Managing the Central Configuration

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

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.
- 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 2.) 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 Licensing, on page 3 and Managing Servers.

Default Ports for Cisco Prime IP Express Services

The following table lists the default ports used for the Cisco Prime IP Express services.

Table 1: Default Ports for Cisco Prime IP Express Services

Port Number	Protocol	Service
53	TCP/UDP	DNS
53	TCP/UDP	Caching DNS
67	UDP	DHCP client to server
67	TCP	Bulk or Active leasequery client to DHCP server
68	UDP	DHCP server to client
80	HTTP	BYOD web server client to server web UI
162	TCP	SNMP traps server to server
389	TCP	DHCP server to LDAP server
443	HTTPS	BYOD web server secure client to server web UI
546	UDP	DHCPv6 server to client
547	UDP	DHCPv6 client to server
647	TCP	DHCP failover server to server

Port Number	Protocol	Service
653	TCP	High-Availability (HA) DNS server to server
1234	TCP	Local cluster CCM server to server
1244	TCP	Regional cluster CCM server to server
4444	TCP	SNMP client to server
5480	HTTPS	Virtual Appliance
8080	HTTP	Local cluster client to server web UI
8090	HTTP	Regional cluster client to server web UI
8443	HTTPS	Local cluster secure client to server web UI
8453	HTTPS	Regional cluster secure client to server web UI

Firewall Considerations

When DNS (caching or authoritative) servers are deployed behind a stateful firewall (whether physical hardware or software, such as conntrack), it is recommended that:

- For at least UDP DNS traffic, stateful support be disabled if possible.
- If it is not possible to disable the stateful support, the number of allowed state table entries may need to be significantly increased.

DNS queries typically arrive from many different clients and requests from the same client may use different source ports. With thousands of queries per second, the number of these different sources can be large and if a firewall is using stateful tracking, it has to keep this state and does so for a period of time. Hence, you need to assure that the firewall can hold sufficient state - given the query traffic rates and the state time interval.

Licensing

Cisco Prime IP Express provides separate license for CCM, Authoritative DNS, Caching DNS, DHCP, and IPAM services or for combinations of these services. For more details on the Licensing, see the “License Files” section in the Overview chapter of the *Cisco Prime IP Express Installation Guide* .

You must have the Central Configuration Management (CCM) license to log into the UI. See Logging In to the Web UIs for entering license data the first time you try to log in. You can add the additional service based licenses in the regional server after you log in.

Whenever you log into a regional or local cluster, the overall licensing status of the system is checked. If there are any violations, you will be notified of the violation and the details. This notification is done only once for each user session. In addition, you will be able to see a message on each page indicating the violation.

Regional Web UI

Choose **Licenses** from **Administration > User Access** to open the List/Add Product Licenses page. Click **Browse** to locate the license file, click the file, then click **Open**. If the license ID in the file is valid, the license key appears in the list of licenses with the message "Successfully added license file "*filename* ." If the ID is not valid, the License field shows the contents of the file and the message "Object is invalid" appears.

The License Utilization section at the top of the page lists the type of license, the number of nodes allowed for the license, and the actual number of nodes used. Expand the section by clicking the plus (+) sign. The license utilization for each licensed service is listed separately in this section.

The Right To Use and the In Use counts are displayed for each licensed service. The Right To Use value will be the aggregation of the counts across all added licenses for that service. The 'total in use' value will be the aggregation of the latest utilization numbers obtained from all the local clusters. Only the services having a positive Right to use or In Use count will be listed in this section.

Licenses and usage count of earlier versions of Cisco IP Express will be listed under a separate section "ip-node".

The **Expert** mode attribute lets you specify how often license utilization is collected from all the local clusters. Changes to this setting require a server restart to take effect. You can set this attribute at the Edit CCM Server page. The default value is 4 hours.

CLI Commands

Use **license** *file* **create** to register licenses that are stored in file. The file referenced should include its absolute path or path relative to where you execute the commands. For example:

```
nrcmd-R> license "C:\licenses\product.licenses" create
```

Use **license list** to list the properties of all the created licenses (identified by key), and **license listnames** to list just the keys. Use **license** *key* **show** to show the properties of a specific license key.

Adding License

Cisco will e-mail you one or more license files after you register the Cisco Prime IP Express Product Authorization Key (PAK) on the web according to the Software License Claim Certificate shipped with the product. Cisco administers licenses through a FLEXlm system. Once you have the file or files:

Regional Web UI

Step 1 Locate the license file or files in a directory (or on the desktop) that is easy to find.

Step 2 On the List/Add Product Licenses page, browse for each file by clicking the **Choose File** button.

Note The List/Add Product Licenses option is only available at the Regional.

Step 3 In the Choose file window, find the location of the initial license file, then click **Open**.

Step 4 If the license key is acceptable, the Add Superuser Administrator page appears immediately.

Step 5 To add further licenses, from **Administration** menu choose **Licenses** under the **User Access** submenuto open the List/Add Product Licenses page. Click **Browse** to open the Choose file window, locate the additional license file, then click **Open**. If the key in the file is acceptable, the key, type, count, and expiration date appear, along with whether it is an evaluation key. If the key is not acceptable, the page shows the license text along with an error message. For the list of license types, see Licensing, on page 3.
Above the table of licenses is a License Utilization area that, when expanded, shows the license types along with the total nodes that you can use and those actually used.

If Cisco Prime IP Express is installed as a distributed system, the license management is done from the regional cluster. You will not have the option of adding licenses in local cluster.

Registering a Local Cluster that is Behind a NAT

License management is done from the regional cluster when Cisco Prime IP Express is installed. You must install the regional cluster first, and load all licenses in the regional cluster. A local cluster can register with a regional either by registering with the regional cluster during the installation process. However, if the local cluster is behind a NAT instance, then the registration may fail because the initial request does not reach the regional cluster.

In Cisco Prime IP Express 8.3 and later, you can register a local cluster that is behind a NAT instance by initiating the registration from the local cluster. To register a local cluster that is spanned by a NAT instance, you must ensure that Cisco Prime IP Express 8.3 or later is installed on both the regional and local clusters. You can also verify the license utilization for the local cluster.

Note To register a local cluster when the regional cluster is behind a NAT instance, you need to register the local cluster from the regional server by registering the local cluster from the regional server, selecting the services and re-synchronizing the data.

To register a local cluster that is behind a NAT instance, do the following:

Local Web UI

Step 1 From **Administration** menu , choose **Licenses** under the **User Access** submenu to open the List Licenses page. On the List Licenses page, add the details of the regional cluster.

a) Enter the IP address of the regional cluster.
b) Enter the SCP port of the regional cluster (1244 is the preset value).
c) Select the IP address of the local cluster that you want to register.
d) Select the component services that you want to register for the local cluster.

Step 2 Click **Register**.

Note The regional CCM server maintains the license utilization history for all the local clusters in the Cisco Prime IP Express system for all counted services (DHCP, DNS, and CDNS).

To view the license utilization for the local cluster, click **Check Poll Status**.

CLI Commands

Use the following command to register or re-register a local cluster:

```
nrcmd> license register [cdns|dns|dhcp[,...]] [<regional-ip>] [<regional-port>]
```

License History

The License History page allows you to view the licenses utilized in the specified time frame.

Regional Web UI

Step 1 Log into the regional cluster as superuser.

Step 2 From the **Administration** menu, choose **Administrators** to open the List/Add Administrators page for the local cluster version of this page, which is essentially identical.

Step 3 Click the Add Administrators icon in the Administrators pane, enter **example-regional-admin** in the Name field, then **examplereg** in the Password field in the Add Administrator dialog box, then click Add Administrator.

Step 4 Multiselect **central-cfg-admin-group** (for cluster administration) and **regional-admin-group** (for user administration) in the Groups drop-down list.

Step 5 Click **Save**.

CLI Command

Use **license showUtilHistory –full** view the number of utilized IP nodes against the RTUs (Right-to-Use) (see the **license** command in the CLIGuide.html file in the /docs directory for syntax and attribute descriptions).

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 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.

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 re-synchronization, 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 Listing Related Servers for DHCP, DNS, and TCP Listener Servers, on page 9). These servers can be DNS, or DHCP failover servers.

Related Topics

Adding Local Clusters, on page 7

Editing Local Clusters, on page 8

Listing Related Servers for DHCP, DNS, and TCP Listener Servers, on page 9

Connecting to Local Clusters, on page 18

Synchronizing with Local Clusters, on page 18

Replicating Local Cluster Data, on page 18

Viewing Replica Data, on page 19

Polling Lease History Data, on page 32

Deactivating, Reactivating, and Recovering Data for Clusters, on page 20

Enabling Lease History Collection, on page 33

Adding Local Clusters

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

To enable subnet utilization and lease history data collection, see Polling Lease History Data, on page 32.

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.

- Purge replica to clear the bad replica data without deleting/re-adding the cluster. Whenever you perform purge replica, you must perform manual replication to the get the replica data again.

Note This option appears only in Expert mode.

- Query subnet utilization data from a local cluster. This function appears only if you 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 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 Create the Local Clusters, on page 54. 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 Configuring Clusters in the Local Web UI 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 Configuring Clusters in the Local Web UI 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 in Cisco PrimeIP Express 8.3 DHCP User Guide*), 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. The following table 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 2: 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 3: Attributes for DHCP Related Failover Servers , on page 10).
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.

Related Server Attribute	Description
Update Response Complete	For DHCP failover only, the percentage of completed update responses, valid only if there are outstanding update responses.

Table 3: 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.
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.

Failover Partner Attribute	Description
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.
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.

Failover Partner Attribute	Description
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.
ip6address	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.

Failover Partner Attribute	Description
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.
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)	

Failover Partner Attribute	Description
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.
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.

Failover Partner Attribute	Description
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 4: 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	

Failover Partner Attribute	Description
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 5: Attributes for TCP Listener Related Servers

Failover Partner Attribute	Description
General attributes	
comm-state	None.
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). The *poll-replica-rrs* attribute controls the replication of RR data without disabling other data replication. This attribute is present in Manage Servers and Manage Clusters page and has the values - none, all, and protected. If *poll-replica-rrs* is set to none, no RR data will be replicated for this cluster. If unset, the CCM server setting will apply.

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.

Purging Replica Data

In the Regional web UI (Expert mode), you can clear the bad replica data without deleting/re-adding the clusters by clicking the **Purge Replica** icon on the List/Add Remote Clusters page. Whenever you perform purge replica, you must perform manual replication to get the replica data again.

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 Modifying the cnr.conf File)
- 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 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 Reactivate button to change back to the Deactivate button and show the status as active.

Central Configuration Management Server

The CCM servers at the local and regional clusters provide the infrastructure for Cisco Prime IP Express operation and user interfaces. The CCM Server reads, writes, and modifies the Cisco Prime IP Express database (CCM DB). The main purpose of the CCM Server is to store and propagate data from the user to the protocol servers, and from the servers back to the user.

The change set is the fundamental unit of change to a data store. It sends incremental changes to a replicating server and provides an audit log for changes to the data store. Change sets consist of lists of change entries that are groups of one or more changes to a single network object. The web UI provides a view of the change sets for each data store.

Managing CCM Server

You can view logs and startup logs; edit the server attributes.

To view logs and startup logs, in the local cluster web UI, from the **Operate** menu, choose **Manage Servers** to open the Manage Servers page.

Editing CCM Server Properties

You can edit the CCM server properties using the Edit CCM Server page.

Local Basic or Advanced Web UI

Step 1 To access the CCM server properties, choose **Manage Servers** under **Operate** menu to open the Manage Servers page.

Step 2 Click **Local CCM Server** in the Manage Servers pane on the left. The Edit Local CCM Server page appears. This page displays all the CCM server attributes.

Step 3 Modify the settings as per your requirement.

Step 4 Click **Save** to save the CCM server attribute modifications.

Simple Network Management

The Cisco Prime IP Express Simple Network Management Protocol (SNMP) notification support allows you to query the DHCP and DNS counters, be warned of error conditions and possible problems with the DNS and DHCP servers, and monitor threshold conditions that can indicate failure or impending failure conditions.

Cisco Prime IP Express implements SNMP Trap Protocol Data Units (PDUs) according to the SNMPv2c standard. Each trap PDU contains:

- Generic-notification code, if enterprise-specific.
- A specific-notification field that contains a code indicating the event or threshold crossing that occurred.

- A variable-bindings field that contains additional information about certain events.

Refer to the Management Information Base (MIB) for the details. The SNMP server supports only reads of the MIB attributes. Writes to the attributes are not supported.

The following MIB files are required:

- **Traps—**CISCO-NETWORK-REGISTRAR-MIB.my and CISCO-EPM-NOTIFICATION-MIB.my
- **DNS server—**CISCO-DNS-SERVER-MIB.my

Note The Caching DNS server requires only a subset of the DNS MIB when it is operating. Caching DNS server only supports the *server-start* and *server-stop* notification events.

- **DHCPv4 server—**CISCO-IETF-DHCP-SERVER-MIB.my
- **DHCPv4 server capability—**CISCO-IETF-DHCP-SERVER-CAPABILITY.my
- **DHCPv4 server extensions—**CISCO-IETF-DHCP-SERVER-EXT-MIB.my
- **DHCPv4 server extensions capability—**CISCO-IETF-DHCP-SERVER-EXT-CAPABILITY.my
- **DHCPv6 server—**CISCO-NETREG-DHCPV6-MIB.my (experimental)

Note The MIB, CISCO-NETREG-DHCPV6-MIB is defined to support query of new DHCP v6 related statistics and new DHCP v6 traps.

These MIB files are available in the /misc directory of the Cisco Prime IP Express installation path.

The following URL includes all files except the experimental CISCO-NETREG-DHCPV6-MIB.my file:

ftp://ftp.cisco.com/pub/mibs/supportlists/cnr/cnr-supportlist.html

The following dependency files are also required:

- **Dependency for DHCPv4 and DHCPv6—**CISCO-SMI.my
- **Additional dependencies for DHCPv6—**INET-ADDRESS-MIB.my

These dependency files are available along with all the MIB files at the following URL:

ftp://ftp.cisco.com/pub/mibs/v2/

To get the object identifiers (OIDs) for the MIB attributes, go to the equivalently named .oid file at:

ftp://ftp.cisco.com/pub/mibs/oid/

Related Topics

How Notification Works, on page 24

Handling SNMP Queries, on page 28

Setting Up the SNMP Server

To perform queries to the SNMP server, you need to set up the server properties.

Local Basic or Advanced and Regional Web UI

Step 1 From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page (see Managing Servers).

Step 2 Click the **Local SNMP Server** link to open the Edit Local SNMP Server page.

Step 3 The *Community string* attribute is the password to access the server. (The community string is a read community string only.) The preset value is **public**.

Step 4 You can specify the Log Settings, Miscellaneous Options and Settings, and Advanced Options and Settings:

- **trap-source-addr—**Optional sender address to use for outgoing traps.
- **trap-source-ip6address—** Optional sender IPv6 address to use for outgoing traps.
- **server-active—**Determines whether the SNMP server is active for queries. The default value is true. If set to false, the server will run, but is not accessible for queries and does not send out traps.
- **cache-ttl—**Determines how long the SNMP caches responds to queries, default to 60 seconds.

Step 5 To manage the SNMP server interfaces in the Advanced mode, click the **Network Interfaces** tab. You can view the default configured network interfaces, and create and edit additional ones. To create and edit them, you must be assigned the server-management subrole of the ccm-admin role.

Step 6 To manage trap recipients for the server:

a) Click the **Trap Recipients** tab.
b) Enter the name of the trap recipient.
c) Enter the IPv4 and/or IPv6 address of a trap recipient.
d) Click **Add Trap Recipient**.
e) Repeat for each additional trap recipient.
f) To set the port, community string, and agent address for a trap recipient, click its name on the Trap Recipients tab to open the Edit Trap Recipient page, then set the values.

Step 7 Complete the SNMP server setup by clicking **Save**.

CLI Commands

To set the community string in the CLI so that you can access the SNMP server, use **snmp set community=***name* . Use **snmp set trap-source-addr** to set the trap source IPv4 address. Use **snmp set trap-source-ip6address** to set the trap source IPv6 address. Use **snmp disable server-active** to deactivate the SNMP server and **snmp set cache-ttl=***time* to set the cache time-to-live.

To set trap recipients, use **trap-recipient**, in the following syntax to include the IP address:

```
nrcmd> trap-recipient name create ip-addr=
nrcmd> trap-recipient name create ip6address=
```

You can also add the *agent-address* , *community* , and *port-number* values for the trap recipient.

Other SNMP-related commands include **snmp disable server-active** to prevent the server from running when started and the **snmp-interface** commands to configure the interfaces.

How Notification Works

Cisco Prime IP Express SNMP notification support allows a standard SNMP management station to receive notification messages from the DHCP and DNS servers. These messages contain the details of the event that triggered the SNMP trap.

Cisco Prime IP Express generates notifications in response to predetermined events that the application code detects and signals. Each event can also carry with it a particular set of parameters or current values. For example, the *free-address-low-threshold* event can occur in the scope with a value of 10% free. Other scopes and values are also possible for such an event, and each type of event can have different associated parameters.

The following table describes the events that can generate notifications.

Table 6: SNMP Notification Events

Event	Notification
Address conflict with another DHCP server detected (*address-conflict*)	An address conflicts with another DHCP server.
DNS queue becomes full (*dns-queue-size*)	The DHCP server DNS queue fills and the DHCP server stops processing requests. (This is usually a rare internal condition.)
Duplicate IP address detected (*duplicate-address* and *duplicate-address6*)	A duplicate IPv4 or IPv6 address occurs.
Duplicate IPv6 prefix detected (*duplicate-prefix6*)	A duplicate IPv6 prefix occurs.
Failover configuration mismatch (*failover-config-error*)	A DHCP failover configuration does not match between partners.
Free-address thresholds (*free-address-low* and *free-address-high* ; or *free-address6-low* and *free-address6-high*)	The high trap when the number of free IPv4 or IPv6 addresses exceeds the high threshold; or a low trap when the number of free addresses falls below the low threshold after previously triggering the high trap.
High-availability (HA) DNS configuration mismatch (*ha-dns-config-error*)	An HA DNS configuration does not match between partners.
HA DNS partner not responding (*ha-dns-partner-down*)	An HA DNS partner stops responding to the DNS server.

Event	Notification
HA DNS partner responding (*ha-dns-partner-up*)	An HA DNS partner responds after having been unresponsive.
DNS masters not responding (*masters-not-responding*)	Master DNS servers stop responding to the DNS server.
DNS masters responding (*masters-responding*)	Master DNS servers respond after having been unresponsive.
Other server not responding (*other-server-down*)	A DHCP failover partner, or a DNS or LDAP server, stops responding to the DHCP server.
Other server responding (*other-server-up*)	DHCP failover partner, or a DNS or LDAP server, responds after having been unresponsive.
DNS secondary zones expire (*secondary-zone-expired*)	A DNS secondary server can no longer claim authority for zone data when responding to queries during a zone transfer.
Server start (*server-start*)	The DHCP or DNS server is started or reinitialized.
Server stop (*server-stop*)	The DHCP or DNS server is stopped.

Handling SNMP Notification Events

When Cisco Prime IP Express generates a notification, it transmits a single copy of the notification as an SNMP Trap PDU to each recipient. All events (and scopes or prefixes) share the list of recipients and other notification configuration data, and the server reads them when you initialize the notification.

You can set SNMP attributes in three ways:

- For the DHCP server, which includes the traps to enable and the default free-address trap configuration if you are not specifically configuring traps for scopes or prefixes (or their templates).
- On the scope or prefix (or its template) level by setting the *free-address-config* attribute.
- For the DNS server, which includes a *traps-enabled* setting.

To use SNMP notifications, you must specify trap recipients that indicate where trap notifications should go. By default, all notifications are enabled, but you must explicitly define the recipients, otherwise no notifications can go out. The IP address you use is often **localhost**.

The DHCP server provides special trap configurations so that it can send notifications, especially about free addresses for DHCPv4 and DHCPv6. You can set the trap configuration name, mode, and percentages for the low threshold and high threshold. The mode determines how scopes aggregate their free-address levels.

DHCP v4 Notification

The DHCP v4 modes and thresholds are (see also Handling Deactivated Scopes or Prefixes, on page 26):

- **scope mode—**Causes each scope to track its own free-address level independently (the default).
- **network mode—**Causes all scopes set with this trap configuration (through the scope or scope template *free-address-config* attribute) to aggregate their free-address levels if the scopes share the same *primary-subnet* .
- **selection-tags mode—**Causes scopes to aggregate their free-address levels if they share a primary subnet and have a matching list of selection tag values.
- **low-threshold—**Free-address percentage at which the DHCP server generates a low-threshold trap and re-enables the high threshold. The free-address level for scopes is the following calculation:

```
100 * available-nonreserved-leases
total-configured-leases
```

- **high-threshold—**Free-address percentage at which the DHCP server generates a high-threshold trap and re-enables the low threshold.

DHCP v6 Notification

The DHCP v6 modes and thresholds are (see also Handling Deactivated Scopes or Prefixes, on page 26):

- **prefix mode—**Causes each prefix to track its own free-address level independently.
- **link mode—**Causes all prefixes configured for the link to aggregate their own free-address levels if all prefixes share the same link.
- **v6-selection-tags mode—**Causes prefixes to aggregate their free-address levels if they share a link and have a matching list of selection tag values.
- **low-threshold—**Free-address percentage at which the DHCP server generates a low-threshold trap and re-enables the high threshold. The free-address level for prefixes is the following calculation:

```
100 * max-leases - dynamic-leases
max-leases
```

- **high-threshold—**Free-address percentage at which the DHCP server generates a high-threshold trap and re-enables the low threshold.

Handling Deactivated Scopes or Prefixes

A deactivated scope or prefix never aggregates its counters with other scopes or prefixes. For example, if you configure a prefix with **link** or **v6-selection-tags** trap mode, and then deactivate the prefix, its counters disappear from the total count on the aggregation. Any changes to the leases on the deactivated prefix do not apply to the aggregate totals.

Therefore, to detect clients for deactivated scopes or prefixes, you must set the event mode to **scope** or **prefix**, and not to any of the aggregate modes (**network**, **selection-tags**, **link**, or **v6-selection-tags**).

The use case for setting traps on deactivated prefixes, for example, is network renumbering. In this case, you might want to monitor both the new prefixes (as an aggregate, ensuring that you have enough space for all

the clients) and old prefixes to ensure that their leases are freed up. You would probably also want to set the high threshold on an old prefix to 90% or 95%, so that you get a trap fired when most of its addresses are free.

Local Basic or Advanced Web UI

Access the SNMP attributes for the DHCP server by choosing **Manage Servers** from the **Operate** menu, then click **Local DHCP Server** in the left pane. You can view the SNMP attributes under SNMP (in Basic mode) or SNMP Settings (in Advanced mode) in the Edit DHCP Server page.

The four *lease-enabled* values (free-address6-low, free-address6-high, duplicate-address6, duplicate-prefix6) pertain to DHCPv6 only. Along with the traps to enable, you can specify the default free-address trap configuration by name, which affects all scopes and prefixes or links not explicitly configured.

To add a trap configuration, do the following:

Step 1 In Advanced mode, from the **Deploy menu** choose **Traps** under the DHCP submenu to access the DHCP trap configurations. The List/Add Trap Configurations page appears.

Step 2 Click the **Add Traps** icon in the left pane to open the Add AddrTrapConfig page.

Step 3 Enter the name, mode, and threshold percentages, then click **Add AddrTrapConfig**.

To edit a trap configuration, do the following:

Step 1 Click the desired trap name in the Traps pane to open the Edit Trap Configuration page

Step 2 Modify the name, mode, or threshold percentages.

Step 3 Click the **on** option for the *enabled* attribute to enable the trap configuration.

Step 4 Click **Save** for the changes to take effect.

Deleting Trap Configuration

To delete a trap configuration, select the trap in the Traps pane and click the **Delete** icon, then confirm or cancel the deletion.

Regional Basic or Advanced Web UI

In the regional web UI, you can add and edit trap configurations as in the local web UI. You can also pull replica trap configurations and push trap configurations to the local cluster on the List/Add Trap Configurations page.

Server Up/Down Traps

Every down trap must be followed by a corresponding up trap. However, this rule is not strictly applicable in the following scenarios:

1. If a failover partner or LDAP server or DNS server or HA DNS partner is down for a long time, down traps will be issued periodically. An up trap will be generated only when that server or partner returns to service.
2. If the DHCP or DNS server is reloaded or restarted, the prior state of the partner or related servers is not retained and duplicate down or up traps can result.

Note Other failover partner or LDAP server or DNS server or HA DNS partner up or down traps occur only to communicate with that partner or server, and therefore may not occur when the other partner or server goes down or returns to service.

CLI Commands

To set the trap values for the DHCP server at the local cluster, use **dhcp set traps-enabled=***value* . You can also set the *default-free-address-config* attribute to the trap configuration. For example:

```
nrcmd> dhcp set traps-enabled=server-start,server-stop,free-address-low,free-address-high

nrcmd> dhcp set default-free-address-config=v4-trap-config
```

Note If you do not define a *default-free-address-config* (or *v6-default-free-address-config* for IPv6), Cisco Prime IP Express creates an internal, unlisted trap configuration named **default-aggregation-addr-trap-config**. Because of this, avoid using that name for a trap configuration you create.

To define trap configurations for DHCPv4 and DHCPv6, use **addr-trap** *name* **create** followed by the *attribute* =*value* pairs for the settings. For example:

```
nrcmd> addr-trap v4-trap-conf create mode=scope low-threshold=25% high-threshold=30%

nrcmd> addr-trap v6-trap-conf create mode=prefix low-threshold=20% high-threshold=25%
```

Handling SNMP Queries

You can use SNMP client applications to query the following MIBs:

- CISCO-DNS-SERVER-MIB.my
- CISCO-IETF-DHCP-SERVER-MIB.my
- CISCO-IETF-DHCP-SERVER-EXT-MIB.my
- CISCO-NETREG-DHCPV6-MIB.my (experimental)

When the SNMP server receives a query for an attribute defined in one of these MIBs, it returns a response PDU containing that attribute value. For example, using the NET-SNMP client application (available over the Internet), you can use one of these commands to obtain a count of the DHCPDISCOVER packets for a certain address:

```
C:\net-snmp5.2.2\bin>snmpget -m ALL -v 2c -c public
192.168.241.39:4444.iso.org.dod.internet.private.enterprises.cisco.ciscoExperiment.
ciscoIetfDhcpSrvMIB.ciscoIetfDhcpv4SrvMIBObjects.cDhcpv4Counters.cDhcpv4CountDiscovers
```

```
CISCO-IETF-DHCP-SERVER-MIB::cDhcpv4CountDiscovers.0 = Counter32: 0
C:\net-snmp5.2.2\bin>snmpget -m ALL -v 2c -c public
192.168.241.39:4444
 1.3.6.1.4.1.9.10.102.1.3.1

CISCO-IETF-DHCP-SERVER-MIB::cDhcpv4CountDiscovers.0 = Counter32: 0
```

Both commands return the same results. The first one queries the full MIB attribute name, while the second one queries its OID equivalent (which can be less error prone). As previously described, the OID equivalents of the MIB attributes are located in the relevant files at the following URL:

ftp://ftp.cisco.com/pub/mibs/oid/

For example, the CISCO-IETF-DHCP-SERVER-MIB.oid file includes the following OID definition that corresponds to the previous query example:

```
"cDhcpv4CountDiscovers" "1.3.6.1.4.1.9.10.102.1.3.1"
```

Here are some possible SNMP query error conditions:

- The community string sent in the request PDU does not match what you configured.
- The version in the request PDU is not the same as the supported version (SNMPv2).
- If the object being queried does not have an instance in the server, the corresponding variable binding type field is set to SNMP_NOSUCHINSTANCE. With a GetNext, if there is no next attribute, the corresponding variable binding type field is set to SNMP_ENDOFMIBVIEW.
- If no match occurs for the OID, the corresponding variable binding type field is set to SNMP_NOSUCHOBJECT. With a GetNext, it is set to SNMP_ENDOFMIBVIEW.
- If there is a bad value returned by querying the attribute, the error status in the response PDU is set to SNMP_ERR_BAD_VALUE.

Integrating Cisco PrimeIP Express SNMP into System SNMP

You can integrate the Cisco Prime IP Express SNMP server into the SNMP server for the system it runs on. The integration can be done in a way where the system will respond to queries for Cisco Prime IP Express MIB entries. On systems using NET-SNMP (and compatible servers) this is done by adding the following entries to the /etc/snmp/snmpd.conf configuration file

```
view systemview included .1.3.6.1.4.1.9.9
view systemview included .1.3.6.1.4.1.9.10

proxy -v 2c -c public 127.0.0.1:4444 .1.3.6.1.4.1.9.9
proxy -v 2c -c public 127.0.0.1:4444 .1.3.6.1.4.1.9.10
```

The community string **public** and the port number **4444** may have to be replaced if the Cisco Prime IP Express SNMP server has been configured with different values for those settings.

NET-SNMP is commonly available on Linux and other Unix-like systems. On other systems, similar mechanisms may also be available.

Bring Your Own Device Web Server

The BYOD web server at the regional cluster provides the infrastructure for Cisco Prime IP Express BYOD operation. The main purpose of the BYOD Web Server is to authenticate the user against AD and collect the device metadata by registering the user's own device in Cisco Prime IP Express.

Managing BYOD Web Server

You can view logs and startup logs; edit the server attributes.

To view logs and startup logs, in the regional cluster web UI, from the **Operate menu,** choose **Manage Servers** under the **Server** submenu to open the Manage Servers page.

Editing BYOD Web Server Properties

You can edit the BYOD web server properties using the Edit Local BYOD Web Server page.

Regional Basic or Advanced or Expert Web UI

Step 1 To access the BYOD web server properties, choose **Manage Servers** under **Operate** menu to open the Manage Servers page.

Step 2 Click Local **BYOD Web Server** in the Manage Servers pane on the left. The Edit Local BYOD Web Server page appears. This page displays the BYOD web server attributes.

- **KeyStore Settings:** Redirects the "http call" of the BYOD web server to secure "https" with a combination of key store file and key store password.
- **LDAP Settings:** Specifies the remote LDAP server used for client registration.
- **Additional Attributes (Auto- start)**: Indicates if the BYOD server should be started automatically after every server agent restart.

Step 3 Modify the settings as per your requirement.

Step 4 Click **Save** to save the BYOD web server attribute modifications.

Step 5 Click **Start Server** or **Restart Server** to apply the modifications to the BYOD web server.

Setting Up BYOD Theme and Content

You can create the content and multiple BYOD themes at the regional cluster which can be applied to BYOD web server interface.

Adding and Previewing BYOD Themes

You can create your own themes on the regional cluster using the BYOD Theme page and apply the created theme to the BYOD web server so that the logo, background, font, and other properties of the BYOD interface are displayed as per your customization. The created theme can be previewed prior to publishing it to the BYOD web server.

To add and preview a theme:

Regional Advanced or Expert Web UI

Step 1 From the **Deploy** menu, choose **Theme** under the **BYOD** submenu to open the List/Add Custom Theme page.

Step 2 Click the **Add Theme** icon in the Theme pane.
The **Add Custom Theme** window appears.

Step 3 Enter the Theme Name in the Add Custom Theme window.

Step 4 Click **Add Custom Theme** to create a new BYOD Theme.

Step 5 Update the Edit Custom Theme page with required theme attributes.

Step 6 Click the **Review Theme** icon in the top right corner of the List/Add Custom Theme page.
The Theme Preview window appears displaying the BYOD page with the newly added theme.

Note You can navigate between the BYOD pages with **Register** and **Reboot** to view how the theme is applied to the BYOD pages. By default, the **Theme preview** window loads the BYOD Device Registration page.

Step 7 Click **Reboot** to preview your theme in the Device Activation page.

Note You must close the Theme Preview window after preview to return to the List/Add Custom Theme page in the regional server.

Step 8 Click **Save** in the List/Add Custom Theme page in the regional server to apply the theme to the BYOD web server or click **Revert** to change the attribute values prior to saving the Custom Theme.

Note You can modify and preview the theme any number of times. Only the recently saved theme is applied to the BYOD web server.

Adding and Previewing BYOD Content

You can create the BYOD web server contents such as login page message, about, terms of services, contact details, and help message on the BYOD content page of the regional cluster, and preview it prior to publishing it to the BYOD web server. These contents can be published in the BYOD web server interface for the device registration and login pages.

To add and review content:

Regional Advanced or Expert Web UI

Step 1 From the **Deploy** menu, choose **Content** under the **BYOD** submenu to open the Edit BYOD content page.

Step 2 Upload the file or enter relevant text in the Edit BYOD content page.

Note You must upload only .html , .htm or .txt files.

Step 3 Click **Review** to preview the content in the Edit BYOD content page before saving. A **Content Review** window containing the contents appears.

Step 4 Click on**About**/**Terms of Service/Contact/Help** in the content review page to preview the content added in the EDIT BYOD content page of the regional server.

Step 5 Click **Save** to publish the added BYOD content to the BYOD web server.

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 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.

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 in Cisco PrimeIP Express 8.3 DHCP User Guide*).

Related Topics

Polling Process, on page 32

Adjusting the Polling Intervals, on page 32

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. **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 the table below, 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 the table below using the **ccm** command.

Note If lease history collection is not explicitly turned on at the local cluster DHCP server (see Enabling Lease History Collection, on page 33), no data is collected, even though polling is on by default.

Table 7: 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 7: Lease History Polling Regional Attributes , on page 33.

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 in Cisco PrimeIP Express 8.3 DHCP User Guide*. 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, on page 34

Pulling Scope Templates from Replica Data, on page 34

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 in Cisco PrimeIP Express 8.3 DHCP User Guide*. 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, on page 35

Pulling Policies from Replica Data, on page 36

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 the *"Managing Client-Classes and Clients" chapter in Cisco PrimeIP Express 8.3 DHCP User Guide*. 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, on page 37

Pushing Client-Classes to Local Clusters, on page 37

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 in Cisco PrimeIP Express 8.3 DHCP User Guide*. 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, on page 38

Pulling VPNs from Replica Data, on page 39

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 in Cisco PrimeIP Express 8.3 DHCP User Guide*. 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, on page 40
DHCPv6 Reservations, on page 40

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.

DHCPv6 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 in Cisco PrimeIP Express 8.3 DHCP User Guide*.

Monitoring Resource Limit Alarms

Resource limit alarms enable you to monitor Cisco Prime IP Express system resources and provide an indication when one or more product resources has entered potentially dangerous level and requires attention. Resource limit alarms are designed to convey the resource limit information in an organized and consolidated way.

Note The log messages related to resource limits are logged to the ccm_monitor_log files. For more information on log files, see Log Files.

You can reset the predefined threshold levels for both critical and warning levels for each monitored resource.

Cisco Prime IP Express reports the current status, the current value, and the peak value of the monitored resources in the web UI and CLI. The peak value is compared to the configured warning or critical limit for the resource limit alarm and the status of the resource limit alarm is displayed as OK, Warning, or Critical. Cisco Prime IP Express displays the alarms on the web UI and CLI until the resulting condition no longer occurs and the peak value is reset.

The resource limit alarms are updated at regular intervals based on the polling interval you configure. For more information on setting up the polling interval, see Setting Resource Limit Alarms Polling Interval, on page 43.

If SNMP traps are enabled for the resource limit alarms, Cisco Prime IP Express generates SNMP traps when the monitored resources exceed the critical or warning levels. SNMP traps are generated whenever the current value exceeds the configured warning or critical level.

The resource limit alarms can be configured both at the regional and in the local cluster. The resource limit alarms data is consolidated at the individual local cluster level. The resource limits alarms available on the regional cluster level pertain to only the regional cluster. The table below lists the types of resource limit alarms that are available on the regional or the local cluster.

Table 8: Resource Limit Alarms

	Regional Cluster	Local Cluster
Data Free Space in../Data Partition	✓	✓
Shadow Backup Time	✓	✓
CCM Memory	✓	✓
CNR Server Agent Memory	✓	✓
Tomcat Memory	✓	✓
DHCP Memory	x	✓
CDNS Memory	x	✓
DNS Memory	x	✓
SNMP Memory	✓	✓
Lease Count	x	✓
Zone Count	x	✓
Resource Records Count	x	✓

Configuring Resource Limit Alarm Thresholds

You can configure the warning and critical limits for the resource limit alarms using the **Edit CCM Server** page.

Local and Regional Web UI

Step 1 To access the CCM server properties, choose **Manage Servers** under the **Operate** menu to open the Manage Servers page.

Step 2 Click **Local CCM Server** in the Manage Servers pane on the left. The Edit Local CCM Server page appears. This page displays all the CCM server attributes.

Step 3 Click the **Configure Resource Limits** tab.

Step 4 Modify the settings as per your requirement.

Note To enable the SNMP traps for the resource limit alarms, select the Enable Traps option in the Trap Configuration group.

Step 5 Click **Save** to save the CCM server attribute modifications.

CLI Commands

To set the resource limit alarms on the local or regional cluster, use **resource set** *attribute=value*. Use **resource show** to review the current setting and use **resource report** ***<all | full| level>*** command to report on the resources.

To view the defined warning and critical levels, use **resource report** levels command.

A 109 status message is reported (if at least one resource is in the critical or warning state) under the following scenarios.

- Execute 'resource report' command.
- Connect to a cluster via CLI.
- Exit from CLI.

Setting Resource Limit Alarms Polling Interval

You can set how often Cisco Prime IP Express polls for alarm data from the server and updates the web UI data. The *stats-history-sample-interval* controls the CCM server system polling rate.

Step 1 To edit the alarm poll interval, you need to edit the user preferences by going to **User Preferences** under the **admin** menu (at the top of the main page).

Step 2 After making the user preference settings, click **Modify User Preferences**.

Viewing Resource Limit Alarms

Resource limit alarms are displayed on the Alarms toolbar. To see a summary of the alarms, in the Cisco Prime IP Express web UI, click the **Alarms** toolbar on the bottom of the web UI. This opens the Alarms toolbar overlay which displays the status, resource values (current, configured warning, and critical value), and the peak value for each resource limit alarm. Based on the peak value for each resource limit, the status of resource limit is displayed as OK, Warning, or Critical on the web UI and CLI. The alarms are updated at regular intervals based on the polling interval you configure. For more information on setting up the polling interval, see Setting Resource Limit Alarms Polling Interval, on page 43.

Note When a resource is in a warning or critical state, the resource limit alarm is also displayed on the Configuration Summary page.

Resetting Resource Limit Alarms Peak Value

Cisco Prime IP Express maintains the peak values for each resource limit. The peak value is updated only when the current value exceeds the peak value. The peak value is compared to the configured warning or critical limit for the resource limit alarm and the status of the resource limit alarm is displayed as OK, Warning, or Critical.

When the peak value exceeds the configured warning or critical limit the status of the resource limit alarm is shown as Warning or Critical (on the web UI and CLI) respectively until the peak value is explicitly reset. To reset the peak value, perform the following steps:

Step 1 On the **Alarms** toolbar, select the Alarm for which you want to reset the peak value.

Step 2 Click **Reset Alarm** to clear the peak value.

CLI Commands

To reset the peak value on the local or regional cluster, use **resource reset** *name*.

Note If no resource name is provided, all are reset.

Export Resource Limit Alarms Data

You can export the resource limit alarms data to a CSV file. To export the resource limit alarms:

Step 1 Click **Alarms** in the alarms toolbar at the bottom of the web UI.

Step 2 Click **Export to CSV**.

Step 3 The File Download pop-up window displays. Click **Save**.

Step 4 In the Save As pop-up window, choose the location you want to save the file to and click **Save**.

Local Cluster Management Tutorial

This tutorial describes a basic scenario on a local cluster of the Example Company. Administrators at the cluster are responsible for users, zone data, DHCP data, address space data, and the servers in general. The task is to set up two zones (example.com and boston.example.com), hosts in the zones, and a subnet. The local cluster must also create a special administrator account so that the regional cluster in San Jose can perform the central configuration and replicate the local cluster administrators and address space at another cluster, as described in Regional Cluster Management Tutorial, on page 52.

Related Topics

Administrator Responsibilities and Tasks, on page 45
Create the Administrators, on page 45
Create the Address Infrastructure, on page 46
Create the Zone Infrastructure, on page 47
Create a Host Administrator Role with Constraints, on page 49
Create a Group to Assign to the Host Administrator, on page 51
Test the Host Address Range, on page 51

Administrator Responsibilities and Tasks

The local cluster administrators have the following responsibilities and tasks:

- **example-cluster-admin—**Created by the superuser:
  - At the Boston cluster, creates the other local administrators (example-zone-admin and example-host-admin).
  - Creates the basic network infrastructure for the local clusters.
  - Constrains the example-host-role to an address range in the boston.example.com zone.
  - Creates the example-host-group (defined with the example-host-role) that the example-zone-admin will assign to the example-host-admin.
- **example-zone-admin**:
  - Creates the example.com and boston.example.com zones, and maintains the latter zone.
  - Assigns the example-host-group to the example-host-admin.
- **example-host-admin—**Maintains local host lists and IP address assignments.

Create the Administrators

For this example, the superuser in Boston creates the local cluster, zone, and host administrators, as described in the Administrator Responsibilities and Tasks, on page 45.

Local Basic Web UI

Step 1 At the Boston local cluster, log in as superuser (usually **admin**).

Step 2 In Basic mode, from the **Administration** menu, choose **Administrators**.

Step 3 Add the local cluster administrator (with superuser access)—On the List/Add Administrators page:

a) Click the **Add Administrators** icon in the Administrators pane, enter **example-cluster-admin** in the Name field.
b) Enter **exampleadmin** in the Password field, then click **Add Admin**.
c) Check the Superuser check box.
d) Do not choose a group from the Groups list.
e) Click **Save**.

Step 4 Add the local zone administrator on the same page:

a) Click the **Add Administrators** icon in the Administrators pane, enter **example-zone-admin** in the Name field, and **examplezone** in the Password field, then click **Add Admin**.
b) Multiselect **ccm-admin-group**, **dns-admin-group**, and **host-admin-group** in the Groups drop-down list. The dns-admin-group is already predefined with the dns-admin role to administer DNS zones and servers. The ccm-admin-group guarantees that the example-zone-admin can set up the example-host-admin with a constrained role later on. The host-admin-group is mainly to test host creation in the zone.
c) Click **Save**.

Step 5 Add the local host administrator on the same page:

a) Click the **Add Administrators** icon in the Administrators pane, enter **example-host-admin** in the Name field, and **examplehost** in the Password field, then click **Add Admin**.
b) Do not choose a group at this point. (The example-zone-admin will later assign example-host-admin to a group with a constrained role.)
c) Click **Save**.

Note For a description on how to apply constraints to the administrator, see the Create a Host Administrator Role with Constraints, on page 49.

Create the Address Infrastructure

A prerequisite to managing the zones and hosts at the clusters is to create the underlying network infrastructure. The network configuration often already exists and was imported. However, this tutorial assumes that you are starting with a clean slate.

The local example-cluster-admin next creates the allowable address ranges for the hosts in the boston.example.com zone that will be assigned static IP addresses. These addresses are in the 192.168.50.0/24 subnet with a range of hosts from 100 through 200.

Local Advanced Web UI

Step 1 At the local cluster, log out as superuser, then log in as the **example-cluster-admin** user with password **exampleadmin**. Because the administrator is a superuser, all features are available.

Step 2 Click **Advanced** to go to Advanced mode.

Step 3 Click **Design**, then **Subnets** under DHCPv4 submenu.

Step 4 On the List/Add Subnets page, enter the boston.example.com subnet address:

a) Click the **Add Subnets** icon in the Subnets pane, enter **192.168.50** in the Address field.
b) Choose **24** in the mask drop-down list—This subnet will be a normal Class C network.
c) Leave the Owner, Region, and Address Type fields as is. Add description if desired.

d) Click **Add Subnet**.

Step 5 Click the 192.168.50.0/24 address to open the Edit Subnet page.

Step 6 In the IP Ranges fields, enter the static address range:

a) Enter **100** in the Start field. Tab to the next field.
b) Enter **200** in the End field.
c) Click **Add IP Range**. The address range appears under the fields.

Step 7 Click **Save**.

Step 8 Click **Address Space** to open the View Unified Address Space page. The 192.168.50.0/24 subnet should appear in the list. If not, click the **Refresh** icon.

Create the Zone Infrastructure

For this scenario, example-cluster-admin must create the Example Company zones locally, including the example.com zone and its subzones. The example-cluster-admin also adds some initial host records to the boston.example.com zone.

Related Topics

Create the Forward Zones, on page 47
Create the Reverse Zones, on page 48
Create the Initial Hosts, on page 48

Create the Forward Zones

First, create the example.com and boston.example.com forward zones.

Local Basic Web UI

Step 1 At the local cluster, log in as the **example-zone-admin** user with password **examplezone**.

Step 2 From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu. This opens the List/Add Forward Zones page.

Step 3 Create the example.com zone (tab from field to field):

a) Click the **Add Forward Zone** icon in the Forward Zones pane, enter **example.com** in the Name field.
b) In the Nameserver FQDN field, enter **ns1**.
c) In the Contact E-Mail field, enter **hostmaster**.
d) In the Serial Number field, enter the serial number.
e) Click **Add Zone**.

Step 4 Create the **boston.example.com** zone in the same way, using the same values as in the previous steps:

a) Creating a zone with a prefix added to an existing zone opens the Create Subzone in Parent Zone page, because the zone can be a potential subzone. Because you do want to create this zone as a subzone to example.com, click **Create as Subzone** on the Create Subzone in Parent Zone page.
b) Because nameservers are different in each zone, you must create a glue Address (A) record to tie the zones together. Enter 192.168.50.1 in the A record field, then click **Specify Glue Records**. Then click **Report**, **Run**, and **Return**.
c) The List/Add Zones page should now list example.com and boston.example.com.

Step 5 Click **Advanced**, then **Show Forward Zone Tree** to show the hierarchy of the zones. Return to list mode by clicking **Show Forward Zone List**.

Create the Reverse Zones

Next, create the reverse zones for example.com and boston.example.com. This way you can add reverse address pointer (PTR) records for each added host. The reverse zone for example.com is based on the 192.168.50.0 subnet; the reverse zone for boston.example.com is based on the 192.168.60.0 subnet.

Local Basic Web UI

Step 1 At the local cluster, you should be logged in as the example-zone-admin user, as in the previous section.
Step 2 From the **Design** menu, choose **Reverse Zones** under the **Auth DNS** submenu.
Step 3 On the List/Add Reverse Zones page, click the **Add Reverse Zone** icon in the Reverse Zones pane, enter **50.168.192.in-addr.arpa** in the Name field. (There is already a reverse zone for the loopback address, 127.in-addr.arpa.)
Step 4 Enter the required fields to create the reverse zone, using the forward zone values:
a) **Nameserver—**Enter **ns1.example.com.** (be sure to include the trailing dot).
b) **Contact E-Mail—**Enter **hostmaster.example.com.** (be sure to include the trailing dot).
c) **Serial Number—**Enter the serial number.
Step 5 Click **Add Reverse Zone** to add the zone and return to the List/Add Reverse Zones page.
Step 6 Do the same for the boston.example.com zone, using **60.168.192.in-addr.arpa** as the zone name and the same nameserver and contact e-mail values as in **Step 4**. (You can cut and paste the values from the table.)

Create the Initial Hosts

As a confirmation that hosts can be created at the Boston cluster, the example-zone-admin tries to create two hosts in the example.com zone.

Local Advanced Web UI

Step 1 As the example-zone-admin user, click **Advanced** to enter Advanced mode.

Step 2 From the **Design** menu, choose **Hosts** under the **Auth DNS** submenu. This opens the List/Add Hosts for Zone page. You should see boston.example.com and example.com in the Select Zones box on the left side of the window.

Step 3 Click example.com in the list of zones.

Step 4 Add the first static host with address 192.168.50.101:

a) Enter **userhost101** in the Name field.
b) Enter the complete address **192.168.50.101** in the IP Address(es) field. Leave the IPv6 Address(es) and Alias(es) field blank.
c) Ensure that the Create PTR Records? check box is checked.
d) Click **Add Host**.

Step 5 Add the second host, **userhost102**, with address **192.168.50.102**, in the same way. The two hosts should now appear along with the nameserver host on the List/Add Hosts for Zone page.

Create a Host Administrator Role with Constraints

In this part of the tutorial, the Boston example-cluster-admin creates the example-host-role with address constraints in the boston.example.com zone.

Local Advanced Web UI

Step 1 Log out as the example-zone-admin user and log in as the **example-cluster-admin** user (with password **exampleadmin**).

Step 2 Click **Advanced** to enter Advanced mode.

Step 3 From the **Administration** menu, choose **Roles** under User Access submenu to open the List/Add Administrator Roles page.

Step 4 Add the example-host-role:

a) Click the **Add Role** icon in the Roles pan to open the Add Roles dialog box.
b) Enter **example-host-role** in the Name field.
c) Click **Add Role**. The example-host-role should now appear in the list of roles on the List/Add Administrator Roles page.

Step 5 Add the constraint for the role:

a) Click **Add Constraint**.
b) On the Add Role Constraint for Role page, scroll down to Host Restrictions.
c) For the *all-forward-zones* attribute, click the **false** radio button.
d) For the *zones* attribute, enter **boston.example.com**.
e) For the *ipranges* attribute, enter the range **192.168.50.101–192.168.50.200**

f) The *zone-regexpr* and *host-regexpr* attribute fields are for entering regular expressions to match zones and hosts, respectively, in regex syntax. (See the following table for the commonly used regex values.)

Table 9: Common Regex Values

Value	Matches
. (dot)	Any character (a wildcard). Note that to match a literal dot character (such as in a domain name), you must escape it by using a backslash (\), such that **\.com** matches .com.
char	Literal character (*char*) that follows, or the *char* has special meaning. Used especially to escape metacharacters such as the dot (**.**) or another backslash. Special meanings include **\d** to match decimal digits, **\D** for nondigits, **\w** for alphanumerics, and **\s** for whitespace.
*char***?**	Preceding *char* once or not at all, as if the character were optional. For example, **example\.?com** matches example.com or examplecom.
*char******	Preceding *char* zero or more times. For example, **ca*t** matches ct, cat, and caaat. This repetition metacharacter does iterative processing with character sets (see **[***charset***]**).
*char***+**	Preceding *char* one or more times. For example, **ca+t** matches cat and caaat (but not ct).
[*charset***]**	Any of the characters enclosed in the brackets (a character set). You can include character ranges such as **[a–z]** (which matches any lowercase character). With the ***** repetition metacharacter applied, the search engine iterates through the set as many times as necessary to effect a match. For example, **a[bcd]*b** will find abcbd (by iterating through the set a second time). Note that many of the metacharacters (such as the dot) are inactive and considered literal inside a character set.
[^*charset***]**	Anything but the *charset* , such that **[^a-zA-Z0-9]** matches any nonalphanumeric character (which is equivalent to using **\W**). Note that the caret outside a character set has a different meaning.
^	Beginning of a line.
$	End of a line.

g) Click **Add Constraint**. The constraint should have an index number of 1.

Step 6 Click **Save**.

Create a Group to Assign to the Host Administrator

The Boston example-cluster-admin next creates an example-host-group that includes the example-host-role so that the example-zone-admin can assign this group to the example-host-admin.

Local Advanced Web UI

Step 1 As example-cluster-admin, still in Advanced mode, from the **Administration** menu, choose **Groups** submenu to open the List/Add Administrator Groups page.

Step 2 Create the example-host-group and assign the example-host-role to it:

a) Click the **Add Groups** icon in the Groups pane, enter **example-host-group** in the Name field.
b) From the Base Role drop-down list, choose **example-host-role**.
c) Click **Add Group**.
d) Add a description such as **Group for the example-host-role**, then click **Save**.

Step 3 Log out as example-cluster-admin, then log in as the **example-zone-admin** user (with password **examplezone**).

Step 4 As example-zone-admin, assign the example-host-group to the example-host-admin:

a) In Basic mode, from the **Administration** menu, choose **Administrators**.
b) On the List/Add Administrators page, click example-host-admin to edit the administrator.
c) On the Edit Administrator page, choose **example-host-group** in the Available list, then click << to move it to the Selected list.
d) Click **Save**. The example-host-admin should now show the example-host-group in the Groups column on the List/Add Administrators page.

Test the Host Address Range

The example-host-admin next tests an out-of-range address and then adds an acceptable one.

Local Advanced Web UI

Step 1 At the local cluster, log out as example-zone-admin, then log in as **example-host-admin** (with password **examplehost**).

Step 2 Click **Advanced** to enter Advanced mode.

Step 3 From the **Design** menu, choose **Hosts** from the **Auth DNS** submenu.

Step 4 On the List/Add Hosts for Zone page, try to enter an out-of-range address (note the range of valid addresses in the Valid IP Ranges field):

a) Enter **userhost3** in the Name field.
b) Deliberately enter an out-of-range address (**192.168.50.3**) in the IP Address(es) field.
c) Click **Add Host**. You should get an error message.

Step 5 Enter a valid address:

a) Enter **userhost103**.
b) Enter **192.168.50.103** in the IP Address(es) field.
c) Click **Add Host**. The host should now appear with that address in the list.

Regional Cluster Management Tutorial

This tutorial is an extension of the scenario described in the Local Cluster Management Tutorial, on page 44. In the regional cluster tutorial, San Jose has two administrators—a regional cluster administrator and a central configuration administrator. Their goal is to coordinate activities with the local clusters in Boston and Chicago so as to create DNS zone distributions, router configurations, and DHCP failover configurations using the servers at these clusters. The configuration consists of:

- One regional cluster machine in San Jose.
- Two local cluster machines, one in Boston and one in Chicago.
- One Cisco uBR7200 router in Chicago.

Related Topics

Administrator Responsibilities and Tasks, on page 53
Create the Regional Cluster Administrator, on page 53
Create the Central Configuration Administrator, on page 53
Create the Local Clusters, on page 54
Add Zone Management to the Configuration Administrator, on page 55
Create a Zone for the Local Cluster, on page 55
Pull Zone Data and Create a Zone Distribution, on page 56
Create a Subnet and Pull Address Space, on page 56

Push a DHCP Policy, on page 57
Create a Scope Template, on page 58
Create and Synchronize the Failover Pair, on page 58

Administrator Responsibilities and Tasks

The regional administrators have the following responsibilities and tasks:

- **example-regional-admin—**Created by the superuser at the San Jose regional cluster, who creates the example-cfg-admin.
- **example-cfg-admin**:
  - Defines the Boston and Chicago clusters and checks connectivity with them.
  - Adds a router and modifies a router interface.
  - Pulls zone data from the local clusters to create a zone distribution.
  - Creates a subnet and policy, and pulls address space, to configure DHCP failover pairs in Boston and Chicago.

Create the Regional Cluster Administrator

The regional superuser first creates the example-regional-administrator, defined with groups, to perform cluster and user administration.

Regional Web UI

Step 1 Log into the regional cluster as superuser.

Step 2 From the **Administration** menu, choose **Administrators** to open the List/Add Administrators page for the local cluster version of this page, which is essentially identical.

Step 3 Click the Add Administrators icon in the Administrators pane, enter **example-regional-admin** in the Name field, then **examplereg** in the Password field in the Add Administrator dialog box, then click Add Administrator.

Step 4 Multiselect **central-cfg-admin-group** (for cluster administration) and **regional-admin-group** (for user administration) in the Groups drop-down list.

Step 5 Click **Save**.

Create the Central Configuration Administrator

As part of this tutorial, the example-regional-admin next logs in to create the example-cfg-admin, who must have regional configuration and address management capabilities.

Regional Web UI

Step 1 Log out as superuser, then log in as **example-regional-admin** with password **examplereg**. Note that the administrator has all but host and address space administration privileges.

Step 2 From the **Administration** menu, choose **Administrators** to open the List/Add Administrators page.

Step 3 Click the **Add Administrators** icon in the Administrators pane, enter **example-cfg-admin** in the Name field, then **cfgadmin** in the Password field in the Add Administrator dialog box, then click **Add Administrator**.

Step 4 Multiselect **central-cfg-admin-group** and **regional-addr-admin-group** in the Groups drop-down list.

Step 5 Click **Save**. The example-cfg-admin now appears with the two groups assigned.
You can also add constraints for the administrator. Click **Add Constraint** and, on the Add Role Constraint for Role page, choose the read-only, owner, or region constraints, then click **Add Constraint**.

Create the Local Clusters

The example-cfg-admin next creates the two local clusters for Boston and Chicago.

Regional Web UI

Step 1 Log out as example-regional-admin, then log in as **example-cfg-admin** with password **cfg admin**.

Step 2 From the **Operate** menu, choose **Servers from the Manage Clusters** submenu to open the List/Add Remote Clusters page.

Step 3 Click **the Add Manage Clusters** icon in the**Manage Clusters** pane.

Step 4 On the Add Cluster dialog box, create the Boston cluster based on data provided by its administrator:

a) Enter **Boston-cluster** in the name field.
b) Enter the IP address of the Boston server in the ipaddr field.
c) Enter **example-cluster-admin** in the admin field, then **exampleadmin** in the password field.
d) Enter in the SCPO-port field the SCP port to access the cluster as set at installation (**1234** is the preset value).
e) Click **Add Cluster**.

Step 5 Create the Chicago cluster in the same way, except use **Chicago-cluster** in the name field, enter the remaining values based on data provided by the Chicago administrator, then click **Add Cluster**. The two clusters should now appear on the List/Add Remote Clusters page.

Step 6 Connect to the Boston cluster. Click the **Go Local** icon next to Boston-cluster. If this opens the local cluster Manage Servers page, this confirms the administrator connectivity to the cluster. To return to the regional cluster web UI, click the **Go Regional** icon.

Step 7 Connect to the Chicago cluster to confirm the connectivity in the same way.

Step 8 Confirm that you can replicate data for the two forward zones from the Boston cluster synchronization:

a) From the **Operate** menu, choose **Replica Data** from the **Servers** submenu.
b) On the View Replica Class List page, click Boston-cluster in the Select Cluster list.

c) In the Select Class list, click **Forward Zones**.
d) Click the **Replicate** icon in the Replicate Data column.
e) Click **View Replica Class List**. On the List Replica Forward Zones for Cluster page, you should see the boston.example.com and example.com zones.

Add Zone Management to the Configuration Administrator

Because there are no zones set up at the Chicago cluster, the example-cfg-admin can create a zone at the regional cluster to make it part of the zone distribution. However, the example-regional-admin must first modify the example-cfg-admin to be able to create zones.

Regional Web UI

Step 1 Log out as example-cfg-admin, then log in as **example-regional-admin**.
Step 2 From the **Administration** menu, choose **Administrators**.
Step 3 On the List/Add Administrators page, click example-cfg-admin from the Administrators pane.
Step 4 On the Edit Administrator page, click central-dns-admin-group in the Groups Available list, then move it (using <<) to the Selected list. The Selected list should now have central-cfg-admin-group, regional-addr-admin-group, and central-dns-admin-group.
Step 5 Click **Save**. The change should be reflected on the List/Add Administrators page.

Create a Zone for the Local Cluster

The example-cfg-admin next creates the chicago.example.com zone for the zone distribution with the Boston and Chicago zones.

Regional Web UI

Step 1 Log out as example-regional-admin, then log in as **example-cfg-admin**.
Step 2 From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu.
Step 3 Click the **Add Forward Zones** icon in the **Forward Zones** pane.
Step 4 On the Add Zone dialog box, enter:
a) **Name—chicago.example.com**.
b) **Nameserver FQDN—ns1**.
c) **Contact E-mail—hostmaster**.
d) **Nameservers—ns1** (click **Add Nameserver**).

e) Click **Add Zone**.

Step 5 Click the **Reverse Zones** submenu.

Step 6 On the List Reverse Zones page, create the **60.168.192.in-addr.arpa** reverse zone for the Chicago zone, with the proper attributes set.

Pull Zone Data and Create a Zone Distribution

The example-cfg-admin next pulls zone data from Boston and Chicago and creates a zone distribution.

Regional Web UI

Step 1 As example-cfg-admin, from the **Design** menu, choose **Views** under the **Auth DNS** submenu to view the List/Add Zone Views page.

Step 2 On the List/Add Zone Views page, pull the zone from the replica database:

a) Click the **Pull Replica** icon in the **Views** pane.
b) On the Select Replica Downsize Data to Pull dialog box, leave the Data Synchronization Mode defaulted as Update, then click **Report** to open the Report Pull Replica Zone Data page.
c) Notice the change sets of data to pull, then click **Run**.
d) On the Run Pull Replica Zone Data page, click **OK**.

Step 3 On the List/Add Zone Views page, notice that the Boston cluster zone distribution is assigned an index number (**1**) in the Name column. Click the number.

Step 4 On the Edit Zone Views page, in the Primary Server field, click Boston-cluster. (The IP address of the Boston-cluster becomes the first master server in the Master Servers list.)

Step 5 Because we want to make the Chicago-cluster DNS server a secondary server for the Boston-cluster:

a) Click **Add Server** in the Secondary Servers area.
b) On the Add Zone Distribution Secondary Server page, choose **Chicago-cluster** in the Secondary Server drop-down list.
c) Click **Add Secondary Server**.

Step 6 On the Edit Zone Distribution page, in the Forward Zones area, move **chicago.example.com** to the Selected list.

Step 7 In the Reverse Zones area, move **60.168.192.in-addr.arpa** to the Selected list.

Step 8 Click **Modify Zone Distribution**.

Create a Subnet and Pull Address Space

The example-cfg-admin next creates a subnet at the regional cluster. This subnet will be combined with the other two pulled subnets from the local clusters to create a DHCP failover server configuration.

Regional Web UI

Step 1 As example-cfg-admin, from the **Design** menu,choose **Subnets** under the **DHCPv4** submenu to open the List/Add Subnets page.

Step 2 Create an additional subnet, 192.168.70.0/24 by clicking the **Add Subnets** icon in the Subnets pane:

a) Enter **192.168.70** (the abbreviated form) as the subnet network address in the Address/Mask field.
b) Leave the **24** (255.255.255.0) selected as the network mask.
c) Click **Add Subnet**.

Step 3 Click **Address Space** to confirm the subnet you created.

Step 4 On the View Unified Address Space page, click **Pull Replica Address Space**.

Step 5 On the Select Pull Replica Address Space page, leave everything defaulted, then click **Report**.

Step 6 The Report Pull Replica Address Space page should show the change sets for the two subnets from the clusters. Click **Run**.

Step 7 Click **OK**. The two pulled subnets appear on the List/Add Subnets page.

Push a DHCP Policy

The example-cfg-admin next creates a DHCP policy, then pushes it to the local clusters.

Regional Web UI

Step 1 As example-cfg-admin, from the **Design** menu, choose **Policies** under the**DHCP Settings** submenu.

Step 2 On the List/Add DHCP Policies page, click the **Add Policies** icon in the**Policies** pane.

Step 3 On the Add DHCP Policy dialog box, create a central policy for all the local clusters:

a) Enter **central-policy-1** in the Name field. Leave the Offer Timeout and Grace Period values as is.
b) Enter a lease period. In the DHCP > DHCPv4 > Options drop-down list, choose **dhcp-lease-time [51] (unsigned time)**, then enter **2w** (two weeks) for the lease period in the Value field.
c) Click **Add Option**.
d) Click **Add Policy**. The central-policy-1 should appear on the List/Add DHCP Policies page.

Step 4 Push the policy to the local clusters:

a) Select the policy, central-policy-1 and click the **Push** button.
b) On the Push DHCP Policy Data to Local Clusters page, leave the Data Synchronization Mode as **Ensure**. This ensures that the policy is replicated at the local cluster, but does not replace its attributes if a policy by that name already exists.
c) Click **Select All** in the Destination Clusters section of the page.
d) Click **<<** to move both clusters to the Selected field.
e) Click **Push Data to Clusters**.

f) View the push operation results on the View Push DHCP Policy Data Report page, then click **OK**.

Create a Scope Template

The example-cfg-admin next creates a DHCP scope template to handle failover server pair creation.

Regional Web UI

Step 1 As the example-cfg-admin user, from the Design menu, choose **Scope Templates** under the**DHCPv4** submenu.

Step 2 On the List/Add DHCP Scopes page, click the **Add Scopes** icon in the **Scope Templates** pane. Enter **scope-template-1** in the Name field, then click **Add Scope Templates**.

Step 3 The template should appear on the List/Add DHCP Scopes page. Set the basic properties for the scope template**—**Enter or choose the following values in the fields:

a) **Scope Name Expression—**To autogenerate names for the derivative scopes, concatenate the example-scope string with the subnet defined for the scope. To do this, enter **(concat** "**example-scope-**" **subnet)** in the field (including the parentheses).

b) **Policy—**Choose **central-policy-1** in the drop-down list.

c) **Range Expression—**Create an address range based on the remainder of the subnet (the second through last address) by entering **(create-range 2 100)**.

d) **Embedded Policy Option Expression—**Define the router for the scope in its embedded policy and assign it the first address in the subnet by entering **(create-option** "**routers**" **(create-ipaddr subnet 1))**.

Step 4 Click **Save**.

Create and Synchronize the Failover Pair

The example-cfg-admin next creates the failover server pair relationship and synchronizes the failover pair. The DHCP server at Boston becomes the main, and the server at Chicago becomes the backup.

Regional Web UI

Step 1 As the example-cfg-admin user, from the **Deploy** menu, choose **Failover** under the **DHCP** submenu.

Step 2 On the List/Add DHCP Failover Pairs page, click the **Add Failover Pair** icon in the **Failover Pairs** pane.

Step 3 On the Add DHCP Failover Pair dialog box, enter or choose the following values:

a) **Failover Pair Name—**Enter **central-fo-pair**.

b) **Main Server—**Click **Boston-cluster**.

c) **Backup Server—**Click **Chicago-cluster**.

d) **Scope Template—**Click **scopetemplate-1**
e) Click **Add Failover Pair**.

Step 4 Synchronize the failover pair with the local clusters:

a) On the List/Add DHCP Failover Pairs page, click the **Report** icon in the Synchronize column.
b) On the Report Synchronize Failover Pair page, accept **Local Server** as the source of network data.
c) Accept **Main to Backup** as the direction of synchronization.
d) Accept the operation **Update**.
e) Click **Report** at the bottom of the page.
f) On the View Failover Pair Sync Report page, click **Run Update**.
g) Click **Return**.

Step 5 Confirm the failover configuration and reload the server at the Boston cluster:

a) On the List/Add DHCP Failover Pairs page, click the **Go Local** icon next to Boston-cluster.
b) On the Manage DHCP Server page, click the **Reload** icon.
c) Click the **Go Regional** icon at the top of the page to return to the regional cluster.

Step 6 Confirm the failover configuration and reload the server at the Chicago cluster in the same way.