



# Maintaining Servers and Databases

---

This chapter explains how to administer and control your local and regional server operations.

- [Managing Servers, on page 1](#)
- [Scheduling Recurring Tasks, on page 3](#)
- [Logs, on page 5](#)
- [Running Data Consistency Rules, on page 10](#)
- [Monitoring and Reporting Server Status, on page 13](#)
- [Troubleshooting DHCP and DNS Servers, on page 30](#)
- [Using the TAC Tool, on page 34](#)
- [Troubleshooting and Optimizing the TFTP Server, on page 37](#)

## Managing Servers

If you are assigned the server-management subrole of the ccm-admin role, you can manage the Cisco Prime Network Registrar servers as follows:

- **Start**—Load the database and start the server.
- **Stop**—Stop the server.
- **Reload**—Stop and restart the server. (Note that you do not need to reload the server for all RR updates, even protected RR updates. For details, see the *"Managing DNS Update"* chapter in *Cisco Prime Network Registrar 10.1 DHCP User Guide*.)
- **Check statistics**—See the [Displaying Statistics, on page 15](#).
- **View logs**—See the [Searching the Logs, on page 8](#).
- **Manage interfaces**—See the specific protocol pages for how to manage server interfaces.

Starting and stopping a server is self-explanatory. When you reload the server, Cisco Prime Network Registrar performs three steps—stops the server, loads configuration data, and restarts the server. Only after you reload the server does it use your changes to the configuration.



---

**Note** The CDNS, DNS, DHCP, and SNMP servers are enabled by default to start on reboot. The TFTP server is not enabled by default to start on reboot. You can change this using `[server] type enable` or `disable start-on-reboot` in the CLI.

---



**Note** If *exit-on-stop* attribute of DHCP, DNS, or TFTP server is enabled, then the statistics and scope utilization data only from the last start (reload) is reported while if the attribute is disabled, information across reloads is displayed.

## Local Basic or Advanced and Regional Web UI

You can manage the protocol servers in the following ways depending on if you are a:

- **Local or regional cluster administrator**—Choose **Manage Servers** from the **Operate** menu to open the Manage Servers page.

The local and regional cluster web UI access to server administration is identical, even though the available functions are different. As a regional administrator, you can check the state and health of the regional CCM server and server agent. However, you cannot stop, start, reload, or view statistics, logs, or interfaces for them.

At the local cluster, to manage the DHCP, DNS, CDNS, TFTP, or SNMP servers, select the server in the Manage Servers pane and do any of the following:

- Click the **Statistics** tab to view statistics for the server. (See the [Displaying Statistics, on page 15.](#))
- Click the **Logs** tab in the **View Log** column to view the log messages for the server. (See the [Searching the Logs, on page 8.](#))
- Click the **Start Server** button to start the server.
- Click the **Stop Server** button stop the server.
- Click the **Restart Server** button to reload the server.

- **Local cluster DNS administrator**—Choose **DNS Server** from the **Deploy** menu to open the Manage DNS Authoritative Server page.

Along with the Statistics, Startup Logs, Logs, HA DNS Server Status, Start Server, Stop Server, and Restart Server functions, you can also perform other functions when you click the **Commands** button to open the DNS Commands dialog box.

The server command functions are:

- **Forcing all zone transfers** (see the *"Enabling Zone Transfers" section in Cisco Prime Network Registrar 10.1 Authoritative and Caching DNS User Guide*)—Click the **Run** icon. This is the equivalent of **dns forceXfer secondary** in the CLI.
- **Scavenging all zones** (see the *"Scavenging Dynamic Records" section in Cisco Prime Network Registrar 10.1 DHCP User Guide*)—Click the **Run** icon. This is the equivalent of **dns scavenge** in the CLI.

- **Local cluster Caching DNS server**—Choose **CDNS Server** from the **Deploy** menu to open the Manage DNS Caching Server page.

Along with the Statistics, Startup Logs, Logs, Start Server, Stop Server, and Restart Server functions, you can also perform other functions when you click the **Commands** button to open the CDNS Commands dialog box.

In Advanced and Expert modes, you can flush Caching CDNS cache and flush the resource records. Click the **Commands** button to execute the commands.

- **Local cluster DHCP administrator**—Click **DHCP Server** from the **Deploy** menu to open the Manage DHCP Server page.

Along with the Statistics, Startup Logs, Logs, Start Server, Stop Server, and Restart Server functions, you can also perform other functions when you click the **Commands** button to open the DHCP Server Commands dialog box.

This page provides the Get Leases with Limitation ID feature, to find clients that are associated through a common limitation identifier (see the *"Administering Option 82 Limitation" section in Cisco Prime Network Registrar 10.1 DHCP User Guide*). Enter at least the IP address of the currently active lease in the IP Address field, then click the **Run** icon. You can also enter the limitation ID itself in the form *nn:nn:nn* or as a string ("*nnnn*"), in which case the IP address becomes the network in which to search. This function is the equivalent of `dhcp limitationList ipaddress [limitation-id] show` in the CLI.

## CLI Commands

In the CLI, the regional cluster allows CCM server management only:

- To start the server, use `server type start` (or simply `type start`; for example, `dhcp start`).
- To stop the server, use `server type stop` (or simply `type stop`; for example, `dhcp stop`). If stopping the server, it is advisable to save it first using the `save` command.
- To reload the server, use `server type reload` (or simply `type reload`; for example, `dhcp reload`). Cisco Prime Network Registrar stops the server you chose, loads the configuration data, and then restarts the server.
- To set or show attributes for the server, use `[server] type set attribute=value` or `[server] type show`. For example:

```
nrcmd> ccm set ipaddr=192.168.50.10
```

## Scheduling Recurring Tasks

In Basic and Advanced user mode in the local cluster web UI, you can schedule a number of recurring tasks. These tasks are:

- Reloading the DHCP server.
- Reloading the DNS server.
- Synchronizing DHCP failover server pairs:
  - If in staged dhcp edit mode, reload the main DHCP server.
  - Synchronize the failover configuration to the backup DHCP server.
  - If in staged dhcp edit mode, reload the backup DHCP server.
- Synchronizing High-Availability (HA) DNS server pairs:
  - If in staged dhcp edit mode, reload the main DNS server.
  - Synchronize the HA DNS configuration to the backup DNS server.
  - If in staged dhcp edit mode, reload the backup DNS server.
- Synchronizing zone distribution maps:
  - If in staged dhcp edit mode, reload the main DNS server.

- If in staged dhcp edit mode, reload the backup HA DNS server.
- Synchronize the zone distribution maps.
- If in staged dhcp edit mode, reload the secondary DNS server or servers.

## Local Basic or Advanced Web UI

To set up one or more of these recurring server tasks:

**Step 1** From the **Operate** menu, choose **Schedule Tasks** under the **Servers** submenu to open the List/Add Scheduled Tasks page.

**Step 2** Click the **Add Scheduled Task** icon in the Scheduled Tasks pane on the left to open the Add Scheduled Task page.

**Step 3** Enter values in the appropriate fields:

- Name of the scheduled task. This can be any identifying text string.
- Pull down from the available list of task types, which are:

- **dhcp-reload**—Reloads the DHCP server
- **dns-reload**—Reloads the DNS server
- **cdns-reload**—Reloads the Caching DNS server
- **sync-dhcp-pair**—Synchronizes the DHCP failover server pair
- **sync-dns-pair**—Synchronizes the HA DNS failover server pair
- **sync-zd-map**—Synchronizes zone distribution maps
- **sync-dns-update-map**—Synchronizes DNS update maps

- Enter the time interval for the scheduled task, such as 15m or 4w2d in the Schedule Interval field.

**Step 4** Click **Add Scheduled Task**.

**Step 5** If you click the name of the task on the List/Add Scheduled Tasks page, on the Edit Scheduled Task page you can view (in the Task Status section) the last status or the list of last errors (if any) that occurred during the task execution. Click **Run Now** to run the task immediately.

**Note** The DNS server startup and background loading slows down when HA is enabled before the HA DNS server communicates to its partner. You need to allow the HA DNS server to communicate with its partner before reloading or restarting the DNS server.

## CLI Commands

The **task** command configures scheduled task objects. These objects can perform periodic operations automatically.

To create a scheduled task, use **task name create task-type interval [sync-obj] [attribute=value]**. *task-type* controls the type of task to be scheduled. Available list of task types are: dhcp-reload, dns-reload, cdns-reload, sync-dhcp-pair, sync-dns-pair, sync-zd-map, and sync-dns-update-map.

To delete a scheduled task, use **task name delete**.

To edit a scheduled task, use **task name set attribute=value [attribute=value ...]**.

# Logs

## Log Files

The following table describes the Cisco Prime Network Registrar log files in the *install-path/logs* directory.

**Table 1: Log Files in .../logs Directory**

Component	File in /logs Directory	Local/Regional	Logs
Installation	install_cnr_log	Both	Installation process
Upgrade	ccm_upgrade_status_log	Both	Upgrade process
	dns_upgrade_status_log	Local	Upgrade process
	dhcp_upgrade_status_log	Local	Upgrade process
Server agent	agent_server_1_log	Both	Server agent starts and stops
Port check	checkports_log	Both	Network ports
DNS server	name_dns_1_log	Local	DNS activity
	dns_startup_log	Local	DNS startup activity
	dns_packet_log	Local	DNS packet logging messages <sup>1</sup>
CDNS server	cdns_log	Local	CDNS activity
	cdns_startup_log	Local	CDNS startup activity
	cdns_query_log	Local	CDNS query log entries <sup>2</sup>
DHCP server	name_dhcp_1_log	Local	DHCP activity
	dhcp_startup_log	Local	DHCP startup activity
TFTP server	file_tftp_1_log file_tftp_1_trace	Local	TFTP activity
	tftp_startup_log	Local	TFTP startup activity
SNMP server	cnrsnmp_log	Both	SNMP activity
CCM database	config_ccm_1_log	Both	CCM configuration, starts, stops
	ccm_startup_log	Both	CCM startup activity
Web UI	cnrwebui_log	Both	Web UI state

Component	File in /logs Directory	Local/Regional	Logs
Tomcat/web UI (in cnrwebui subdirectory)	catalina.date.log.txt jsui_log.date.txt cnrwebui_access_log.date.txt	Both	CCM database for Tomcat server and web UI (Because new files are created daily, periodically archive old log files.)
Resource Limits	ccm_monitor_log	Both	Resource limit activity.

<sup>1</sup> When packet-logging is enabled and "packet" is set as the packet-logging-file, the packet logging messages are logged to the dns\_packet\_log file. Restart the server to see this log file.

<sup>2</sup> When the query log-setting is enabled, the query log entries are logged to the cdns\_query\_log file.

DNS, DHCP, CDNS, CCM, and TFTP servers can generate a number of log files, each with a preconfigured maximum size of 10 MB. This preconfigured value applies to new installs only.



**Note** Upgrades from pre-10.1 versions will use the old preconfigured (or explicitly configured) value of 1,000,000 bytes for log files.

The first log file name has the \_log suffix. When this file reaches its maximum size, it gets the .01 version extension appended to its name and a new log file is created without the version extension. Each version extension is incremented by one for each new file created. When the files reach their configured maximum number, the oldest file is deleted and the next oldest assumes its name. The usual maximum number is 10 for the DNS, DHCP, CDNS, CCM, and TFTP servers.

Cisco Prime Network Registrar also has server\_startup\_log files. This applies to the CCM, DHCP, DNS, and TFTP servers. These files log the start up and shut down phases of the server (the information is similar to the normal log file information). Server startup log files are useful in diagnosing problems that have been reported when the server was last started.

The number of these start-up logs is fixed at four for a server, and the size is fixed at 10 MB per server.



**Note** Some user commands can create *User authentication* entries in the Server Agent log because of separate connections to the cluster. Do not interpret these as a system security violation by another user.

Logging can also be directed to Syslog. See [Modifying the cnr.conf File, on page 30](#).

## CLI Commands

You can check the configured maximums for the DNS, DHCP, and TFTP servers using **[server] type serverLogs show** in the CLI, which shows the maximum number (*nlogs*) and size (*logsize*) of these protocol server log files. You can adjust these parameters using **[server] type serverLogs set nlogs=*nlogs* logsize=*logsize***. You cannot adjust these maximums for any of the other log files.



**Note** A change to the server logs will not take effect until you restart Cisco Prime Network Registrar.

## Logging Server Events

When you start Cisco Prime Network Registrar, it automatically starts logging Cisco Prime Network Registrar system activity. Cisco Prime Network Registrar maintains all the logs by default on:

- **Windows**—*install-path*\logs
- **Linux**—*install-path*/logs (to view these logs, use the **tail -f** command)



---

**Tip** To avoid filling up the Windows Event Viewer and preventing Cisco Prime Network Registrar from running, in the Event Log Settings, check the **Overwrite Events as Needed** box. If the events do fill up, save them to a file, then clear them from the Event Log.

---

## Local Basic or Advanced and Regional Web UI

Server logging is available in the web UI when you open the Manage Servers page for a server (see the [Managing Servers, on page 1](#)), then click the **Logs** tab. This opens the logs for server page. The log is in chronological order with the page with the latest entries shown first. If you need to see earlier entries, click the left arrow at the top or bottom of the page.

## Related Topics

[Searching the Logs, on page 8](#)

[Logging Format and Settings, on page 7](#)

## Logging Format and Settings

The server log entries include the following categories:

- **Activity**—Logs the activity of your servers.
- **Info**—Logs standard operations of the servers, such as starting up and shutting down.
- **Warning**—Logs warnings, such as invalid packets, user miscommunication, or an error in a script while processing a request.
- **Error**—Logs events that prevent the server from operating properly, such as out of memory, unable to acquire resources, or errors in configuration.



---

**Note** Warnings and errors go to the Event Viewer on Windows (see the Tip in [Logging Server Events, on page 7](#)). For a description of the log messages for each server module, see the *install-path/docs/msgid/MessageIdIndex.html* file.

---

## Local Basic or Advanced and Regional Web UI

You can affect which events to log. For example, to set the logging for the local cluster DNS and DHCP server:

- **DNS**—From the **Deploy** menu, choose **DNS Server** under the **DNS** submenu to open the Manage DNS Server page. Click the name of the server to open the Edit DNS Server page. Expand the Log Settings section to view the log settings. Make changes to the attributes as desired, click **Save**, and then reload

the server. (See the table in the *"Troubleshooting DNS Servers"* section of *Cisco Prime Network Registrar 10.1 Authoritative and Caching DNS User Guide* for the log settings to maximize DNS server performance.)

- **DHCP**—From the **Deploy** menu, choose **DHCP Server** under the **DHCP** submenu to open the Manage DHCP Server page. Click the name of the server to open the Edit DHCP Server page. Expand the Logging section to view the log settings. Make changes to the attributes as desired, click **Save**, and then reload the server. (See the table in the *"Tuning the DHCP Server"* section of *Cisco Prime Network Registrar 10.1 DHCP User Guide* for the log settings to maximize DHCP server performance.)
- **CCM**—From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page. Click the name of the server to open the Edit Local CCM Server page. Expand the Logging section to view the log settings. Make changes to the attributes as desired and click **Save**. (See the table in [Managing CCM Server](#) to enable or disable the required log categories.)

## CLI Commands

Use `dns set log-settings=value`, `dhcp set log-settings=value`, `ccm set log-settings=value`, and `tftp set log-settings=value` for the respective servers.

## Searching the Logs

The web UI provides a convenient way to search for entries in the activity and startup log files. You can locate specific message text, log message IDs, and message timestamps using a regular expression string entry. In the text field next to the Search icon at the top or bottom of the page, enter the search string in the regular expression syntax. (For example, entering `name?` searches for occurrences of the string `name` in the log file.) Click the **Search** icon to view the results of log search. Change between table and text view by clicking the Page icon which is available at the top and bottom of the page.

To view the full message text, click the name of the log message. Click **Close** on the Log Search Result page to close the browser window.

## View Change Log

In the web UI, you can view the change logs and tasks associated with configurations you make.

## Local and Regional Web UI

From the **Operate** menu, choose **Change Log**. To view the change log, you must be assigned the database subrole of the `ccm-admin` or `regional-admin` role:

- The View Change Log page shows all the change logs, sorted by DBSN name. To get to the bottom of the list, click the right arrow at the bottom left of the page. Click the DBSN number of the change log entry to open a View Change Set page for it.

On the View Change Log page, you can filter the list, manually trim it, and save it to a file. You can filter the list by:

- Start and end dates
- Administrator who initiated the changes
- Configuration object class



- Specific object
- Object identifier (ID), in the format OID-00:00:00:00:00:00:00:00
- Server
- Database

Click **Filter List** or **Clear Filter** (to clear the filter that persists through the session). You can initiate a trim of the change log by setting how many days old you want the record to get before trimming it, by setting a number of days value in the “older than” field and clicking the **Delete** icon.

To save the change log entries to a comma-separated values (CSV) file, click the **Save to CSV Format** icon.

If a task is associated with a change log, it appears on the View Change Set page. You can click the task name to open the View CCM Task page for it.

## CLI Command

Use **export changeLog filename [attribute=value ...] [-all]** to export the change log records (in CSV format).

## Dynamic Update on Server Log Settings

The DHCP and the DNS servers register the changes on the server logs only during the server configuration, which happens during a reload. Reloading the servers is time consuming. Cisco Prime Network Registrar allows the DHCP and DNS servers to register the changes to log settings, without a reload.

### Local Basic or Advanced Web UI

To dynamically update DHCP server log settings, do the following:

- 
- Step 1** From the **Deploy** menu, choose **DHCP Server** under the **DHCP** submenu. The Manage DHCP Server page appears.
  - Step 2** Click the name of the DHCP server in the left pane to open the Edit DHCP Server page.
  - Step 3** Modify the log settings as desired.
  - Step 4** Click **Save** at the bottom of the page. The new log settings are applied to the DHCP server. The Manage DHCP Server page is displayed with an updated page refresh time.
- 

### Local Basic or Advanced Web UI

To dynamically update DNS server log settings, do the following:

- 
- Step 1** From the **Deploy** menu, choose **DNS Server** under the **DNS** submenu. This opens the Manage DNS Server page.
  - Step 2** Click the name of the DNS server in the left pane to open the Edit DNS Server page.
  - Step 3** Modify the log settings as desired.
  - Step 4** Click **Save** at the bottom of the page. The new log settings are applied to the DNS server. The Manage DNS Server page is displayed with an updated page refresh time.
-

**Note** If the `dhcp-edit-mode` or `dns-edit-mode` is set to `synchronous`, and if the server is running, the change in server log settings is communicated to the server.

## CLI Commands

To dynamically update the DHCP or DNS server log settings using the CLI, you must have the appropriate edit-mode set to `synchronous`. After changing the server log settings, use the `save` command to save the settings.

For example:

```
nrcmd> session set dhcp-edit-mode=synchronous
nrcmd> dhcp set log-settings=new-settings
nrcmd> save
```

## Running Data Consistency Rules

Using consistency rules, you can check data inconsistencies such as overlapping address ranges and subnets. You can set data consistency rules at the regional and local clusters.

The table on the List Consistency Rules page explains these rules. Check the check box next to the rule that you want to run.



**Note** You must set the locale parameters on UNIX to `en_US.UTF-8` when running Java tools that use Java SDK, such as `cnr_rules`.

The List Consistency Rules page includes functions to select all rules and clear selections. You can show the details for each of the rule violations as well as view the output. The rule selections you make are persistent during your user session.

## Local and Regional Web UI

To run consistency rules, do the following:

- 
- Step 1** From the **Operate** menu, choose **Consistency Reports** under the **Reports** submenu. The List Consistency Rules page appears.
- Step 2** Check the check boxes for each of the listed consistency rules that you want to apply.
- To select all the rules, click the **Select All Rules** link.
  - To clear all selections, click the **Clear Selection** link.
- Step 3** Click **Run Rules**.
- The Consistency Rules Violations page appears. The rules are categorized by violation type.
- To show details for the violations, click the **Show Details** link.
  - To show the output, click the page icon.

- Click **Display XML** to show the output in XML format.

**Step 4** Click **Return to Consistency Rules** to return to the List Consistency Rules page.

## CLI Tool

Use the **cnr\_rules** consistency rules tool from the command line to check for database inconsistencies. You can also use this tool to capture the results of the rule in a text or XML file.

The **cnr\_rules** tool is located at:

- **Windows**—...\\bin\cnr\_rules.bat
- **Linux**—.../usrbin/cnr\_rules

To run the **cnr\_rules** tool, enter:

```
> cnr_rules -N username -P password [options]
```

- **-N username** —Authenticates using the specified username.
- **-P password** —Authenticates using the specified password.
- **options** —Describes the qualifying options for the tool, as described in the following table. If you do not enter any options, the command usage appears.

**Table 2: cnr\_rules Options**

Option	Description
<b>-list</b>	<p>Lists the available consistency rules.</p> <p><b>Note</b> The list of available commands is tailored to the permissions of the administrator specified in the value of the <b>-N</b> option.</p> <pre>&gt; cnr_rules -N admin -P changeme -list</pre>
<b>-run [rule-match]</b>	<p>Run the available rules. Optionally, you can run a subset of the available rules by applying a case-insensitive rule-match string.</p> <ul style="list-style-type: none"> <li>• Runs all rules: <pre>&gt; cnr_rules -N admin -P changeme -run</pre> </li> <li>• Runs only the rules whose names contain the string “dhcp”: <pre>&gt; cnr_rules -N admin -P changeme -run dhcp</pre> </li> </ul> <p><b>Tip</b> To match a string containing spaces, enclose the string using double-quotation marks ("). For example: <code>&gt; cnr_rules -N admin -P changeme -run "router interface"</code></p>

Option	Description
<b>-details</b>	<p>Includes details of the database objects that violate consistency rules in the results.</p> <p>Runs the DNS rules, and includes details of the database object in the results:</p> <pre>&gt; cnr_rules -N admin -P changeme -run DNS -details</pre>
<b>-xml</b>	<p>Generates rule results in an XML file.</p> <p><b>Note</b> When using the <b>-xml</b> option, the <b>-details</b> option is ignored because the XML file includes all the detailed information.</p> <pre>&gt; cnr_rules -N admin -P changeme -run -xml</pre>
<b>-path</b> <i>classpath</i>	<p>Changes the Java classpath that is searched to locate the available consistency rules (optional).</p> <p>In order to run a new, custom consistency rule, you can use this option. You must get the support of a support engineer to do this.</p>
<b>-interactive</b>	<p>Runs the tool in an interactive session.</p> <pre>&gt; cnr_rules -N admin -P changeme -run -interactive RuleEngine [type ? for help] &gt; ? Commands:   load &lt;class&gt;      // load the specified rule                     class   run &lt;rule-match&gt; // run rules matching a string,              or '*' for all   list              // list rules by name   xml               // toggle xml mode   detail            // toggle detail mode (non-xml           only)   quit              // quit RuleEngine</pre>
<b>-both</b>	Displays domain names in both Unicode and ASCII.

You can redirect the output of any of these preceding commands to another file. Use the following syntax to capture the rule results in a:

- Text file:

```
> cnr_rules -N username -P password -run -details > filename.txt
```

- XML file:

```
> cnr_rules -N username -P password -run -xml > filename.xml
```

# Monitoring and Reporting Server Status

Monitoring the status of a server involves checking its:

- State
- Health
- Statistics
- Log messages
- Address usage
- Related servers (DNS and DHCP)
- Leases (DHCP)

## Related Topics

[Server States, on page 13](#)

[Displaying Health, on page 13](#)

[Displaying Statistics, on page 15](#)

[Displaying IP Address Usage, on page 26](#)

[Displaying Related Servers, on page 27](#)

[Displaying Leases, on page 29](#)

## Server States

All Cisco Prime Network Registrar protocol servers (DNS, DHCP, SNMP, and TFTP) pass through a state machine consisting of the following states:

- **Loaded**—First step after the server agent starts the server (transitional).
- **Initialized**—Server was stopped or fails to configure.
- **Unconfigured**—Server is not operational because of a configuration failure (transitional).
- **Stopped**—Server was administratively stopped and is not running (transitional).
- **Running**—Server is running successfully.

The two essential states are initialized and running, because the server transitions through the states so quickly that the other states are essentially invisible. Normally, when the server agent starts the server, it tells the server to be up. The server process starts, sets its state to loaded, then moves up to running. If you stop the server, it walks down the states to initialized, and if you restart, it moves up to running again. If it fails to configure for some reason, it drops back to initialized, as if you had stopped it.

There is also an exiting state that the server is in very briefly when the process is exiting. The user interface can also consider the server to be disabled, but this rarely occurs and only when there is no server process at all (the server agent was told not to start one).

## Displaying Health

You can display aspects of the health of a server, or how well it is running. The following items can decrement the server health, so you should monitor their status periodically. For the:

- Server agent (local and regional clusters)
- CCM server (local and regional clusters)
- DNS server (local cluster):
  - Configuration errors
  - Memory
  - Disk space usage
  - Inability to contact its root servers
- Caching DNS server (local cluster)
- DHCP server (local cluster):
  - Configuration errors
  - Memory
  - Disk space usage
  - Packet caching low
  - Options not fitting in the stated packet limit
  - No more leases available
- TFTP server (local cluster):
  - Memory
  - Socket read or write error
  - Exceeding the overload threshold and dropping request packets

## Server Health Status

The server health status varies from the value 0 to 10. The value 0 means the server is not running and 10 means the server is running. Some of the servers report only 0 or 10, and not anything in between. When a server reports a value from 1 to 9, it means that it detected conditions that indicate possible problems. It has nothing to do with the actual performance of the server. So, if the health of the server is a value from 1 to 9, the server log files need to be reviewed to see what errors were logged.




---

**Note** Depending on the level of activity and the size and number of log files, the condition that reduced the server health might not be visible in the log files. It is important to review the log files, but the servers do not log all the conditions that reduce the server health.

---

The following conditions can reduce the DHCP server health:

- Configuration errors (occurs when the server is getting started or restarting)
- When the server detects out-of-memory conditions
- When packet receive failures occur
- When packets are dropped because the server is out of request or response buffers
- When the server is unable to construct a response packet

Similar conditions exist for the TFTP server.



**Tip** Health values range from 0 (the server is not running) to 10 (the highest level of health). It is recommended that the exact value (1 to 10) of the health status be ignored, with the understanding that zero means server is not running and greater than zero means server is running. On Linux, you can run the **cnr\_status** command, in the *install-path/usrbin/* directory, to see if your local cluster server is running. For more information on how to check whether the local cluster server is running, see *Cisco Prime Network Registrar 10.1 Installation Guide*.

### Local Basic or Advanced and Regional Web UI

From the **Operate** menu, select **Manage Servers**. Check the Manage Servers page for the state and health of each server.

### CLI Commands

Use `[server] type getHealth`. The number 10 indicates the highest level of health, 0 that the server is not running.

## Displaying Statistics

To display server statistics, the server must be running.

### Local Basic or Advanced and Regional Web UI

Go to the Manage Servers page, click the name of the server in the left pane, then click the **Statistics** tab, if available. On the Server Statistics page, click the name of the attribute to get popup help.

The DHCP, DNS, and CDNS statistics are each divided into two groups of statistics. The first group is for total statistics and the second group is for sample statistics. The total statistics are accumulated over time. The sample statistics occur during a configurable sample interval. The names of the two categories vary per server and per user interface, and are identified in the following table.

*Table 3: Server Statistics Categories*

Server	User Interface	Total Statistics (Command)	Sample Statistics (Command)
DHCP	Web UI	Total Statistics	Activity Summary
	CLI	Total Counters since the start of the last DHCP server process <b>(dhcp getStats)</b>	Sample counters that were collected during the last sample interval. These are updated once every sample period. <b>(dhcp getStats server sample)</b>

Server	User Interface	Total Statistics (Command)	Sample Statistics (Command)
DNS	Web UI	Total Statistics	Sample Statistics
	CLI	Total Counters since the start of the last server process <b>(dns getStats)</b>	Sample counters that are being collected during the current sample interval. These are updated constantly. <b>(dns getStats performance sample)</b>
CDNS	Web UI	Total Statistics	Sample Statistics
	CLI	Total Counters since the start of the last server process <b>(cdns getStats server total)</b>	Sampled counters since the last sample interval <b>(cdns getStats server sample)</b>

To set up the sample counters, you must activate either the *collect-sample-counters* attribute for the server or a *log-settings* attribute value called *activity-summary*. You can also set a *log-settings* value for the sample interval for each server, which is preset to 5 minutes. The *collect-sample-counters* attribute is preset to true for the DNS server, but is preset to false for the DHCP server. For example, to enable the sample counters and set the interval for DHCP, set the following attributes for the DHCP server:

- Enable *collect-sample-counters* (**dhcp enable collect-sample-counters**)
- Set *log-settings* for *activity-summary* (**dhcp set log-settings=activity-summary**)
- Set *activity-summary-interval* to 5m (**dhcp set activity-summary-interval=5m**)

## CLI Commands

In the CLI, if you use **[server] type getStats**, the statistics are encoded in curly braces followed by sets of fields, as described in [Table 4: DNS Statistics](#) for DNS, [Table 6: DHCP Statistics](#) for DHCP, and [Table 7: TFTP Statistics](#) for TFTP. The **server type getStats all** command is more verbose and identifies each statistic on a line by itself. Using the additional **sample** keyword shows the sample statistics only.

Reset the counters and total statistic by using **dhcp resetStats**, **dns resetStats**, or **cdns resetStats**.

## DNS Statistics

The DNS server statistics in the web UI appear on the DNS Server Statistics page, click on the statistic's name to read its description. You can refresh the DNS Server Statistics.

The DNS server statistics details include the server identifier, recursive service, process uptime, time since reset, server status, counter reset time, sample time, statistics interval, total zone, and total RRs, followed by total and sample statistics mentioned below:

- Performance Statistics—Displays the statistics of the DNS Server performance.
- Query Statistics—Displays the statistics of the queries.



- Update Statistics—Displays the statistics of the DNS updates.
- HA Statistics—Displays the statistics of the HA DNS Server.
- Push Notification Statistics—Displays the statistics of DNS Push Notifications.
- Host Health Check Statistics—Displays the statistics of DNS Host Health Check.
- DB Statistics—Displays the statistics of DNS Database.
- Cache Statistics—Displays the statistics of DNS Query Cache.
- Security Statistics—Displays the statistics of the security.
- IPv6 Statistics—Displays the statistics of the IPv6 packets received and sent.
- Error Statistics—Displays the statistics of the errors.
- Max Counter Statistics—Displays the statistics of the maximum number of concurrent threads, RRs, DNS update latency, concurrent packets, and so on.
- Top Name Statistics—Displays the statistics of the top names.



**Note** To get the most recent data, click the **Refresh Server Statistics** icon at the top left of the Statistics page.

The `dns getStats` command has the following options:

```
dns getStats [<performance [,] query [,] update [,] errors [,] security [,]
maxcounters [,] ha [,] ipv6 [,] dns-pn [,] cache [,] datastore [,] top-names [,]
dns-hhc | all> [total | sample]]
```

The `dns getStats all` command is the most commonly used. The `dns getStats` command without `all` option returns the statistics in a single line of positional values in the following format (the table below shows how to read these values):

```
nrcmd> dns getStats

100 Ok
{1} 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17
```

**Table 4: DNS Statistics**

Field	Statistic	Description
{1}	id	Implementation ID (release and build information).
2	config-recurs	Recursion services—(1) available, (2) restricted, (3) unavailable.
3	config-up-time	Time (in seconds) elapsed since the last server startup.
4	config-reset-time	Time (in seconds) elapsed since the last server reset (restart).

Field	Statistic	Description
5	config-reset	Status or action to reinitializes any name server state—If using the (2) reset action, reinitializes any persistent name server state; the following are read-only statuses: (1) other—server in some unknown state, (3) initializing, or (4) running.
6	counter-auth-ans	Number of queries answered authoritatively.
7	counter-auth-no-names	Number of queries returning authoritative no such name responses.
8	counter-auth-no-data-resps	Number of queries returning authoritative no such data (empty answer) responses. (Deprecated statistics)
9	counter-non-auth-datas	Number of queries answered nonauthoritatively (cached). (Deprecated statistics)
10	counter-non-auth-no-datas	Number of queries answered nonauthoritatively with no data.
11	counter-referrals	Number of queries forwarded to other servers.
12	counter-errors	Number of responses answered with errors (RCODE values other than 0 or 3).
13	counter-rel-names	Number of requests received for names of only one label (relative names).
14	counter-req-refusals	Number of refused queries.
15	counter-req-unparses	Number of unparsable requests.
16	counter-other-errors	Number of aborted requests due to other errors.
17	total-zones	Total number of configured zones.

## CDNS Statistics

The CDNS server statistics in the web UI appear on the DNS Caching Server Statistics page, click on the name of the statistics to read its description. You can refresh the CDNS Server Statistics.

The CDNS server statistics details include the server identifier, recursive service, current time, process up, server restart time, counter reset time, sample time, statistics interval, time elapsed, and so on, followed by total and sample statistics mentioned below:

- Query Details—Displays the statistics of the queries.
- Answer Details—Displays the statistics related to CDNS query responses.
- Performance—Displays the statistics of the DNS Server performance.
- DNS64—Displays the statistics of DNS64.
- Firewall—Displays the statistics of DNS firewall.

- Rate Limiting—Displays the statistics related to rate limiting.
- Top Name Statistics—Displays the statistics of the top names.



**Note** To get the most recent data, click the **Refresh Server Statistics** icon at the top left of the Statistics page.

The **cdns getStats** command has the following options:

```
cdns getStats [<server | top-names | rate-limit | all> [total | sample]]
```

Both **cdns getStats** and **cdns getStats server** commands are same as **cdns getStats server total**.

The **cdns getStats top-names** and **cdns getStats rate-limit** commands always report the “sample” data, they ignore the mode parameter (there is no “total” data to report).

The **cdns getStats** and **cdns getStats all** commands return the statistics mentioned in the following table.

**Table 5: CDNS Statistics**

Statistic	Description
name	Name identifying the DNS Caching Server.
restart-time	Time when the DNS Caching Server was last restarted or reloaded.
config-recurs	Recursion services offered by this name server.
time-current	Current time given by the CDNS Server.
time-up	Amount of time the server has been up and running.
time-elapsed	Time elapsed since last statistics poll.
reset-time	The most recent time the stats were reset (that is, <b>cdns resetStats</b> in CLI).
sample-time	Time the server collected the last set of sample statistics.
sample-interval	Sample interval used by the server when collecting sample statistics.
queries-total	Total number of queries received by the CDNS Server.
queries-over-tcp	Total number of queries received over TCP by the CDNS Server.
queries-over-ipv6	Total number of queries received over TCP by the CDNS Server.
queries-with-edns	Number of queries with EDNS OPT RR present.
queries-with-edns-do	Number of queries with EDNS OPT RR with DO (DNSSEC OK) bit set.
queries-type-A	Number of A queries received.
queries-type-AAAA	Number of AAAA queries received.

Statistic	Description
queries-type-ANY	Number of ANY queries received.
queries-type-CNAME	Number of CNAME queries received.
queries-type-PTR	Number of PTR queries received.
queries-type-NS	Number of NS queries received.
queries-type-SOA	Number of SOA queries received.
queries-type-MX	Number of MX queries received.
queries-type-DS	Number of DS queries received.
queries-type-DNSKEY	Number of DNSKEY queries received.
queries-type-RRSIG	Number of RRSIG queries received.
queries-type-NSEC	Number of NSEC queries received.
queries-type-NSEC3	Number of NSEC3 queries received.
queries-type-other	Number of queries received of type 256+.
queries-with-flag-QR	Number of incoming queries with QR (query response) flag set. These queries are dropped.
queries-with-flag-AA	Number of incoming queries with AA (auth answer) flag set. These queries are dropped.
queries-with-flag-TC	Number of incoming queries with TC (truncation) flag set. These queries are dropped.
queries-with-flag-RD	Number of incoming queries with RD (recursion desired) flag set.
queries-with-flag-RA	Number of incoming queries with RA (recursion available) flag set.
queries-with-flag-Z	Number of incoming queries with Z flag set.
queries-with-flag-AD	Number of incoming queries with AD flag set.
queries-with-flag-CD	Number of incoming queries with CD flag set.
queries-failing-acl	Number of queries being dropped or refused due to ACL failures.
cache-hits	Total number of queries that were answered from cache.
cache-misses	Total number of queries that were not found in the cache.
cache-prefetches	Number of prefetches performed.
mem-query-cache-exceeded	Number of times the message cache has gone over the configured limit. This indicates that the configured limit may be undersized for its environment.

<b>Statistic</b>	<b>Description</b>
mem-cache-exceeded	Number of times the RRSet cache has gone over the configured limit. This indicates that the configured limit may be undersized for its environment.
remote-ns-cache-exceeded	Number of times the remote name server cache has gone over the configured limit. This indicates that the configured limit may be undersized for its environment.
key-cache-exceeded	Number of times the key cache has gone over the configured limit. This indicates that the configured limit may be undersized for its environment.
smart-cache	Total number of times the CDNS Server employed a smart-cache response, when smart-cache is enabled.
requestlist-total	Total number of queued requests waiting for recursive replies.
requestlist-total-user	Total number of queued user requests waiting for recursive replies.
requestlist-total-system	Total number of queued system requests waiting for recursive replies.
requestlist-total-average	Average number of requests on the request list.
requestlist-total-max	Maximum number of requests on the request list.
requestlist-total-overwritten	Number of requests on the request list that were overwritten by newer entries.
requestlist-total-exceeded	Number of requests dropped because the request list was full.
recursive-replies-total	Total number of recursive queries replies.
recursive-time-average	Average time to complete a recursive query.
recursive-time-median	The median time to complete a recursive query.
mem-process	An estimate of the memory in bytes of the CDNS process.
mem-cache	Memory in bytes of RRSet cache.
mem-query-cache	Memory in bytes of incoming query message cache.
mem-iterator	Memory in bytes used by the CDNS iterator module.
mem-validator	Memory in bytes used by the CDNS validator module.
answers-with-NOERROR	Number of answers from cache or recursion that result in rcode of NOERROR being returned to client.
answers-with- NXDOMAIN	Number of answers from cache or recursion that result in rcode of NXDOMAIN being returned to client.
answers-with-REFUSED	Number of answers from cache or recursion that result in rcode of REFUSED being returned to client.

Statistic	Description
answers-with-SERVFAIL	Number of answers from cache or recursion that result in rcode of SERVFAIL being returned to client.
answers-with-FORMERR	Number of answers from cache or recursion that result in rcode of FORMERR being returned to client.
answers-with-NOTAUTH	Number of answers from cache or recursion that result in rcode of NOTAUTH being returned to client.
answers-with-NOTIMP	Number of answers from cache or recursion that result in rcode of NOTIMP being returned to client.
answers-with-NODATA	Number of answers that result in pseudo rcode of NODATA being returned to client.
answers-with-other-errors	Number of answers that result in pseudo rcode of NODATA being returned to client.
answers-secure	Number of answers that correctly validated.
answers-unsecure	Number of answers that did not correctly validate.
answers-rrset-unsecure	Number of RRsets marked as bogus by the validator.
answers-unwanted	Number of replies that were unwanted or unsolicited. High values could indicate spoofing threat.
dns64-a2aaaa-conversions	Number of times DNS64 has converted a type A RR to a type AAAA RR.
dns64-ptr-conversions	Number of times DNS64 has converted an IPv4 PTR RR to an IPv6 PTR RR.
firewall-dropped	Number of times DNS Firewall dropped a query.
firewall-redirected	Number of times DNS Firewall redirected a query.
firewall-refused	Number of times DNS Firewall refused a query.
firewall-redirect-nxdomain	Number of times DNS Firewall redirected a query with an NXDOMAIN answer.
firewall-rpz	Number of times DNS Firewall RPZ rules matched an incoming query.
exceeded-max-target-count	Number of queries that exceeded the maximum number of name servers glue lookups allowed.
client-rate-limit	Number of times a client has been rate limited, when <i>client-rate-limiting</i> is enabled.
domain-rate-limit	Number of times a zone has been rate limited, when <i>domain-rate-limiting</i> is enabled.

## DHCP Statistics

The DHCP server statistics in the web UI appear on the DHCP Server Statistics page, click on the statistic's name to read its description.

The DHCP server statistics details include the server start time, server reload time, server up time, and statistics reset time followed by statistics in the following sections:

- Total Statistics—Displays the total statistics of the scopes, request buffers, response buffers, packets and so on.
- Lease Counts (IPv4)—Displays the statistics of the IPv4 lease counts such as active leases, configured leases, reserved leases, and reserved active leases.
- Packets Received (IPv4)—Displays the statistics of the IPv4 packets received.
- Packets Sent (IPv4)—Displays the statistics of the IPv4 packets sent.
- Packets Failed (IPv4)—Displays the statistics of the failed IPv4 packets.
- Failover Statistics—Displays the statistics of the DHCP failover server.
- IPv6 Statistics—Displays the statistics of the IPv6 prefixes configured, timed-out IPv6 offer packets and so on.
- Lease Counts (IPv6)—Displays the statistics of the IPv6 lease counts of active leases, configured leases, reserved leases, and reserved active leases.
- Packets Received (IPv6)—Displays the statistics of the IPv6 packets received.
- Packets Sent (IPv6)—Displays the statistics of the IPv6 packets sent.
- Packets Failed (IPv6)—Displays the statistics of the failed IPv6 packets.

Additional Attributes include Top Utilized Aggregations and Activity Summary.



**Note** To get the most recent data, click the **Refresh Server Statistics** icon at the top left of the Statistics page.

The **dhcp getStats** command has the following options:

```
dhcp getStats [<all | server [, failover [, dhcpv6 [, top-utilized>
[total | sample]]]
```

The **dhcp getStats all** command is the most commonly used. The **dhcp getStats** command without **all** option returns the statistics in a single line of positional values in the following format (the table below shows how to read these values):

```
nrcmd> dhcp getStats
100 Ok
{1} 2 3 4 5 6 7 8
```

**Table 6: DHCP Statistics**

Field	Statistic	Description
{1}	start-time-str	Date and time of last server reload, as a text string.

Field	Statistic	Description
2	total-discovers	Number of DISCOVER packets received.
3	total-requests	Number of REQUEST packets received.
4	total-releases	Number of RELEASED packets received.
5	total-offers	Number of OFFER packets sent.
6	total-acks	Number of acknowledgement (ACK) packets sent.
7	total-naks	Number of negative acknowledgement (NAK) packets sent.
8	total-declines	Number of DECLINE packets received.

## TFTP Statistics

The TFTP server statistics in the web UI appear on the TFTP Server Statistics page, click on the statistic's name to read its description. The following table shows the TFTP statistics encoded as output to the generic `tftp getStats` command.

When the TFTP server starts up, it allocates sessions (`tftp-max-sessions`) and packets (`tftp-max-packets`) for its use. The TFTP session represents the communication between the TFTP client and TFTP server.

When a read request reaches the TFTP server, the server assigns a packet for the request, increments the `total-packets-in-use` and `total-read-requests` values by one, and responds to the user with a data packet. The TFTP server backs up the latest communication packet to resend, if needed. The TFTP server picks another packet from the pool to use it as data packet. When the TFTP server receives an acknowledgment for the block of data sent to the client, it sends the next data block. The TFTP server queues up packets associated with a session, if the session is not able to work on the packets immediately.

The TFTP server statistics details are available for:

- **Attribute**—Displays the server statistics such as port number, default device, home directory, use home directory as root, and so on.
- **Log Settings**—Displays the statistics of the log level, log settings, and packet trace level.



**Note** To get the most recent data, click the **Refresh Server Statistics** icon at the top left of the page.

TFTP statistics is encoded as an output to the generic `tftp getStats` command in the following format:

```
nrcmd> tftp getStats

100 Ok
{1} 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
```

**Table 7: TFTP Statistics**

Field	Attribute	Description
{1}	id	Implementation ID (release and build information).



Field	Attribute	Description
2	server-state	State of the server (up or down).
3	server-time-since-start	Running time since last start.
4	server-time-since-reset	Running time since last reset.
5	total-packets-in-pool	Number of packets in the pool.
6	total-packets-in-use	Number of packets the server is using.
7	total-packets-received	Number of packets received since the last start or reload.
8	total-packets-sent	Number of packets sent since the last start or reload.
9	total-packets-drained	Number of packets read and discarded since the last start or reload.
10	total-packets-dropped	Number of packets dropped since the last start or reload.
11	total-packets-malformed	Number of packets received that were malformed since the last start or reload.
12	total-read-requests	Number of packets read since the last start or reload.
13	total-read-requests-completed	Number of read packets completed since the last start or reload.
14	total-read-requests-refused	Number of read packets refused since the last start or reload.
15	total-read-requests-ignored	Number of read packets ignored since the last start or reload.
16	total-read-requests-timed-out	Number of read packets that timed out since the last start or reload.
17	total-write-requests	Number of read packets that were write requests since the last start or reload.
18	total-write-requests-completed	Number of write requests completed since the last start or reload.
19	total-write-requests-refused	Number of write requests refused since the last start or reload.
20	total-write-requests-ignored	Number of write requests ignored since the last start or reload.
21	total-write-requests-timed-out	Number of write requests that timed out since the last start or reload.

Field	Attribute	Description
22	total-docsis-requests	Number of DOCSIS requests received since the last start or reload.
23	total-docsis-requests-completed	Number of DOCSIS requests completed since the last start or reload.
24	total-docsis-requests-refused	Number of DOCSIS requests refused since the last start or reload.
25	total-docsis-requests-ignored	Number of DOCSIS requests ignored since the last start or reload.
26	total-docsis-requests-timed-out	Number of DOCSIS requests that timed out since the last start or reload.
27	read-requests-per-second	Number of read requests per second.
28	write-requests-per-second	Number of write requests per second.
29	docsis-requests-per-second	Number of DOCSIS requests per second.

## Displaying IP Address Usage

Displaying IP address usage gives an overview of how clients are currently assigned addresses.

### Local Advanced and Regional Web UI

You can look at the local or regional cluster address space, or generate a DHCP utilization or lease history report at the regional cluster, to determine IP address usage. These functions are available in the **Design > DHCPv4** menu, if you have address space privileges at the local or regional cluster.

You can determine the current address space utilization by clicking the Current Usage tab for the unified address space, address block, and subnet (see the *"Viewing Address Utilization for Address Blocks, Subnets, and Scopes"* section in *Cisco Prime Network Registrar 10.1 DHCP User Guide*). You can also get the most current IP address utilization by querying the lease history (see the *"Querying Leases"* section in *Cisco Prime Network Registrar 10.1 DHCP User Guide*). In the latter case, the regional CCM server references the appropriate DHCP server directly. To ensure this subnet-to-server mapping, you must update the regional address space view so that it is consistent with the relevant local cluster. Do this by pulling the replica address space, or reclaiming the subnet to push to the DHCP server (see the *"Reclaiming Subnets"* section in *Cisco Prime Network Registrar 10.1 DHCP User Guide*). Also ensure that the particular DHCP server is running.

### CLI Commands

You can generate an IP address usage report using the **report** command. The command has the following syntax:

```
report [column-separator=string]
      [dhcp-only]
      [dhcpv4]
      [dhcpv6]
      [file=outputfile]
      [vpn=name]
```

The column-separator specifies the character string that separates the report columns (the preset value is the space character). If you want to include more than one space, precede them with the backslash (\) escape character (enclosed in quotation marks). You can specify DHCPv4 or DHCPv6 addresses (**dhcp-only** is the same as **dhcpv4**). Not specifying the VPN returns the addresses in the current VPN only.

## Displaying Related Servers

Cisco Prime Network Registrar displays the relationship among servers in a DNS zone distribution or a DHCP failover configuration. In the web UI, you can view a related servers page when you click the **Related Servers** icon on various pages. You can use the display of related servers to diagnose and monitor misconfigured or unreachable servers.

### Related Topics

[Monitoring Remote Servers Using Persistent Events, on page 27](#)

[DNS Zone Distribution Servers, on page 28](#)

[DHCP Failover Servers, on page 29](#)

## Monitoring Remote Servers Using Persistent Events

To service clients that require updates to DNS and LDAP related servers, the DHCP server uses a persistent event algorithm to ensure updates to related servers when a related server is temporarily unavailable. In addition, the algorithm prevents a misconfigured or offline DNS server from using up all the available update resources.

At startup, the DHCP server calculates the number of related servers in the configuration that require persistent events. A preconfigured *Maximum Pending Events* attribute (an Expert mode attribute that specifies the number of in-memory events that is preset to 40,000) is divided by the number of servers to obtain a limit on the number of events permitted for each remote server. This calculation covers related DNS and LDAP servers (DHCP failover does not use persistent storage for events). The DHCP server uses this calculation to issue log messages and take the action described in the following table. The table shows a hypothetical case of a DHCP server with four related DNS servers each having a limit of 10K events.

**Table 8: Persistent Event Algorithm**

Event Reached	
50% of the calculated per-server limit (Maximum Pending Events value divided by the number of total related servers); for example, 5K events on a related server out of a total of 40K maximum pending events	<p>Issues an INFO log message every 2 minutes, as long as the limits are exceeded:</p> <pre>The queue of events for the <i>name</i> remote server at <i>address</i> has <i>x</i> events, and has reached the info limit of <i>y/2</i> events out of an upper limit of <i>y</i> events per remote server. The remote server may be misconfigured, inoperative, or unreachable.</pre>

Event Reached	
100% of the calculated per-server limit and less than 50% of the Maximum Pending Events value; for example, 10K events on a related server, with fewer than 10K total maximum pending events	<p>Issues a WARNING log message every 2 minutes, as long as the limits are exceeded:</p> <p>The queue of events for the <i>name</i> remote server at address has <i>x</i> events, has exceeded the limit of <i>y</i> events per remote server, but is below the limit of <i>z</i> total events in memory. The remote server may be misconfigured, inoperative, or unreachable.</p>
100% of the calculated per-server limit and 50% or more of the Maximum Pending Events value; for example, 10K events on a related server, with 20K total maximum pending events	<p>Issues an ERROR log message every 2 minutes, as long as the limits are exceeded:</p> <p>The queue of events for the <i>name</i> remote server at address has <i>x</i> events, and has grown so large that the server cannot continue to queue new events to the remote server. The limit of <i>y</i> events per remote server and <i>z/2</i> total events in memory has been reached. This and future updates to this server will be dropped. The current eventID <i>n</i> is being dropped.</p> <p>The server drops the current triggering event and all subsequent events with that server.</p>
100% of the Maximum Pending Events value; for example, 40K events across all related servers	<p>Issues an ERROR log message:</p> <p>The queue of pending events has grown so large that the server cannot continue to queue new events. The queue's size is <i>z</i>, and the limit is <i>z</i>.</p> <p>The server drops all subsequent events with all related servers.</p>

SNMP traps and DHCP server log messages also provide notification that a related server is unreachable.

## DNS Zone Distribution Servers

A DNS zone distribution simplifies creating multiple zones that share the same secondary server attributes. You can view and set the primary and secondary DNS servers in a zone distribution.

### Local Basic or Advanced Web UI

From the **Deploy** menu, click **Zone Distribution** under the **DNS** submenu. This opens the List/Add Zone Distributions page. The local cluster allows only one zone distribution, the default. Click this zone distribution name to open the Edit Zone Distribution page, which shows the authoritative and secondary servers in the zone distribution.

### Regional Web UI

From the **Deploy** menu, choose **Zone Distribution** under the **DNS** submenu. This opens the List/Add Zone Distributions page. The regional cluster allows creating more than one zone distribution. Click the zone distribution name to open the Edit Zone Distribution page, which shows the name of the zone distribution map, primary, authoritative, and secondary servers in the zone distribution.



**Note** Default zone distribution names are not editable. However, non-default zone distribution names are editable and can be saved.

### CLI Commands

Create a zone distribution using **zone-dist name create primary-cluster [attribute=value]**, then view it using **zone-dist list**. For example:

```
nrcmd> zone-dist distr-1 create Boston-cluster
```

```
nrcmd> zone-dist list
```

## DHCP Failover Servers

Related servers in a DHCP failover pair relationship can show the following information:

- **Type**—Main or backup DHCP server.
- **Server name**—DNS name of the server.
- **IP address**—Server IP address in dotted octet format.
- **Requests**—Number of outstanding requests, or two dashes if not applicable.
- **Communication status**—OK or INTERRUPTED.
- **Cluster state**—Failover state of this DHCP server.
- **Partner state**—Failover state of its partner server.

For details on DHCP failover implementation, see the *"Managing DHCP Failover" section in Cisco Prime Network Registrar 10.1 DHCP User Guide*

### Local Basic or Advanced Web UI

From the **Deploy** menu, choose **Failover Pairs** under the **DHCP** submenu. The List/Add DHCP Failover Pairs page shows the main and backup servers in the failover relationship.

### CLI Commands

Use **dhcp getRelatedServers** to display the connection status between the main and partner DHCP servers. If there are no related servers, the output is simply 100 Ok.

## Displaying Leases

After you create a scope, you can monitor lease activity and view lease attributes.

### Local Basic or Advanced Web UI

From the **Design** menu, choose **Scopes** under the **DHCPv4** submenu; or from the **Design** menu, choose **Prefixes** under the **DHCPv6** submenu. On the List/Add DHCP Scopes or List/Add DHCPv6 Prefixes page, click the **Leases** tab to view the leases.

## Local Advanced and Regional Advanced Web UI

From the **Operate** menu, choose **DHCPv4 Lease History** or **DHCPv6 Lease History** under the **Reports** submenu. Set the query parameters and then query the lease history. (See the *"Querying Leases"* section in *Cisco Prime Network Registrar 10.1 DHCP User Guide*.)

## Troubleshooting DHCP and DNS Servers

The following sections describe troubleshooting the configuration and the DNS, DHCP, and TFTP servers.

### Related Topics

[Immediate Troubleshooting Actions, on page 30](#)

[Modifying the cnr.conf File, on page 30](#)

[Troubleshooting Server Failures, on page 33](#)

[Troubleshooting and Optimizing the TFTP Server, on page 37](#)

[Linux Troubleshooting Tools, on page 34](#)

[Using the TAC Tool, on page 34](#)

### Immediate Troubleshooting Actions

When facing a problem, it is crucial not to cause further harm while isolating and fixing the initial problem. Here are things to do (or avoid doing) in particular:

- Have 512 MB or more of memory and 2.5 GB or more of a data partition.
- Do not reboot a cable modem termination system (CMTS).
- Enable DHCP failover.
- Do not reload, restart, or disrupt Cisco Prime Network Registrar with failover resynchronization in progress.

### Modifying the cnr.conf File

Cisco Prime Network Registrar uses the **cnr.conf** file for basic configuration parameters. This file is normally located in the *install-path/conf* directory. Cisco Prime Network Registrar creates the file during installation and processes it line by line.

You can edit this file if configuration parameters change. Note that during normal operation, you would not want to change the values. However, certain conditions might require you to modify certain values, such as when you move the data files for disk space reasons.

The format of the **cnr.conf** file consists of parameter name-value pairs, one per line; for example, for a Windows local cluster installation:

```
cnr.rootdir=C:\\Program Files\\Network Registrar\\Local
cnr.ccm-port=1234
cnr.cisco-gss-appliance-integration=n
cnr.datadir=C:\\NetworkRegistrar\\Local\\data
cnr.java-home=C:\\Program Files\\Java\\jre1.5.0_12
cnr.logdir=C:\\NetworkRegistrar\\Local\\logs
```

```

cnr.https-port=8443
cnr.tmpdir=C:\\NetworkRegistrar\\Local\\temp
cnr.http-port=8080
cnr.ccm-mode=local
cnr.ccm-type=cnr
cnr.http-enabled=y
cnr.https-enabled=n
cnr.keystore-file=C:
cnr.keystore-password=unset
cnr.backup-time=23:45

```

Directory paths must be in the native syntax for the operating system. The format allows the use of colons (:) in directory paths, but not as name-value pair separators; it does not allow line continuation or embedded unicode characters. Other modifications to the file might include the location of the log directory (see [Log Files, on page 5](#)) or the time `cnr_shadow_backup` backups should occur (see [Setting Automatic Backup Time](#)).

In rare cases, you might want to modify the file; for example, to exclude certain data from daily backups due to capacity issues. To do this, you need to add the appropriate settings manually.



### Caution

We recommend that you use the default settings in this file. If you must change these settings, do so only in consultation with the Cisco Technical Assistance Center (TAC) or the Cisco Prime Network Registrar development team.

The following settings are supported:

- `cnr.backup-dest`—Specify the destination to place backed up databases. Defaults to `cnr.datadir` if not specified.
- `cnr.backup-dbs`—Provide a comma-separated list of the databases you want to backup. For a local cluster the default is `cdns,ccm,dhcp,dns,mcd,cnrsnmp`. For a regional cluster it is `ccm,dns,leasehist,lease6hist,subnetutil,replica`.
- `cnr.backup-files`—Provide a comma-separated list of files and the complete path to the files that you want copied as part of the backup. Files are copied to `cnr.backup-dest`.
- `cnr.dbrecover-backup`—Specify whether to run db recover and db verify on a backed up Oracle Berkeley database. The default is `true`. This setting is used for daily backups only. Manual backups ignore this setting. Disabling the automatic operation means that you must run the operation manually, preferably on a separate machine, or at a time when the Cisco Prime Network Registrar servers are relatively idle.
- `cnr.daily-backup`—Specify whether to run the daily back up. The default is `true`.

## Syslog Support

Cisco Prime Network Registrar supports logging to a Syslog server (on Linux). The Syslog support is not enabled by default. To configure which messages need to be logged, based on logging levels, the `cnr.conf` file must be updated.

In addition, on Windows, event logging for Warnings and Errors is enabled by default (for Windows Event log). In this release, you can log more (or less) to the event log by changing the log settings.

The following `cnr.conf` configuration parameters are supported:

- `cnr.syslog.enable`—Specifies whether logging to Syslog server or Windows Event log is enabled for Prime Network Registrar servers.
  - To disable all logging, the value can be `0`, `off`, or `disabled`.

- To enable all logging, the value can be 1, on, or enabled.
- By default, this parameter is disabled for Linux and enabled for Windows.
- `cnr.syslog.levels`—Specifies the severity levels to be logged to Syslog or Windows Event log. If Syslog is enabled, this defaults to warning and error. The value can be a case-blind, comma separated, list of the following keywords: error, warning, activity, info, and debug. This parameter is ignored if Syslog is disabled.




---

**Caution** While it is possible to enable all of the severity levels and thus all messages written to the server log files are also logged to Syslog, this is not recommended. The performance impact on Syslog and the servers may vary greatly depending on how logging is configured. Syslog may rate limit the messages, so useful messages may also be lost.

Cisco highly recommends reviewing the Syslog settings and messages in order to minimize the number of messages written. Writing too many messages to Syslog will cause a performance impact on the Cisco Prime Network Registrar servers and Syslog.

---

- `cnr.syslog.facility`—Specifies the facility under which Syslog logs (Linux OS). This parameter is ignored for Windows. The valid facility keywords are daemon (the default), local0, local1, local2, local3, local4, local5, local6, local7.
- `cnr.syslog.ids`—Specifies the individual messages to be logged (or not logged) as either a comma separated list of message IDs or message ID ranges (*x-y*). If a message ID or range is preceded by a minus sign (hyphen) or ! (exclamation mark), the message ID or range of IDs will explicitly not be logged. The explicitly referenced message IDs are logged or not logged regardless of any other Syslog settings (including the `.enable` setting).

Refer to the `/opt/nwreg2/local/docs/msgid/*.html` files (or the actual server log files) for determining the message IDs.

For example:

```
cnr.syslog.ids=4000-4100,-4101-4200,4300
```

This would cause messages 4000-4100 and 4300 to be logged to Syslog (or Windows Event logging) and messages 4101-4200 to NOT be logged (regardless of any other Syslog settings).




---

**Note**

- These parameters apply to all Cisco Prime Network Registrar servers (`cnrservagt`, `ccm`, `cdns`, `cnrsnmp`, `dns`, `dhcp`, and `tftp`).
  - To apply any change to the `cnr.conf` parameters, Cisco Prime Network Registrar must be restarted.
- 

The following `cnr.conf` configuration parameters allow server-specific overrides of the above parameters. server is one of `cnrservagt`, `ccm`, `cdns`, `cnrsnmp`, `dns`, `dhcp`, and `tftp`.

- `cnr.syslog.server.enable`—Specifies whether Syslog or Windows Event logging is enabled for the specified server (`cnr.syslog.enable` is ignored for that server).



- `cnr.syslog.server.levels`—Specifies the severity levels for the specified server (`cnr.syslog.levels` is ignored for that server).
- `cnr.syslog.server.facility`—Specifies the Syslog facility for the specified server (`cnr.syslog.facility` is ignored for that server).

The server specific configuration value is used, if specified. Otherwise, all parameters of the server are used. For example, to enable Syslog only for DHCP, add the following to the `cnr.conf` file:

```
cnr.syslog.dhcp.enable=1
```

As an example of setting Syslog setting for all servers:

```
cnr.syslog.enable=1
cnr.syslog.levels=error,warning,activity
```

To enable Syslog only for the Authoritative DNS server:

```
cnr.syslog.dns.enable=1
cnr.syslog.dns.levels=error,warning,activity
```




---

**Tip** Syntax or other errors in the `cnr.conf` parameters are not reported and are ignored (that is, if a `levels` keyword is mistyped, that keyword is ignored). Therefore, if a configuration change does not work, check if the parameter(s) have been specified correctly.

---




---

**Note** Most Syslog implementations implement rate limiting and the Cisco Prime Network Registrar server logging can easily trigger this causing a loss of log data to Syslog. If this is occurring, you will typically see “Suppressed *number* messages from ....” messages, usually in `/var/log/messages`. Most Syslog implementations have settings to control the rate that triggers this or even disable the action (though disabling is not recommended). You may need to make these adjustments or consider reducing what is logged to Syslog (it is not a best idea to log everything, especially for high levels of activity). Typically, this means adjusting the `/etc/systemd/journald.conf`’s `RateLimitInterval` and `RateLimitBurst` settings.

---

## Troubleshooting Server Failures

The server agent processes (`nwreglocal` and `nwregregion`) normally detect server failures and restart the server. You can usually recover from the failure and the server is not likely to fail again immediately after restarting. On rare occasions, the source of the server failure prevents the server from successfully restarting, and the server fails again as soon as it restarts. In such instances, perform the following steps:

---

**Step 1** If the server takes a significantly long time to restart, stop and restart the server agent. On:

- Windows:

```
net stop nwreglocal or nwregregion
net start nwreglocal or nwregregion
```

- Linux:

```
/etc/rc.d/init.d/nwreglocal stop or nwregregion stop
/etc/rc.d/init.d/nwreglocal start or nwregregion start
```

- Step 2** Keep a copy of all the log files. Log files are located in the *install-path/logs* directory on Linux, and the *install-path\logs* folder on Windows. The log files often contain useful information that can help isolate the cause of a server failure.
- Step 3** Use the TAC tool, as described in [Using the TAC Tool, on page 34](#), or save the core or user.dmp file, if one exists, depending on the operating system:
- **Windows**—The user.dmp file is located in the system directory, which varies depending on the Windows system. Search for this file and save a renamed copy.
  - **Linux**—The core file is located in the *install-path*. Save a renamed copy of this file that Cisco Prime Network Registrar does not overwrite.
- Step 4** On Windows, use the native event logging application to save the System and Application event logs to files. You can do this from the Event Viewer. These event logs often contain data that helps debug Cisco Prime Network Registrar server problems. For a description of the log messages for each server module, see the *install-path/docs/msgid/MessageIdIndex.html* file.

## Linux Troubleshooting Tools

You can also use the following commands on Linux systems to troubleshoot Cisco Prime Network Registrar. To:

- See all Cisco Prime Network Registrar processes:  
`ps -leaf | grep nwr`
- Monitor system usage and performance:  
`top`  
`vmstat`
- View login or bootup errors:  
`grep /var/log/messages*`
- View the configured interfaces and other network data:  
`ifconfig -a`

## Using the TAC Tool

There may be times when any amount of troubleshooting steps will not resolve your problem and you have to resort to contacting the Cisco Technical Assistance Center (TAC) for help. Cisco Prime Network Registrar provides a tool so that you can easily assemble the server or system error information, and package this data for TAC support engineers. This eliminates having to manually assemble this information with TAC assistance. The resulting package from this tool provides the engineers enough data so that they can more quickly and easily diagnose the problem and provide a solution.

The `cnr_tactool` utility is available in the bin directory of the Windows, and usrbin directory of the UNIX or Linux, installation directories. Execute the `cnr_tactool` utility:

```
> cnr_tactool -N username -P password [-d output-directory] [-n]
```

The output directory is optional and normally is the temp directory of the installation directories (in the /var path on Linux). You may specify the **-n** option to indicate that when the `cnr_exim` tool is run, it is run without exporting any resource records (this specifies the `-a none` option to `cnr_exim`). If you do not supply the username and password on the command line, you are prompted for them:

```
> cnr_tactool
user:
password:
[processing messages...]
```

The tool generates a packaged tar file whose name includes the date and version. The tar file contains all the diagnostic files.

## Using the statscollector Utility

Cisco Prime Network Registrar includes the **statscollector** utility, which reads the statistics collected by the CCM server on the local clusters. It has several options:

- Obtain the CCM server history from a cluster. You can get the history currently available and optionally continue collecting new history as it becomes available, and write this into a file. The file can later be processed or appended to. Note that, by default, the statscollector will continue running "forever" collecting the history. You can ask it to just get the current history and exit by specifying **-i 0**. If the file exists, it will read that history data to determine the "last" sample collected and start collecting more data from there (so, you can run it with **-i 0** every so often to just get the "new" history). Note that, if you use a file against a different cluster, you will get a mix of two clusters data which probably is not useful.

Example:

```
statscollector -C cluster -N user -P password stats.bin
```

- Generate XML (for import to Excel or other tools) of the statistics data either previously collected into a file or obtained from a cluster.

- Example (using pre-existing file):

```
statscollector -e stats.xml stats.bin
```

- Example (collecting from cluster):

```
statscollector -C cluster -N user -P password -e stats.xml
```

- Generate HTML (using Google Charts API) of the statistics data either previously collected into a file or obtained from a cluster. You can use the built-in or define your own charts, and have them plotted.

- Example (using pre-existing file):

```
statscollector -h stats.html stats.bin
```

- Example (collecting from cluster):

```
statscollector -C cluster -N user -P password -h stats.html
```

In Linux, you can run the statscollector from the following location:

**/opt/nwreg2/local/usrbin**

The following options are available:

Table 9: statscollector Options

Option	Description
<b>-C</b> <i>cluster:[port]</i>	The local cluster to connect to (default: localhost).
<b>-N</b> <i>admin</i>	The administrator account name.
<b>-P</b> <i>password</i>	The administrator password.
<b>-i</b> <i>interval</i>	The interval at which to poll for new statistics or 0 to exit after one read (default: 60 secs)
<b>-e</b> <i>file.xml</i>	Exports binary data file to XML format. <b>Note:</b> <b>-i</b> controls minimum interval between samples of data used (default: 1 sec).
<b>-h</b> <i>file.html</i>	Creates HTML file with statistics charts. <b>Note:</b> <b>-i</b> controls minimum interval between samples of data used (default: 1 sec).
<b>-c</b> <i>charts.txt</i>	For <b>-h</b> , optional chart definitions file.
<b>-s</b> <i>date time</i>	Ignores statistics samples before the specified date and time.
<b>-f</b> <i>date time</i>	Ignores statistics samples after the specified date and time.
<b>-w</b> <i>width X height</i>	Specifies chart width and height, in pixels (default: 800 X 400).
<b>-j</b> <i>name,value</i>	Specifies a Google annotation chart option.
<b>-t</b> <i>"title"</i>	Chart title (overrides default).
<b>-u</b> <i>infile.html</i>	Updates charts in the source HTML file. Requires the <b>-h</b> option.
<i>file</i>	The binary data file (required unless <b>-e</b> or <b>-h</b> specified). Data is appended to the file, if it exists.



**Note** When exporting to XML or HTML, the resulting files could be very large depending on the statistics collected. You may want to limit the data by using the **-i**, **-s**, and **-f** options. For example, **-i 300** means that the exported data is only reported every 5 minutes. But, **-s** and **-f** may be more effective to view the data for a specific (shorter) time interval.

# Troubleshooting and Optimizing the TFTP Server

You can set certain attributes to troubleshoot and optimize TFTP server performance.

## Related Topics

[Tracing TFTP Server Activity, on page 37](#)

[Optimizing TFTP Message Logging, on page 37](#)

[Enabling TFTP File Caching, on page 38](#)

## Tracing TFTP Server Activity

To trace TFTP server activity, set the *packet-trace-level* attribute to a value of 1 through 4, depending on the level of verbosity you want the TFTP server to use to write messages to the trace file. The trace files are located in the /logs subdirectory of the installation directory. Windows tracing goes to the file\_tftp\_1\_log file; Linux tracing goes to the /var/nwreg2/{local | regional}/logs/file\_tftp\_1\_log and file\_tftp\_1\_trace files.

Here are the trace levels, with each higher level being cumulative:

- **0**—Disables all server tracing (the default).
- **1**—Displays all the log messages in the trace file.
- **2**—Displays the client IP address and port number for all packets.
- **3**—Displays the packet header information.
- **4**—Displays the first 32 bytes of the packet.



---

**Note** Setting and getting the trace level only works if the TFTP server is started. Turn on packet tracing only for debugging purposes, and then not for any extended time, for performance reasons.

---

## Optimizing TFTP Message Logging

You can improve TFTP server performance by restricting logging and tracing. By default, the server logs error, warning, and informational messages to file\_tftp\_1\_log files. You can set the log levels using a few TFTP server parameters:

- **Log level** (use the *log-level* attribute)—Primary controller of server logging, which is preset to, and is best left at, level 3 (logs all error, warning, and informational messages). As with packet tracing, the higher logging levels are cumulative. If set to 0, no server logging occurs.
- **Log settings** (use the *log-settings* attribute)—This is the second level of logging control and takes only two values, *default* or *no-success-messages*. The *default* log setting does not alter the default value of log level 3 (error, warning, and informational messages). However, you may want to disable writing success informational messages, and thereby improve server performance, by changing the log settings to *no-success-messages*.

- **Log file count and size** (use the *log-file-count* attribute)—Sets how many log files to maintain and how large to allow them to get in the /logs directory. The default value is to maintain a maximum of ten files of 10 MB each.




---

**Note** Reload the TFTP server after changing these values.

---

## Enabling TFTP File Caching

You can improve TFTP server performance significantly by enabling file caching on the server. You must do this explicitly, because it is preset to disabled. You must also create and point to a file cache directory, and you can set the maximum size of this directory. Here are the steps:

- 
- Step 1** Determine where you want the TFTP cache files to go. This becomes a subdirectory of the TFTP home directory, which is preset to *install-path/data/tftp* (on Linux, it is */var/nwreg2/{local | regional}/data/tftp*). If you want a different location, set the *home-directory* attribute.
  - Step 2** Change to the TFTP home directory and create the cache directory, such as CacheDir, in the home directory, using the **mkdir CACHEDIR** command. Note that Cisco Prime Network Registrar ignores any files in any subdirectories of this cache directory.
  - Step 3** Use the *file-cache-directory* attribute to set up the TFTP server to point to the cache directory. You cannot use absolute path or relative path in the directory name. The *file-cache-directory* name is either appended to the path given in the *home-directory* or the default home directory path (if you do not specify one).
  - Step 4** Use the *file-cache-max-memory-size* attribute to set the maximum memory size, in bytes, of the cache. The preset value is 32 KB. Cisco Prime Network Registrar loads all files into cache that cumulatively fit this memory size. If you set the value to 0, Cisco Prime Network Registrar does not cache any data, even if you enable file caching.
  - Step 5** Copy all of the files you want cached into the cache directory, and not into any subdirectory. Because all files in this directory are loaded into cache, do not include large files.
  - Step 6** Enable the *file-cache* attribute to enable file caching, then reload the server. Cisco Prime Network Registrar logs the name of each cached file, and skips any it cannot load. It reads in all files as binary data and translates them as the TFTP client requests. For example, if a client requests a file as NetASCII, the client receives the cached data in that form.
  - Step 7** Writing to cache is not allowed. If you need to update a cache file, overwrite it in the cache directory, then reload the server.
-