" type="radio"/> Replace <input type="radio"/> Exact
Chicago-cluster		Pull All Zone Templates Mode: <input type="radio"/> Ensure <input checked="" type="radio"/> Replace <input type="radio"/> Exact

Cancel

10 Change Page Size

This page shows a tree view of the regional server’s replica data for the local clusters’ zone templates. The tree has two levels, one for the local clusters and one for the zone templates in each cluster. You can pull individual zone templates or you can pull all of them. To pull individual zone templates, expand the tree for the cluster, then click **Pull Zone Template** next to its name. To pull all the zone templates, click **Pull All Zone Templates from Cluster**. To pull the zone templates, you must choose a synchronization mode:

- **Ensure**—Ensures that the regional cluster has new data without affecting any existing data.
- **Replace**—The default: replaces data without affecting other objects unique to the regional cluster.
- **Exact**—Available for “pull all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the regional cluster.