



Maintaining Servers

This chapter explains how to administer and control your local and regional servers' operations through Cisco CNS Network Registrar's Web-based user interface (Web UI) and CLI (the **nrcmd** program).

Controlling Servers

You can control the Network Registrar servers as follows:

- Start—Load the database and start the server.
- Stop—Stop the server.
- Reload—Stop and restart the server.

Starting and stopping a server is self-explanatory. When you reload the server, 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.

In the Web UI, you can start, stop, and reload the protocol servers in one of two ways:

- If you are a local CCM or regional cluster administrator, click **Administration** on the Primary Navigation bar, then **Servers** on the Secondary Navigation bar, to open the Manage Servers page (see [Figure 6-1](#) for a local cluster example).

Figure 6-1 Manage Servers Page

Name	IP Address	Type	State	Health	Statistics	View Log	Start/Stop/Reload
Local Server Agent	127.0.0.1	CNRAGENT	running	10	[N/A]		[N/A] [N/A] [N/A]
Local CCM Server	127.0.0.1	CCM	running	9	[N/A]		[N/A] [N/A] [N/A]
Local DHCP Server	127.0.0.1	DHCP	running	10			
Local DNS Server	127.0.0.1	DNS	running	10			
Local TFTP Server	127.0.0.1	TFTP	disabled	0			

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, server agent, and Router Interface Configuration (RIC) server. However, you cannot stop, start, reload, or view statistics for them.

At the regional cluster, click the Start icon (➕) to start the server, the Stop icon (⊖) to stop it, or the Reload icon (🔄) to reload it.

- If you are a local cluster zone administrator, click **Zone** on the Primary Navigation bar, then **DNS Server** on the Secondary Navigation bar, to open the Manage DNS Server page. Click the Start icon (➕) to start the server, the Stop icon (⊖) to stop it, or the Reload icon (🔄) to reload it.
- If you are a local cluster DHCP administrator, click **DHCP** on the Primary Navigation bar, then **DHCP Server** on the Secondary Navigation bar, to open the Manage DHCP Server page. Click the Start icon (➕) to start the server, the Stop icon (⊖) to stop it, or the Reload icon (🔄) to reload it. This page also lets you view the related server properties for the DHCP server by clicking the Related Servers icon (🔗). (See the “[Listing Related Servers for DHCP Servers](#)” section on page 5-3.) These servers can be DNS, TFTP, or DHCP failover servers.

In the CLI for the local cluster:

- To start the server, use the **server type start** command (or simply **type start**, such as **dhcp start**).
- To stop the server, use the **server type stop** command (or simply **type stop**, such as **dhcp stop**). If stopping the server, it is advisable to save it first using the **save** command.
- To reload the server, use the **server type reload** command (or simply **type reload**, such as **dhcp reload**). Network Registrar stops the server you chose, loads the configuration data, and then restarts the server.



Note

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

Logging Server Events

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

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



Tip

To avoid filling up the Windows Event Viewer and preventing Network Registrar from running, in the Event Log Settings, check the **Overwrite Events as Needed** box.

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.

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

- **DNS**—Click **Zone** on the Primary Navigation bar, then **DNS Server** on the Secondary Navigation bar, to open the Manage DNS Server page. Click the name of the server to open the Edit DNS Server page. Expand the Logging attributes section of the page to view the log settings. Make changes to these settings as desired, click **Modify Server**, then reload the server.
- **DHCP**—Click **DHCP** on the Primary Navigation bar, then **DHCP Server** on the Secondary Navigation bar, to open the Manage DHCP Server page. Click the name of the server to open the Edit DHCP Server page. Expand the Logging section (the first one on the page) to view the log settings. Make changes to these settings as desired, click **Modify Server**, then reload the server.

In the CLI, use the **dns set log-settings**, **dhcp set log-settings**, and **tftp set log-settings** commands for the respective servers.



Note

Warnings and errors go to Syslog on Solaris or the Event Viewer on Windows. See the Caution on page 6-2. For a description of the log messages for each server module, see the *install-path/docs/msgid/MessageIdIndex.html* file.

Log Files

Table 6-1 describes the Network Registrar log files in the *install-path/logs* directory.

Table 6-1 Log Files in .../logs Directory

Component	File in /logs Directory	Local/Regional	Content
Installation	install_cnr_log	Both	Log of installation process
Server agent	agent_server_1_log	Both	Log of server agent starts and stops
Port check	checkports_log	Both	Log of network ports
DNS server	name_dns_1_log	Local	Log of DNS activity
DHCP server	name_dhcp_1_log	Local	Log of DHCP activity
TFTP server	file_tftp_1_log file_tftp_1_trace	Local	Log of TFTP activity
RIC server	ric_server_log	Regional	Log of RIC server activity
CCM database	config_ccm_1_log	Both	Log of CCM configuration, starts, stops
Web UI	cnrwebui_log	Both	Log of Web UI state
Tomcat/Web UI (in cnrwebui subdirectory)	catalina_log.date.txt jsui_log.date.txt localhost_access_log.date.txt	Both	Log of CCM database for Tomcat server and Web UI; new files created daily, so to monitor disk usage, periodically archive old log files

Each component can generate a number of log files, each with a preconfigured maximum size of 1 MB. 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 four for the DNS, DHCP, and TFTP servers.

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

**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.

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)

Displaying State

A Network Registrar server can be in one of the following states for any sustained time:

- Running—Server is successfully running.
- Disabled—Server is disabled and not running.
- Stopped—Server was administratively stopped and is not running.
- Unconfigured—Server is not operational because of a configuration failure.

Displaying Health

You can display aspects of a server's health, or how well it is running. The following items can decrement the server's 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
- 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
- RIC server (regional cluster)

**Tip**

Use the existence of any descending health values as a reminder to check the log files for the server.

In both the local and regional cluster Web UIs, on the Primary Navigation bar, click **Administration.**, then **Servers** on the Secondary Navigation bar. Check the Manage Servers page for the state and health of each server (see [Figure 6-1 on page 6-1](#) for a DHCP server example).

In the local cluster CLI, use the `[server] type getHealth` command. The number 10 indicates the highest level of health, 0 that the server is not running.

**Tip**

On Solaris or Linux, you can run the `cnr_status` command, in the `install-path/usrbin/` directory, to see if your local cluster server is running. See the *Network Registrar Installation Guide*.

Displaying Statistics

To display server statistics, the server must be running. In the local cluster Web UI, go to the Manage DNS Server page or Manage DHCP Server page and click the Statistics icon (). On the Statistics for Server page (see [Figure 6-2](#) for a DHCP example), click the name of the statistic attribute to get popup help for its description.

Figure 6-2 Statistics for Server Page

Attribute	Value
start-time	Sat Jan 17 09:35:15 2004
total-discovers	0
total-requests	0
total-releases	0
total-offers	0
total-acks	0
total-naks	0
total-declines	0

In the CLI, use the `[server] type getStats` command. The statistics for the DNS and DHCP servers are encoded in curly braces followed by sets of digits, as described in [Table 6-2](#) for DNS, [Table 6-3](#) for DHCP, and [Table 6-4](#) for TFTP.

Table 6-2 DNS Statistics

Position	Attribute	Description
{ 1 }	id	Implementation ID (release and build information)
2	config-recurs	Recursion services—(1)available, (2)restricted, (3)unavailable
3	config-up-time	Process up time (seconds)
4	config-reset-time	Time since the last reset (seconds)
5	config-reset	Number of server resets
6	counter-auth-no-names	Number of queries returning “authoritative no such name” responses
7	counter-auth-no-data-resps	Number of queries returning “authoritative no such data” responses
8	counter-non-auth-datas	Number of queries answered nonauthoritatively (cached)
9	counter-non-auth-no-datas	Number of queries answered nonauthoritatively with no data
10	counter-referrals	Number of forwarded queries
11	counter-errors	Number of error responses
12	counter-rel-names	Number of requests received of names of only one label (relative names)
13	counter-req-refusals	Number of refused queries
14	counter-req-unparses	Number of unparseable requests
15	counter-other-errors	Number of aborted requests due to other errors

Table 6-3 DHCP Statistics

Position	Attribute	Description
{1}	start-time	Date and time of last server reload
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-nacks	Number of negative acknowledgement (NACK) packets sent
8	total-declines	Number of DECLINE packets received

Table 6-4 TFTP Statistics in the CLI


Position	Attribute	Description
{1}	id	Implementation ID (release and build information)
2	server-state	State of the server
3	server-start-time	Start date and time (in seconds)
4	server-reset-time	Reset date and time
5	server-time-since-start	Running time since last start
6	server-time-since-reset	Running time since last reset
7	total-packets-in-pool	Number of packets in the pool
8	total-packets-in-use	Number of packets the server is using
9	total-packets-received	Number of packets received since the last start or reload
10	total-packets-sent	Number of packets sent since the last start or reload
11	total-packets-drained	Number of packets read and discarded since the last start or reload
12	total-packets-dropped	Number of packets dropped since the last start or reload
13	total-packets-malformed	Number of packets received that were malformed since the last start or reload
14	total-read-requests	Number of packets read since the last start or reload
15	total-read-requests-completed	Number of read packets completed since the last start or reload
16	total-read-requests-refused	Number of read packets refused since the last start or reload
17	total-read-requests-ignored	Number of read packets ignored since the last start or reload
18	total-read-requests-timed-out	Number of read packets that timed out since the last start or reload

Table 6-4 TFTP Statistics in the CLI (continued)

Position	Attribute	Description
19	total-write-requests	Number of read packets that were write requests since the last start or reload
20	total-write-requests-completed	Number of write requests completed since the last start or reload
21	total-write-requests-refused	Number of write requests refused since the last start or reload
22	total-write-requests-ignored	Number of write requests ignored since the last start or reload
23	total-write-requests-timed-out	Number of write requests that timed out since the last start or reload
24	total-docsis-requests	Number of DOCSIS requests received since the last start or reload
25	total-docsis-requests-completed	Number of DOCSIS requests completed since the last start or reload
26	total-docsis-requests-refused	Number of DOCSIS requests refused since the last start or reload
27	total-docsis-requests-ignored	Number of DOCSIS requests ignored since the last start or reload
28	total-docsis-requests-timed-out	Number of DOCSIS requests that timed out since the last start or reload
29	read-requests-per-second	Number of read requests per second
30	write-requests-per-second	Number of write requests per second
31	docsis-requests-per-second	Number of DOCSIS requests per second


Displaying IP Address Usage

In the Web UI, you can look at the local or regional cluster's address space, or generate a subnet utilization or lease history report on the regional cluster, to determine IP address usage. These functions are available in both Web UIs by clicking **Address Space** on the Primary Navigation bar, if you have address space privileges on the local or regional cluster.

You can determine the current address space utilization by clicking the View icon () in the Current Usage column for the unified address space, address block, and subnet (see the [“Viewing Address Utilization for Address Blocks, Subnets, and Scopes”](#) section on page 18-11). You can also get the most current IP address utilization by querying the lease history (see the [“Querying IP Lease History Using the Web UI”](#) section on page 12-14). 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 on page 18-9). Also ensure that the particular DHCP server is running.

In the CLI, you can generate an IP address usage report using the **report** command. See the *Network Registrar CLI Reference* for additional options you can set.

Displaying Related Servers

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.

DNS Zone Distribution Servers

A DNS zone distribution simplifies creating multiple zones that share the same secondary server attributes. In the Web UI, you can view and set the primary and secondary DNS servers in a zone distribution:

- On the local cluster, click **Zone** on the Primary Navigation bar, then **Zone Distribution** on the Secondary Navigation bar. This opens the List 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.
- On the regional cluster, click **DNS Configuration** on the Primary Navigation bar, then **Zone Distributions** on the Secondary Navigation bar. 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 primary, authoritative, and secondary servers in the zone distribution.

You cannot use the CLI to view or set zone distributions.

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's 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.

In the Web UI:

- On the local cluster, on the Primary Navigation bar, click **DHCP**, then on the Secondary Navigation bar, click **Failover**. The List DHCP Failover Pairs page shows the main and backup servers in the failover relationship.
- On the regional cluster, on the Primary Navigation bar, click **DHCP Configuration**, then on the Secondary Navigation bar, click **Failover**. The List DHCP Failover Pairs page shows the main and backup servers in the failover relationship.

In the CLI, use the **dhcp getRelatedServers** command 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.

In the Web UI:

- On the local cluster, on the Primary Navigation bar, click **DHCP**, then on the Secondary Navigation bar, click **Scopes**. Click a scope name on the List/Add DHCP Scopes page to open the Edit DHCP Scope page. Halfway down the page, click **List Leases** to open the List DHCP Lease for Scope page.
- On the regional cluster, you can view the lease history. On the Primary Navigation bar, click **Address Space**, then on the Secondary Navigation bar, click **Lease History**. Set the query parameters, then click **Query Lease History**. (See the “[Running IP Lease Histories](#)” section on page 12-13.)

In the CLI, use the **lease list** command to view the properties of all the available leases.

Troubleshooting Servers

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

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 Network Registrar with failover resynchronization in progress.

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
```
- Solaris:


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

- Linux:


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

- Step 2** Keep a copy of all the log files. Log files are located in the *install-path/logs* directory on Solaris and 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 the “Using the TAC Tool” section on page 6-13, or save the core or user.dmp file, if one exists, depending on the operating system:
- On 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.
 - On Solaris and Linux—The core file is located in the *install-path*. Save a renamed copy of this file that 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 Network Registrar server problems. For a description of the log messages for each server module, see the *install-path/docs/msgid/MessageIdIndex.html* file.
-

Troubleshooting and Optimizing the TFTP Server

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

Tracing TFTP Server Activity

The TFTP server has two CLI commands that help create more output to logs and can be useful in troubleshooting, although this usually impacts performance. These commands set up server packet tracing. The **tftp getTraceLevel** command identifies the current trace level, which by default is 0, or no tracing. The **tftp setTraceLevel** command sets the packet tracing to a value between 0 and 4. The trace files are located in the /logs subdirectory of the installation directory. Windows tracing goes to the file_tftp_1_log file; Solaris and 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’s 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)—Master controller of server logging, which defaults 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 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 is to maintain a maximum of four files of one megabyte 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 disabled by default. 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 to store the TFTP cache files. This cache directory becomes a subdirectory of the TFTP home directory, which by default is *install-path/data/tftp* (on Solaris and 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 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 relative paths in the directory name, such as `.../cachedir`. If the directory does not exist, file caching cannot occur.
 - Step 4** Use the *file-cache-max-memory-size* attribute to set the maximum memory size, in bytes, of the cache. The default is 32 K bytes. Network Registrar loads all files into cache that cumulatively fit this memory size. If you set the value to 0, 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. 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.
-

Solaris and Linux Troubleshooting Tools

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

- See all Network Registrar processes:
`ps -leaf | grep nwr`
- Monitor system usage and performance:
`top`
`vmstat`
- View login or bootup errors:
 - On Solaris—`grep /var/adm/messages*`
 - On Linux—`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. 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]
```

The output directory is optional and normally is the temp directory of the installation directories (in the /var path on Solaris and Linux). If you do not supply the username and password on the command line, you are prompted for them:

```
> cnr_tactool
username:
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.

