



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, replicating their data, and polling subnet utilization and lease history data from them.
- Setting up routers (see [Chapter 10, “Managing Router Interface Configurations”](#)).
- Managing network objects such as DHCP scope templates, policies, client-classes, options, networks, and virtual private networks (VPNs).
- Managing DHCP failover server pairs.


These functions are available only to administrators assigned the central-cfg-admin role. (The full list of functions for the central-cfg-admin are listed in [Table 4-2 on page 4-3](#).) 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](#).

Configuring 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 at 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**. Once the page is populated with clusters, it shows some rich information and provides some useful functions. 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.

The View Tree of Clusters page might have been populated by manually adding clusters on the List Server Clusters page, or 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**. 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, 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 5-8.

To add a cluster, click **Add Cluster**. This opens the Add Server Cluster page (see [Figure 4-16](#) 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 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.

In the command line interface (CLI), use **cluster name create address** to give the cluster a name and address and set the important attributes. For example:

```
nrcmd> cluster example-cluster create 192.168.100.101 admin=admin password=changeme
```

Note that the administrator must be a superuser to fully synchronize at the local cluster. The local cluster SCP port defaults to 1234 and the HTTP port defaults to 8080.

Editing Local Clusters

To edit a local cluster, click its name on the View Tree of Server Clusters page or List Server Clusters page to open the Edit Server Cluster page. This 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**.

In the CLI, use **cluster name set attribute** to set or reset the attributes. For example:

```
nrcmd> cluster Example-cluster set poll-replica-interval=8h
```

Listing Related Servers for DHCP Servers


If you have related DNS, TFTP, or DHCP failover servers (see the “[Creating and Synchronizing Failover Server Pairs](#)” section on page 26-4), you can access the attributes for these servers from the View Tree of Server Clusters page once you create the clusters. 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 communication 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

Related Server Attribute	Description
Related Server Type	Type of related server: DHCP, DNS, or LDAP.
Related Server IP 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.
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.

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

Failover Partner Attribute	Description
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.
update-request-done-time	Time at which the last of any update request completed.

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



Failover Partner Attribute	Description
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 the List Related Servers for DHCP Server page or View Failover Related Server page, click **OK**.

In the CLI, use **dhcp getRelatedServers** to list the related servers in a failover pair.


Connecting to Local Clusters

If you have an equivalent administrator account at 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 at 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 at 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.

When you upgrade the local cluster, you should also resynchronize the cluster.



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 needs to occur before you can pull DHCP object data into the regional server’s database. 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.

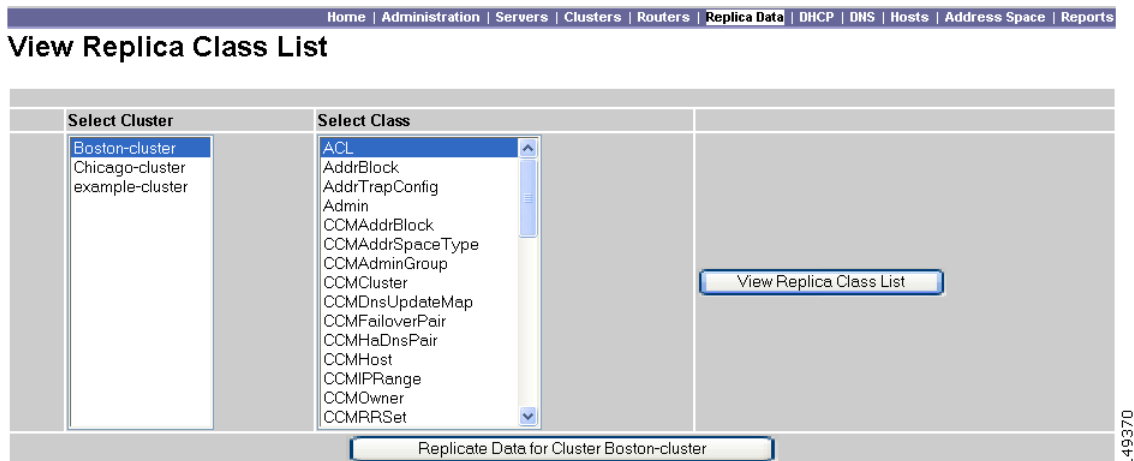
Replication happens at a given time interval. You can also force an immediate replication by clicking the Replicate icon (🔄) next to the cluster name on the View Tree of Server Clusters page or List Server Clusters page.

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 set the fixed time of day to poll replica data by using the *poll-replica-offset* attribute; its default is zero hours (no offset).

Viewing Replica Data


You can view the replica data cached in the replica database at the regional cluster by clicking **Replica Data**. This opens the View Replica Class List page (see Figure 5-1).

Figure 5-1 View Replica Class List Page (Regional)






On this page, select the:

1. Cluster in the Select Cluster list.
2. Object class in the Select Class list.

3. Replicate the data for the cluster and class chosen—Click the Replicate icon ()
4. View the replica data—Click **View Replica Class List**, which opens a List Replica Data for Cluster page for the cluster and specific class of object you choose. On this page, you can:
 - Click the name of an object to open a View page at 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 at the local cluster. Return to the List Replica *object* page by clicking the Go Regional icon ()

Click **Return** on the List Replica Data for Cluster page to return to the View Replica Class List page.

Deactivating, Reactivating, and Recovering Data for Clusters

Deactivating a cluster might be necessary if you suspect that a hard disk error occurred where configuration data could have been lost. You can deactivate the cluster, remedy the problem, recover cluster data from the replica database, then reactivate the cluster. This saves you from having to delete and then recreate the cluster with all of its data lost in the process.

Deactivating, reactivating, and recovering the data for a cluster is available only in the Web UI, and you must be an administrator assigned the central-config-admin role.



Data that is not recovered (and that you need to manually restore) includes:


- Contents of the cnr.conf file
- Web UI configuration files
- Unprotected DNS resource records
- Product licenses
- Administrator accounts
- Lease history
- Extension scripts





Note


Restoring the data to a different IP address requires some manual reconfiguration of such things as DHCP failover server pair and High-Availability (HA) DNS server pair addresses.

Deactivate a cluster in the Web UI by clicking the Activated icon () in the Activation column for the cluster. This immediately changes the icon to the Deactivated () icon to show the status of the cluster. Deactivating a cluster disables deleting, synchronizing, replicating data, and polling subnet utilization and lease history. These operations are not available while the cluster is deactivated.

Deactivating the cluster also displays the Recover icon () in the Recover Data column of the cluster. Click this icon to recover the replica data. This opens a separate “in process” status window that prevents any operations on the Web UI pages while the recovery is in process. As soon as the recovery is successful, the disabled functions are again enabled and available.

To reactivate the cluster, click the Deactivated icon () to change back to the Activated icon () and show the status as active.

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** (see the [“Generating Subnet Utilization History Reports”](#) section on page 8-13, or the [“Running IP Lease Histories”](#) section on page 21-15).

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 at 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, using a network time service for all 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. These attributes are set in three places at the regional cluster, with the following priority:

1. Failover pair (see the [“Managing DHCP Failover Pairs”](#) section on page 5-15)—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 occurs if the main server 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.In the CLI, set the attributes listed in [Table 5-3 on page 5-9](#) using the **failover-pair** command.
2. 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. In the CLI, set the attributes listed in [Table 5-3 on page 5-9](#), using the **cluster** command.
3. Regional CCM server (the default polling interval is 4 hours)—This is set on the Edit CCM Server page, accessible by clicking **Servers**, then the **Local CCM Server** link. In the CLI, set the attributes listed in [Table 5-3 on page 5-9](#) using the **ccm** command.

**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), no data is 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
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
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)
Polling priority for the regional failover pair—Pull data from the main or backup server first (failover pair setting only)	<i>poll-subnet-util-server-first</i> choose mainserver (default) or backupserver	<i>poll-lease-history-server-first</i> choose mainserver (default) or backupserver


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



- 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).

If you are configuring simple DHCP 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).

In the CLI, set the attributes using the **dhcp set** command.
- Step 3** Reload the local cluster DHCP server.
- Step 4** At 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 on page 5-9](#).

- Step 7** Click **Modify Cluster**.
- Step 8** 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

- 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).
- In the CLI, set the attributes using the **dhcp set** command.
- Step 3** Reload the local cluster DHCP server.
- Step 4** At 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 5-9](#).
- 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.
-

Managing 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.

For details on creating and editing scope templates, and applying them to scopes, see the [“Creating and Applying Scope Templates” section on page 19-2](#). The regional cluster Web UI has the added feature of pushing scope templates to local clusters and pulling them from local clusters.

Pushing Scope Templates to Local Clusters

You can 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.

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 (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 Replicate 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.

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 Scope Template** 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 (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.

Managing 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.

For details on creating and editing DHCP policies, and applying them to scopes, see the “[Configuring DHCP Policies](#)” section on page 20-1. The regional cluster Web UI has the added feature of pushing policies to, and pulling them from, the local clusters.

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** (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. Then click **Push Data to Clusters** to open the View Push Policy Data Report page.



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.

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 Replicate 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** (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.

Managing 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 hostnames—What name to assign clients.
- Denial of service—Whether unauthorized clients should be offered leases.

For details on creating and editing client-classes, see [Chapter 23, “Configuring Client-Classes and Clients.”](#) The regional cluster Web UI has the added feature of pushing client-classes to, and pulling them from, the local clusters.

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.

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 (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. Then click **Push Data to Clusters** to open the View Push Client-Class Data Report page.



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.

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 Replicate 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.

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 (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.

Managing Virtual Private Networks

A virtual private network (VPN) is a specialized address space identified by a key. 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.

For details on creating and editing VPNs, and applying them to various network objects, see the [“Configuring Virtual Private Networks Using DHCP” section on page 22-12](#). The regional cluster Web UI has the added feature of pushing VPNs to local clusters and pulling them from local clusters.

Pushing VPNs to Local Clusters

You can 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.

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 (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. Then click **Push Data to Clusters** to open the View Push VPN Data Report page.



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.

Pulling VPNs from Replica Data


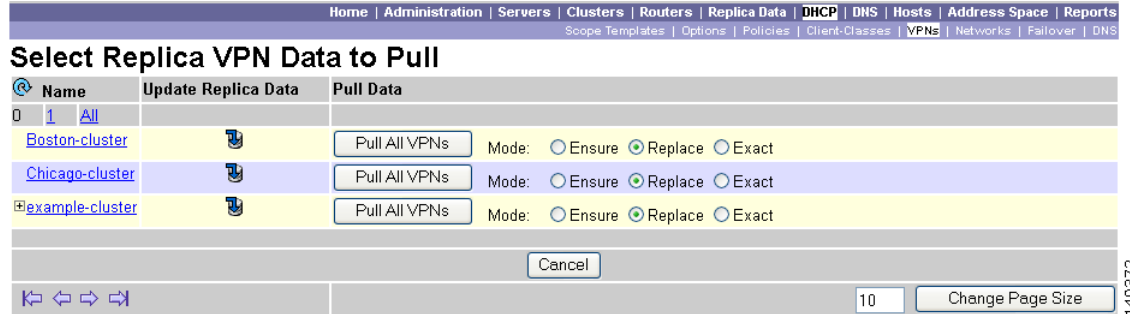
Instead of explicitly creating VPNs, you can pull them from the local clusters. (You may first want to update the VPN replica data by clicking the Replica icon  next to the cluster name.) To pull the replica data, click **Pull Replica VPNs** to open the Select Replica VPN Data to Pull page (see [Figure 5-2](#)).

Figure 5-2 Select Replica VPN Data to Pull Page (Regional)



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 (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.

Managing DHCP Failover Pairs

With DHCP failover, a backup DHCP server can take over for a main server if the latter comes off the network for any reason. You can use failover to configure two 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 responds to their lease request. 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**, then **Failover**. This functionality is available only to administrators who are assigned the dhcp-management subrole of the central-cfg-admin role.

For details on creating and editing failover pairs, see the “[Creating and Synchronizing Failover Server Pairs](#)” section on page 26-4. The regional cluster Web UI has the added feature of pulling addresses from local clusters to create the failover pairs.

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.
-

