Managing the Central Configuration

This chapter explains how to manage the central configuration at the Cisco Prime IP Express regional cluster.

Related Topics

Central Configuration Tasks

Central configuration management at the regional cluster can involve:

- Setting up server clusters, replicating their data, and polling lease history data from them.
- Setting up routers.
- 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 5-2 on page 5-4.) 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 5-15 and "Managing Servers" section on page 7-1.

Configuring Server Clusters

Server clusters are groupings of CCM, DNS, CDNS, and DHCP 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 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.

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 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/Add Remote 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, DNS, and TCP Listener Servers" section on page 6-4). These servers can be DNS, or DHCP failover servers.

Related Topics

Adding Local Clusters, page 6-2
Editing Local Clusters, page 6-3
Listing Related Servers for DHCP, DNS, and TCP Listener Servers, page 6-4
Connecting to Local Clusters, page 6-10
Synchronizing with Local Clusters, page 6-10
Replicating Local Cluster Data, page 6-11
Viewing Replica Data, page 6-11
Deactivating, Reactivating, and Recovering Data for Clusters, page 6-12
Polling Lease History Data, page 6-13
Enabling Lease History Collection, page 6-14

Adding Local Clusters

Adding local clusters to the regional cluster is the core functionality of the central-cfg-admin role.

To enable lease history data collection, see the "Polling Lease History Data" section on page 6-13.

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 CNRDB database is located. Obtain the SCP and HTTP ports, username, and password from the local cluster administrator. The preset value at Cisco Prime IP Express 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 Cisco Prime IP Express Communications Security Option installed to be effective.

Regional Web UI

From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu. This opens the Manage Servers page. View the local clusters on this page. You can also add server clusters on the List/Add Remote Clusters page. The List/Add Remote Clusters page 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 replica database.
- Query lease history data from a local cluster. This function appears only if you are assigned the regional-addr-admin role with at least the lease-history subrole.

To add a cluster, click the **Add Manage Clusters** icon in the Manage Clusters pane. This opens the Add Cluster dialog box. For an example of adding a local cluster, see the "Create the Local Clusters" section on page 5-34. Click **Add Cluster** to return to the List/Add Remote Clusters page.

Local Web UI

You can also manage clusters in the local web UI. See the "Configuring Clusters in the Local Web UI" section on page 2-9 for details.

CLI Commands

To add a cluster, 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.

Editing Local Clusters

Editing local clusters at the regional cluster is the core functionality of the central-cfg-admin role.

Regional Web UI

To edit a local cluster, click its name on the Manage Clusters pane to open the Edit Remote Cluster page. This page is essentially the same as the List/Add Remote Clusters page, except for an additional attribute unset function. You can choose the service (dhcp, dns, cdns, or none) that you want to run in the local by checking/unchecking the check boxes provided in the **Local Services** area. Make your changes, then click **Save**.

Local Web UI

You can also edit clusters in the local web UI. See the "Configuring Clusters in the Local Web UI" section on page 2-9 for details.

CLI Commands

To edit a local cluster, 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, DNS, and TCP Listener Servers

If you have related DNS, or DHCP failover servers (see the "Setting Up Failover Server Pairs" section on page 28-3), you can access the attributes for these servers.

Regional Web UI

On the Failover Pairs or HA DNS Server Pair page, click the Manage Failover Servers tab and then click Related Servers tab to open the DHCP Related Server Attributes page. This page shows the communication and failover states the servers are in. Table 6-1 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 6-1 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 6-2 on page 6-4).
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—None, Startup, Normal, Communications-interrupted, Partner-down, Potential-conflict, Recover, Paused, Shutdown, or Recover-done. For High-Availability (HA) DNS—Send-Update, Probe, or ha-state-unknown. Only the server that is successfully updating can be in Send-Update state. The partner server not sending updates is then always in Probe or unknown state. When the DHCP server comes up if there is no client activity, both DNS servers are often in the unknown state. This changes when the DHCP server tries to do DNS updates.
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 6-2 Attributes for DHCP Related Failover Servers

Failover Partner Attribute	Description
General attributes	
failover-pair-name	The name of the failover pair object used to manage this server.
current-time	Current time on the server returning this object.

Table 6-2 Attributes for DHCP Related Failover Servers (continued)

Failover Partner Attribute	Description
comm-state	None, OK, or Interrupted.
smoothed-time-delta	The time difference between the local server and the partner server. If the local server time is ahead of the partner server time, the attribute value is positive. If the local server time is behind the partner server time, the attribute value is negative. If the servers are not communicating, the last known attribute value is recorded.
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.
load-balancing-backup-pct	The current failover load balancing backup percentage. If the backup percentage is zero, failover load balancing is not in use (disabled).
Local server information	
our-ipaddr	IPv4 address of the interface to this server.
our-ip6address	IPv6 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 *update-request-in-progress* 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 *use-other-available* is false or unset, the time when *use-other-available* 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.
load-balancing-local-hba	The current hash bucket assignment of the local server, usually shown as a range of the hash bucket numbers. (See RFC 3074.)
request-buffers-in-use	The number of failover request buffers the DHCP server is using at the time the statistics are calculated.
decaying-max-request-buffers-in-use	The maximum number of failover request buffers that have recently been in use.
request-buffers-allocated	The number of request buffers that the server has allocated to support the failover capability.

Table 6-2 Attributes for DHCP Related Failover Servers (continued)

Failover Partner Attribute	Description
connection-start-time	The time at which the most recent connection started. This value is set whenever a connection is started, and it not cleared when a connection ended.
connection-end-time	The time at which the most recent connection ended. This value is set whenever a connection is ended, and it not cleared when a new connection starts.
Partner server information	
ipaddr	IP address of the partner server.
Iip6address	IPv6 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 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 current failover state began.
est-partner-end-recover-time	If the *partner-state* 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.
load-balancing-partner-hba	The current hash bucket assignment of the partner server, usually shown as a range of the hash bucket numbers. (See RFC 3074.)
partner-vendor-major-version	The vendor ID major version from the partner server.
partner-vendor-minor-version	The vendor ID minor version from the partner server.
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 *update-request-outstanding* request was started.
update-request-done-time	Time at which the last of any update request completed.
v6-update-response-in-progress	The type and origin of the response.
v6-update-response-percent-complete	The percent complete of the current IPv6 update response.
v6-update-response-start-time	The time that the IPv6 update response mentioned in v6-update-response-in-progress was started.
v6-update-response-done-time	The time that the most recent IPv6 update response sent an update done to the partner server.

Table 6-2 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 *update-response-outstanding* appears, the percent complete of the current update response.
update-response-start-time	Time that the update response mentioned in *update-response-in-progress* was started.
update-response-done-time	Time that the most recent update response sent an update done to the partner server.
Load Balancing Counters	
load-balancing-processed-requests	The number of server processed requests, both IPv4 and IPv6, subject to load balancing. This counter includes only the requests made after the latest transition of server to normal state.
load-balancing-dropped-requests	The number of server dropped requests, both IPv4 and IPv6, subject to load balancing. This counter includes only the requests made after the latest transition of server to normal state.
load-balancing-processed-total	The number of server processed requests, both IPv4 and IPv6, subject to load balancing. This counter includes the requests since this server was last started or reloaded.
load-balancing-dropped-total	The number of server dropped requests, both IPv4 and IPv6, subject to load balancing. This counter includes the requests since this server was last started or reloaded.
Binding Update or Ack Counters (this connection)	
binding-updates-sent	The number of binding update (BNDUPD) messages sent to the failover partner.
binding-acks-received	The number of binding acknowledgement (BNDACK) messages received from the failover partner.
binding-updates-received	The number of binding update (BNDUPD) messages received from the failover partner.
binding-acks-sent	The number of binding acknowledgement (BNDACK) messages sent to the failover partner.
v6-binding-updates-sent	The number of IPv6 binding updates (BNDUPD6) messages received from the failover partner since the start of the most recently established connection.
v6-binding-acks-received	The number of IPv6 binding acknowledgements (BNDACK6) messages received from the failover partner since the start of the most recently established connection.
v6-binding-updates-received	The number of IPv6 binding updates (BNDUPD6) messages received from the failover partner since the start of the most recently established connection.

Table 6-2 Attributes for DHCP Related Failover Servers (continued)

Failover Partner Attribute	Description
v6-binding-acks-sent	The number of IPv6 binding acknowledgements (BNDACK6) messages sent to the failover partner since the start of the most recently established connection.
Binding Update/Ack Counters Totals	
binding-updates-sent-total	The number of IPv4 binding updates (BNDUPD) messages sent to the failover partner since the most recent statistics reset.
binding-acks-received-total	The number of IPv4 binding acknowledgements (BNDACK) messages received from the failover partner since the most recent statistics reset.
binding-updates-received-total	The number of IPv4 binding updates (BNDUPD) messages received from the failover partner since the most recent statistics reset.
binding-acks-sent-total	The number of IPv4 binding acknowledgements (BNDACK) messages sent to the failover partner since the most recent statistics reset.
v6-binding-updates-sent-total	The number of IPv6 binding updates (BNDUPD6) messages sent to the failover partner since the most recent statistics reset.
v6-binding-acks-received-total	The number of IPv6 binding acknowledgements (BNDACK6) messages received from the failover partner since the most recent statistics reset.
v6-binding-updates-received-total	The number of IPv6 binding updates (BNDUPD6) messages received from the failover partner since the most recent statistics reset.
v6-binding-acks-sent-total	The number of IPv6 binding acknowledgements (BNDACK6) messages sent to the failover partner since the most recent statistics reset.
Flow Control Counters (this connection)	
current-binding-updates-in-flight	The current number of binding updates (both IPv4 and IPv6) that are currently in-flight (sent).
current-binding-updates-queued	The current number of binding updates (both IPv4 and IPv6) that are queued at present.
maximum-binding-updates-in-flight	The maximum number of binding updates (both IPv4 and IPv6) that were in-flight (sent) at one time.
maximum-binding-updates-queued	The maximum number of binding updates (both IPv4 and IPv6) that were queued at one time.
last-binding-update-sent-time	The time the last binding update (either IPv4 or IPv6) was sent.
last-binding-ack-received-time	The time the last IPv4 or IPv6 binding acknowledgement (whether NAKed or not) was received.
last-binding-update-received-time	The time the last binding update (either IPv4 or IPv6) was received.
last-binding-ack-sent-time	The time the last IPv4 or IPv6 binding acknowledgement (whether NAKed or not) was sent.

Table 6-3 Attributes for DNS Related Failover Servers

Failover Partner Attribute	Description
General attributes	
current-time	Current time on the server returning this object.
ipaddr	IP address
comm-state	None.
dns-server-state	PROBE.
probe-polling-event-id	Zero.
requests	Zero.
HA DNS Configuration information	
ha-dns-role	STANDALONE-DNS.
dns-timeout	Number of milliseconds that the DHCP server will wait for a response from the DNS server for a dynamic dns update, before retrying dynamic dns update.
max-dns-retries	Number of times that the DHCP server will try to send dynamic updates to a DNS server.
ha-dns-failover-timeout	Maximum time period, in seconds, the DHCP server will wait for a reply from a DNS server, before the DHCP will failover to use next DNS Server to perform the dynamic-update. Default value is 30 seconds.
ha-dns-probe-timeout	If cnr-ha-dns is enabled, DHCP server will use this timer to co-oridnate and reduce latency in failing over between HA-DNS servers, when HA-DNS servers are in COMMUNIATION-INTERRUPTED state or SYNCHRONIZING. Default value is 3 seconds.
ha-dns-probe-retry	If cnr-ha-dns is enabled, DHCP server will use this retry count and ha-dns-probe-timeout to co-oridnate and reduce latency in failing over between HA-DNS servers, when HA-DNS servers are in COMMUNIATION-INTERRUPTED state or SYNCHRONIZING. Default value is 1 retry attempt.
Current HA DNS State Information	
ha-dns-state	State of HA-DNS Servers interaction.
last-ha-dns-state	Failover role of the partner of the server returning this object—None, Main, or Backup.
last-ha-dns-state-change-time	Time at which the failover role was last changed.
last-reply-received-time	Time at which the last reply was received.
last-ha-dns-role-switch-time	Time at which the failover role was changed from one state to another.

Table 6-4 Attributes for TCP Listener Related Servers

Failover Partner Attribute	Description
General attributes	
comm-state	None.

Table 6-4 Attributes for TCP Listener Related Servers

Failover Partner Attribute	Description
current-connections	Zero
ipaddr	IP address.
ip6addr	IPv6 address.
name	foobar string (w/o null terminator).
port	Port number.
rejected-connections	Zero.
total-connections	Zero.

Other controls are available on these pages:

- To refresh the data on the Related Server tab, click **Refresh Data**.
- On the Related Server tab, 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. Never set both partners to Partner Down mode.
- To return from the List Related Servers for DHCP Server page or View Failover Related Server page, click **Return**.

CLI Commands

To list the related servers for a DHCP server, use **dhcp getRelatedServers**.

Connecting to Local Clusters

In the web UI, if you have an equivalent administrator account at the local cluster, you can single sign-on to the local cluster Manage Servers page by clicking the **Connect** icon on the List/Add Remote Clusters page. To return to the regional cluster web UI, click the **Return** icon at the top right corner of the local cluster page. If you do not have an equivalent account at the local cluster, the Connect icon opens the local cluster 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 re synchronize 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/Add Remote 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. 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.

Note When you resynchronize clusters at the regional cluster, an automatic reinitialization of replica data occurs. The result is that for larger server configurations, resynchronization might take several minutes. The benefit, however, is that you do not need a separate action to update the replica data.

Replicating Local Cluster Data

Replication is copying the configuration data from a local server to the regional cluster replica database. Replication needs to occur before you can pull DHCP object data into the regional server 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 on the List/Add Remote 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 preset at four hours. You can also set the fixed time of day to poll replica data by using the *poll-replica-offset* attribute; its default value is zero hours (no offset).

Caution If the replica database is corrupted in any way, the regional CCM server will not start. If you encounter this problem, stop the regional service, remove (or move) the replica database files located in the *install-path*/regional/data/replica directory (and the log files in the /logs subdirectory), then restart the regional server. Doing so recreates the replica database without any data loss.

Viewing Replica Data

In the web UI, you can view the replica data cached in the replica database at the regional cluster by choosing **View Replica Data** from **Servers** submenu under the **Operate** menu. This opens the View Replica Class List page.

Regional Web UI

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 Data for Cluster** button.

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 **Connect** icon to go to the List page for the object at the local cluster. Return to the List Replica *object* page by clicking the **Return** 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 (see the "Modifying the cnr.conf File" section on page 7-23)
- Web UI configuration files
- Unprotected DNS resource records
- Administrator accounts

  Note If the local secret db is lost, the old references are no longer valid, even though they are restored. To recover your passwords, you have to use central management for your admins, and then push them to your local clusters. Routers, since they have their own secrets, also need to be centrally managed and then should be re-pushed. For the local cluster partner objects, running the sync from regional will create valid objects, but the old cluster objects may need to be deleted first.

- 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.

Regional Web UI

Deactivate a cluster by clicking the Deactivate button for the cluster. This immediately changes the button to Reactivate to show the status of the cluster. Deactivating a cluster disables deleting, synchronizing, replicating data, and polling 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 Reactivate button to change back to the Deactivate button and show the status as active.

Polling Lease History Data

Lease history data is automatically collected at any regional cluster where these feature is 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 by clicking the **Lease History** icon on the List/Add Remote 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 are assigned the regional-addr-admin role with at least the lease-history subrole), you can query the lease history data by choosing Current Utilization or Lease History from **Operate** menu (see the "Running IP Lease Histories" section on page 23-21).

Related Topics

Polling Process, page 6-13
Adjusting the Polling Intervals, page 6-13

Polling Process

When the regional cluster polls the local cluster for 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 time, adjusted for that cluster time zone.

If the times on each server are not synchronized, you might observe odd query results. For example, if the regional cluster time lags behind that of a local cluster, 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 lease history, along with other attributes. These attributes are set in three places at the regional cluster, with the following priority:

1. **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 6-5, using the **cluster** command.

2. **Regional CCM server** (the preset 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 6-5 using the **ccm** command.

Note If lease history collection is not explicitly turned on at the local cluster DHCP server (see the "Enabling Lease History Collection" section on page 6-14), no data is collected, even though polling is on by default.

Table 6-5 Subnet Utilization and Lease History Polling Regional Attributes

Attribute Type	Lease History
Polling interval—How often to poll data	*poll-lease-hist-interval* 0 (no polling) to 1 year, preset to 4 hours for the CCM server
Retry interval—How often to retry after an unsuccessful polling	*poll-lease-hist-retry* 0 to 4 retries
Offset—Hour of the day to guarantee polling	*poll-lease-hist-offset* 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 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***—Enable or disable the lease history database for v4-only (DHCPv4), v6-only (DHCPv6), or both.
- ***ip-history-max-age***—Limit on the age of the history records (preset to 4 weeks).

  In the CLI, set the attributes using the **dhcp set ip-history=<value> (v4-only, v6-only, both, or disable)** command.

Step 3 If in staged dhcp edit mode, reload the local cluster DHCP server.

Step 4 At the regional cluster, create the cluster that includes this DHCP server.

Step 5 In the regional web UI, go to the Lease History Settings section of the List/Add Remote Clusters page.

Step 6 Set the attributes in Table 6-5 on page 6-14.

Step 7 Click **Save**.

Step 8 On the List/Add Remote Clusters page, click the **Replica** icon next to the cluster name.

Step 9 Click the **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/Add DHCP Scope Templates page (choose **Scope Templates** from the **Design > DHCPv4** menu).

For details on creating and editing scope templates, and applying them to scopes, see the "Creating and Applying Scope Templates" section on page 21-3. The regional cluster web UI has the added feature of pushing scope templates to local clusters and pulling them from local clusters.

Related Topics

Pushing Scope Templates to Local Clusters, page 6-15
Pulling Scope Templates from Replica Data, page 6-16

Pushing Scope Templates to Local Clusters

You can push the scope templates you create from the regional cluster to any of the local clusters. In the web UI, go to the List/Add DHCP Scope Templates page, and do any of the following:

- if you want to push a specific template to a cluster, select the scope template from the Scope Templates pane on the left, and click **Push** (at the top of the page). This opens the Push DHCP Scope Template page.
- If you want to push all of the available scope templates, click the **Push All** icon at the top of the Scope Templates pane. This opens the Push Data to Local Clusters page.

Regional Web UI

The Push DHCP Scope Template page and Push Data to Local Clusters page identify 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** (preset value)—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 in the regional web UI, click the **Pull Replica** icon at the top of the Scope Templates pane.

Regional Web UI

The Select Replica DHCP Scope Template Data to Pull page shows a tree view of the regional server 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**.

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** (preset value)—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 22-1. The regional cluster web UI has the added feature of pushing policies to, and pulling them from, the local clusters.

Related Topics

Pushing Policies to Local Clusters, page 6-16
Pulling Policies from Replica Data, page 6-17

Pushing Policies to Local Clusters

You can also push the policies you create from the regional cluster to any of the local clusters. In the regional web UI, go to List/Add DHCP Policies page, and do any of the following:

- If you want to push a specific policy to a cluster, select the policy from the Policies pane on the left, and click **Push** (at the top of the page).
- If you want to push all the policies, click the **Push All** icon at the top of the Policies pane.

Regional Web UI

The Push DHCP Policy Data to Local Clusters 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** (preset value)—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. (In the regional web UI, 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 the **Pull Replica** icon at the top of the Policies pane.

Regional Web UI

The Select Replica DHCP Policy Data to Pull page shows a tree view of the regional server 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**.

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** (preset value)—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 the Cisco Prime IP Express 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 25, "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.

Related Topics

Pushing Client-Classes to Local Clusters, page 6-18
Pulling Client-Classes from Replica Data, page 6-18

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. In the Regional web UI, go to the List/Add DHCP Client Classes page, and do any of the following:

- If you want to push a specific client-class to a cluster in the web UI, select the client-class from the Client Classes pane on the left, and click **Push** (at the top of the page). This opens the Push DHCP Client Class page.
- If you want to push all the client-classes, click the **Push All** icon at the top of the Client Classes pane. This opens the Push Data to Local Clusters page.

Regional Web UI

The Push DHCP Client Class page and Push Data to Local Clusters 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** (preset value)—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. (In the web UI, you might 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 the **Pull Replica** icon at the top of the Client Classes pane.

Regional Web UI

The Select Replica DHCP Client-Class Data to Pull page shows a tree view of the regional server 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**.

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** (preset value)—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 24-18. The regional web UI has the added feature of pushing VPNs to local clusters and pulling them from local clusters.

Related Topics

Pushing VPNs to Local Clusters, page 6-19
Pulling VPNs from Replica Data, page 6-20

Pushing VPNs to Local Clusters

You can push the VPNs you create from the regional cluster to any of the local clusters. In the Regional web UI, go to the List/Add VPNs page, and do any of the following:

- If you want to push a specific VPN to a cluster in the web UI, select the VPN from the VPNs pane on the left, and click **Push** (at the top of the page). This opens the Push VPN page.
- If you want to push all the VPNs, click the **Push All** icon at the top of the VPNs pane. This opens the Push Data to Local Clusters page.

Regional Web UI

The Push VPN page and Push Data to Local Clusters page identify 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** (preset value)—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. (In the regional web UI, 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 the **Pull Replica** icon at the top of the VPNs pane on the left, to open the Select Replica VPN Data to Pull page.

This page shows a tree view of the regional server 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**.

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** (preset value)—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.

In the regional web UI, you can view any created failover pairs on the List/Add DHCP Failover Pairs page. To access this page, click **DHCP**, 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 "Setting Up Failover Server Pairs" section on page 28-3. 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 regional-addr-admin privileges.

Regional Web UI

Step 1 On the List/Add DHCP Failover Pairs page or View Unified Address Space page, click the **Pull Replica** icon in the Failover Pairs pane.

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 the **Report** button in the Synchronize Failover Pair tab and click **Return**.

Step 4 Click **Run** on the Report Pull Replica Address Space page.

Step 5 Click **OK** on the Run Pull Replica Address Space page.

Managing Lease Reservations

You can push lease reservations you create from the regional cluster to any of the local clusters. In the regional cluster web UI, go to the List/Add DHCPv4 Reservations page or List/Add DHCPv6 Reservations page, and click the **Push All** icon in the Reservations pane on the left. Note that you cannot push individual reservations. If the cluster pushed to is part of a DHCP failover configuration, pushing a reservation also pushes it to the partner server.

Related Topics

DHCPv4 Reservations, page 6-21
DHCP v6 Reservations, page 6-21

DHCPv4 Reservations

To create DHCPv4 reservations, the parent subnet object must exist on the regional server. If there are pending reservation edits at regional, these can be pushed to the subnet local cluster or failover pair. If the subnet has never been pushed, the parent scope is added to the local cluster or pair.

Once a subnet is pushed to a local cluster or pair, reservations are pushed to that cluster or pair. To move the scopes and subnet to another local cluster or failover pair, the subnet must first be reclaimed.

DHCP v6 Reservations

To create DHCPv6 reservations, the parent prefix must exist on the regional server. When there are pending reservation or prefix changes, you can push the updates to the local cluster.

Once a prefix is pushed to a local cluster, it can only update that local cluster. To move the prefix to another local cluster, it must first be reclaimed.

Regional Web UI

The ensuing 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**—Ensures that the local cluster has new data without affecting any existing data.
- **Replace** (preset value)—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 Reservations Data Report page. Click **OK** on this page.

You can also pull the replica address space on the List/Add DHCP v6 Reservations page, and opt whether to omit reservations when doing so. You should use this option only to reduce processing time when you are sure that there are no pending changes to reservations to merge. To omit reservations for the pull, check the *Omit Reservations?* check box, then click **Pull Data**.

See the "Managing DHCPv6 Addresses" section on page 27-1.