



Backup and Recovery

This chapter explains how to maintain the Cisco Prime Network Registrar databases.

Related Topics

[Backing Up Databases, page 8-1](#)
[Troubleshooting Databases, page 8-9](#)

Backing Up Databases

Because the Cisco Prime Network Registrar databases do a variety of memory caching and can be active at any time, you cannot rely on third-party system backups to protect the database. They can cause backup data inconsistency and an unusable replacement database.

For this purpose, Cisco Prime Network Registrar provides a shadow backup utility, `cnr_shadow_backup`. Once a day, at a configurable time, Cisco Prime Network Registrar takes a snapshot of the critical files. This snapshot is guaranteed to be a consistent view of the databases.

Related Topics

[Syntax and Location, page 8-2](#)
[Backup Strategy, page 8-2](#)
[Database Recovery Strategy, page 8-4](#)
[Backing Up CNRDB Data, page 8-5](#)
[Backing Up all CNRDBs using tar or similar tools, page 8-5](#)
[Recovering CNRDB Data from Damaged Databases, page 8-6](#)
[Recovering CNRDB Data from Backups, page 8-7](#)
[Recovering all CNRDBs using tar or Similar Tools, page 8-8](#)
[Recovering single CNRDB from tar or similar tools, page 8-9](#)
[Virus Scanning While Running Cisco Prime Network Registrar, page 8-9](#)
[Virus Scanning While Running Cisco Prime Network Registrar, page 8-9](#)

Syntax and Location

Be sure to understand that the notation “.../data/db” in the following sections refers to directories in the Cisco Prime Network Registrar product data location path, depending on the operating system:

- **Windows**—“.../data” means the data directory, which by default is **C:\NetworkRegistrar\{Local | Regional}\data**.
- **Solaris and Linux**—“.../data” means the data directory, which by default is **/var/nwreg2/{local | regional}/data**.

Cisco Prime Network Registrar database utility programs mentioned in the following sections are located in the “.../bin” directory, which you run as its full path name:

- **Windows**—“.../bin/program” means the program file in the bin directory, which by default is **C:\Program Files\Network Registrar\{Local | Regional}\bin\program** for a 32-bit OS and **C:\Program Files (x86)\Network Registrar\{Local | Regional}\bin\program** for a 64-bit OS.
- **Solaris and Linux**—“.../bin/program” means the program file in the bin directory, which by default is **/opt/nwreg2/local/usrbin/program** or **/opt/nwreg2/regional/usrbin/program**.



Note

Use only the approved utilities for each type of database. In Windows, if you want to run the utility from outside the installed path, you must set the CNR_HOME environment variable.

Backup Strategy

The backup strategy involves either:

- Making CCM perform a nightly shadow backup for you (See the [“Setting Automatic Backup Time” section on page 8-3](#)) and using the shadow backups for permanent backup and then doing an explicit backup - either using the `cnr_shadow_backup` utility and backing up the backup files (*.bak DBs) or

Shutting down Cisco Prime Network Registrar and performing a backup using TAR or other similar tools.

Using `cnr_shadow_backup` utility:

Use the `cnr_shadow_backup` utility to back up the following databases:

- **CNRDB databases**—...data/dhcp, ...data/dns/csetdb, ...data/dns/rrdb, ...data/cdns, ...data/leasehist, ...data/lease6hist, ...data/subnetutil, ...data/mcd, ...data/replica, and ...data/ccm/ndb



Note

If you change the location of the data directory, you must edit the **cnr.conf** file, which is located in `.../conf` (see the [“Modifying the cnr.conf File” section on page 7-25](#)). Change the **cnr.datadir** variable to the full path to the data directory. For example, the following is the default value on Windows:

```
cnr.datadir=C:\\NetworkRegistrar\\{Local | Regional}\\data
```

The most basic component of a backup strategy is the daily shadow backup. When problems occur with the operational database, you might need to try recovering based on the shadow backup of the previous day. Therefore, you must recognize and correct any problems that prevent a successful backup.

The most common problem is disk space exhaustion. To get a rough estimate of disk space requirements, take the size of the `.../data` directory and multiply by 10. System load, such as usage patterns, application mix, and the load on Cisco Prime Network Registrar itself, may dictate that a much larger reserve of space be available.

You should regularly archive existing shadow backups (such as to tape, other disks, or other systems) to preserve them for possible future recovery purposes.

**Caution**

Using a utility on the wrong type of database other than the one recommended can cause database corruption. Use only the utilities indicated. Also, never use the database utilities on the operational database, only on a copy.

Related Topics

[Setting Automatic Backup Time, page 8-3](#)

[Performing Manual Backups, page 8-3](#)

[Using Third-Party Backup Programs with `cnr_shadow_backup`, page 8-3](#)

Setting Automatic Backup Time

You can set the time at which an automatic backup should occur by editing the `cnr.conf` file (in `.../conf`). Change the `cnr.backup-time` variable to the hour and minute of the automatic shadow backup, in 24-hour `HH:MM` format, then restart the server agent. For example, the following is the preset value:

```
cnr.backup-time=23:45
```

Performing Manual Backups

You can also initiate a manual backup with the `cnr_shadow_backup` utility, which requires root privileges. Enter the `cnr_shadow_backup` command at the prompt to perform the backup.

**Note**

To restore DHCP data from a failover partner that is more up to date than a backup, see the [“Restoring DHCP Data from a Failover Server”](#) section on page 8-13.

Using Third-Party Backup Programs with `cnr_shadow_backup`

You should avoid scheduling third-party backup programs while `cnr_shadow_backup` is operating. Third-party backup programs should be run either an hour earlier or later than the `cnr_shadow_backup` operation. As described in the [“Setting Automatic Backup Time”](#) section on page 8-3, the default shadow backup time is daily at 23:45.

Configure third-party backup programs to skip the Cisco Prime Network Registrar operational database directories and files, and to back up only their shadow copies.

The operational files are listed in the [“Backup Strategy”](#) section on page 8-2. On Solaris and Linux, Cisco Prime Network Registrar also maintains lock files in the following directories:

- Cisco Prime Network Registrar server processes—`/var/nwreg2/local/temp/np_destiny_trampoline` or `/var/nwreg2/regional/temp/np_destiny_trampoline`

The lock files are recreated during a reboot. These files are important while a system is running. Any maintenance process (such as virus scanning and archiving) should exclude the temporary directories, operational database directories, and files.

Windows does not maintain lock files, but uses named-pipes instead.

Database Recovery Strategy

Cisco Prime Network Registrar uses the CNRDB database. [Table 8-1](#) lists the types of CNRDB database that must be backed up and recovered.

Table 8-1 Cisco Prime Network Registrar Databases for Recovery

Subdirectory	Cluster	Type	Description
mcd	local	CNRDB	MCD change log data. Only exists for upgrades from pre 8.0 databases as long as there is MCD change log history that has not been trimmed.
ccm	local, regional	CNRDB	Central Configuration Management database. Stores local centrally managed cluster and the SNMP server data.
dns	local	CNRDB	DNS database. Stores zone state information, names of protected RRs, and zone configuration data for the DNS server.
cdns	local	CNRDB	Caching DNS database. Stores the initial DNSSEC root trust anchor and root hints.
dhcp	local	CNRDB	DHCP database. Stores lease state data for the DHCP server.
dhcpeventstore	local		Queue that Cisco Prime Network Registrar maintains to interact with external servers, such as for LDAP and DHCPv4 DNS Update interactions. Recovery is not necessary.
tftp	local		Default data directory for the TFTP server. Recovery is not necessary.
replica	regional	CNRDB	Stores replica data for the local clusters.
lease6hist	regional	CNRDB	DHCPv6 lease history database.
leasehist	regional	CNRDB	DHCPv4 lease history database.
subnetutil	regional	CNRDB	Subnet utilization database.

The general approach to recovering a Cisco Prime Network Registrar installation is:

1. Stop the Cisco Prime Network Registrar server agent.
2. Restore or repair the data.
3. Restart the server agent.
4. Monitor the server for errors.

After you are certain that you executed a successful database recovery, always manually execute the `cnr_shadow_backup` utility to make a backup of the current configuration and state.

Backing Up CNRDB Data

In the case of the CNRDB databases, the `cnr_shadow_backup` utility copies the database and all log files to a secondary directory in the directory tree of the installed Cisco Prime Network Registrar product. For:

- **DHCP**—The operational database is in the `.../data/dhcp/ndb` and `.../data/dhcp/clientdb` directories, with the log files in the `.../data/dhcp/ndb/logs` directory. The shadow copies are in the `.../data/dhcp.bak/ndb` directory.
- **DNS**—The operational database is in the `.../data/dns/rrdb` directory. The important operational components are the High-Availability (HA) DNS is in the `.../data/dns/hadb` directory, with log files in the `.../data/dns/hadb/logs` directory. The shadow copies are in the `.../data/dns.bak` directory.
- **CDNS**—The operational database is in the `.../data/cdns` directory. The shadow copies are in the `.../data/cdns.bak` directory.
- **CCM**—The operational database and log files are in the `.../data/ccm/ndb` directory. The shadow copies are in the `.../data/ccm.bak` directory.
- **MCD change log**—The operational database and log files are in the `.../data/mcd/ndb` directory. The shadow copies are in the `.../data/mcd.bak` directory. MCD Change Log database may not exist if there are no change log entries. Also, the database is deleted when the MCD change log history is trimmed or when there is no MCD change log data to begin with.
- **Lease history**—The operational database and log files are in the `.../data/leasehist` and `.../data/lease6hist` directories. The shadow copies are in the `.../data/leasehist.bak` and `.../data/lease6hist.bak` directories.
- **Subnet utilization**—The operational database and log files are in the `.../data/subnetutil` directory. The shadow copies are in the `.../data/subnetutil.bak` directory.
- **Replica**—The operational database and log files are in the `.../data/replica` directory.

The actual file naming convention is:

- **Database**—`dhcp.ndb` and `dns.ndb`.
- **Log files**—`log.0000000001` through `log.9999999999`. The number of files varies with the rate of change to the server. There are typically only a small number. The specific filename extensions at a site vary over time as the database is used. These log files are not humanly readable.

Backing Up all CNRDBs using tar or similar tools

This section describes the procedure for backing up all Cisco Prime Network Registrar databases using tar or similar tools.

-
- Step 1** Shut down Cisco Prime Network Registrar.
- Backups cannot be done using tar or similar tools if Cisco Prime Network Registrar is running.
- Step 2** Back up the entire data directory and subdirectories:
- ```
> /var/nwreg2/local/data or /var/nwreg2/regional/data
> /opt/nwreg2/*/conf
```
- Step 3** Restart Cisco Prime Network Registrar when the backup is complete.
-

**Note**

Technically the backups do not need to include the \*.bak directories (and subdirectories of those directories) as those contain nightly shadow backups. However, unless your available storage space is severely limited, we recommend a full backup of the entire data directory (and subdirectories) including the shadow backups.

## Recovering CNRDB Data from Damaged Databases

This section describes a procedure that recovers any or all CNRDB type databases. Depending on the event that caused the database corruption, you can restore the database to a healthy state by using the current data. This is the best option. Always attempt recovery on a copy of the database file and associated log files, never on the operational files. This is a simple file copy operation, distinct from a shadow backup. Also, never attempt a recovery while Cisco Prime Network Registrar is running.

In most cases, you can use the log files that accompany the databases (such as the DHCP log files in `.../data/dhcp/ndb/logs`) to repair a failed server database. You can do so because the log files journal all database activity. You should never move, rename, or delete these log files, even after successfully completing a recovery. In fact, the recovery process uses copies rather than the originals of these files.

**Caution**

It is possible to damage the CNRDB database files without the damage being immediately obvious. Such damage could occur if you (a) inappropriately delete log files; (b) mix pre- and post-recovery database and log files; or (c) attempt to recover database files currently in use by an application. For the CNRDB database, use the `cnrdb_archive`, `cnrdb_recover`, and `cnrdb_verify` utilities.

Use the `cnrdb_recover` utility (see the “Using the `cnrdb_recover` Utility” section on page 8-12), included in the Cisco Prime Network Registrar product distribution, for database recovery. Use this tool with care. You should never use it directly on an operational database, or on files another application is concurrently accessing. On a successful database recovery, do not intermingle the recovered files (database file and log files) with files from another source, such as the operational database or shadow backups. Recovered database files acquire state information that make them incompatible with older database files.

- 
- Step 1** Stop the Cisco Prime Network Registrar server agent. This stops all the protocol servers. Ensure that enough disk space is available for a copy of the database files, plus a 15% safety margin.
- On Solaris, you can use the `df -k` utility to check your disk space, then `stop` to stop the server agent:
- ```
> df -k
> /etc/init.d/nwreglocal stop
```
- Step 2** Create a backup directory, named `backup`, outside the Cisco Prime Network Registrar installation tree. On Windows, this could be `C:\temp\backup`; on Solaris and Linux, this could be `/tmp/backup`.
- As a precaution, a copy of the directory tree of the current database that you are going to repair will get copied here automatically.
- For example, use `mkdir` on Solaris to create a backup directory:
- ```
> mkdir /tmp/backup
```
- Step 3** Copy the database subdirectories you want to restore under `.../data` to the backup directory. For example, to recover the DHCP database, recursively copy the `.../data/dhcp` directory and its subdirectories to `/tmp/backup`:

```
> cp -rp /var/nwreg2/local/data/dhcp /tmp/backup
```

When the copy is completed, double check that the database file and all log files were copied correctly. Do not allow these files to be modified in any way. Do not run any utilities or servers on these files.

**Note**

The log files must not be copied or moved while trying to repair the database because the "set\_lg\_dir logs" in the DB\_CONFIG file provides the information as to where the log files are located (in the log subdirectory). This enables the CNRDB utilities to find the log files without your having to copy or move the log files to any location. The relative path is used in the DB\_CONFIG file so that it is easier to move the directories around.

**Step 4** Repair the database:

- a. From the database file directory, run the **cnrdb\_recover** program, using the **-c** and **-v** options. It is helpful to use **-v** in that it displays output in the absence of errors (see the [“Using the cnrdb\\_recover Utility”](#) section on page 8-12). For example:

```
> cd /var/nwreg2/local/data/dhcp/ndb
> /opt/nwreg2/local/bin/cnrdb_recover -v -c
db_recover: Finding last valid log LSN: file: 1 offset 95181
db_recover: Recovery starting from [1][28]
db_recover: Recovery complete at Mon Jun 19 18:44:15 2006
db_recover: Maximum transaction ID 800000009 Recovery checkpoint [1][95229]
db_recover: Recovery complete at Mon Jun 19 18:44:15 2006
db_recover: Maximum transaction ID 800000000 Recovery checkpoint [1][95529]
```

- b. Run the **cnrdb\_verify** utility for each of the servers. There is no output if the verification is successful (see the [“Using the cnrdb\\_verify Utility”](#) section on page 8-12). For example:

```
> cd /var/nwreg2/local/data/dhcp/ndb
> /opt/nwreg2/local/bin/cnrdb_verify dhcp.ndb
```

- c. Optionally, for additional confidence, run the **cnrdb\_archive** utility:

- **cnrdb\_archive -l**—Lists all log files
- **cnrdb\_archive -s**—Lists the database file

- d. If there are any indications that an error occurred, proceed to restore the database from a backup, as described in the [“Recovering CNRDB Data from Backups”](#) section on page 8-7.

**Step 5** For DHCP only, delete the files in the `.../data/dhcpeventstore` directory.

**Step 6** Restart Cisco Prime Network Registrar.

## Recovering CNRDB Data from Backups

If there are any indications, such as server log messages or missing data, that database recovery was unsuccessful, you may need to base a recovery attempt on the current shadow backup (in the Cisco Prime Network Registrar installation tree). To do this:

**Step 1** Stop the Cisco Prime Network Registrar server agent.

**Step 2** Move the operational database files to a separate temporary location.

**Step 3** Copy each `.../data/name.bak` directory to `.../data/name`; for example, copy `.../data/ccm.bak` to `.../data/ccm`.




---

**Note** If you set the `cnr.dbrecover` variable to `false` in the `cnr.conf` file to disable recovery during the `cnr_shadow_backup nightly backup`, you must also do a recovery as part of these steps.

---

**Step 4** Rename the files.

The CNRDB database maintains centrally managed configuration data that is synchronized with the server configuration databases.

**Step 5** Create a new data directory and then `untar` or recover the backed up directory.

We recommend that you run the DB directory and recovery tools to ensure that the databases are good.




---

**Note** Ensure that the logs subdirectory is present in the same directory or the logs path is mentioned in the `DB_CONFIG` file.

---

**Step 6** Restart the server agent.

---




---

**Note** If the recovery fails, perhaps because the current shadow backup is simply a copy of corrupted files, use the most recent previous shadow backup. This illustrates the need to regularly archive shadow backups. You cannot add operational log files to older shadow backup files. All data added to the database since the shadow backup was made will be lost.

---

After a successful database recovery, initiate an immediate backup and archive the files using the `cnr_shadow_backup` utility (see the [“Performing Manual Backups”](#) section on page 8-3).

## Recovering all CNRDBs using tar or Similar Tools

This section describes the procedure for recovering all Cisco Prime Network Registrar databases using `tar` or similar tools.

---

**Step 1** Shut down Cisco Prime Network Registrar. Run `/etc/init.d/nwreglocal stop` to ensure that Cisco Prime Network Registrar is down.

**Step 2** Rename the active data directory (such as `mv data old-data`).




---

**Note** You must have sufficient disk space for twice the size of the data directory (and all the files in it and its subdirectories). If you do not have sufficient disk space, move the active data directory to another drive.

---

**Step 3** Create a new data directory and then `untar` or recover the backed up directory.

We recommend that you run the CNRDB directory and recovery tools to ensure that the databases are good.

**Step 4** Start Cisco Prime Network Registrar.

---



**Note**

Technically the restores do not need to include the \*.bak directories (and subdirectories of those directories) as those contain nightly shadow backups. However, unless your available storage space is severely limited, we recommend a full restore of the entire data directory (and subdirectories) including the shadow backups.

## Recovering single CNRDB from tar or similar tools

This section describes the procedure for recovering single database using tar or similar tools.

- Step 1** Shut down Cisco Prime Network Registrar. Run `/etc/init.d/nwreglocal stop` to ensure that Cisco Prime Network Registrar is down.
- Step 2** Rename the active data directory (such as `mv data old-data`).

**Note**

You must have sufficient disk space for twice the size of the data directory (and all the files in it and its subdirectories). If you do not have sufficient disk space, move the active data directory to another drive.

- Step 3** Create a new data directory and then untar or recover only the files in that directory (and its subdirectories) from the backup.
- We recommend that you run the CNRDB integrity and recovery tools to ensure that the CNRDB are good.
- Step 4** Repeat [Step 2](#) to [Step 3](#) for other DBs that have to be recovered.
- Step 5** Start Cisco Prime Network Registrar.

## Virus Scanning While Running Cisco Prime Network Registrar

If you have virus scanning enabled on your system, it is best to configure it to exclude certain Cisco Prime Network Registrar directories from being scanned. Including these directories might impede Cisco Prime Network Registrar operation. The ones you can exclude are the `.../data`, `.../logs`, and `.../temp` directories and their subdirectories.

## Troubleshooting Databases

The following sections describe troubleshooting the Cisco Prime Network Registrar databases.

### Related Topics

- [Using the `cnr\_exim` Data Import and Export Tool, page 8-10](#)
- [Using the `cnrdb\_recover` Utility, page 8-12](#)
- [Using the `cnrdb\_verify` Utility, page 8-12](#)
- [Using the `cnrdb\_checkpoint` Utility, page 8-13](#)
- [Restoring DHCP Data from a Failover Server, page 8-13](#)

## Using the `cnr_exim` Data Import and Export Tool

The `cnr_exim` data import and export tool now supports the following for a user not constrained to a specific tenant:

- Exporting all the data
- Exporting the data specific to a tenant either with or without the core data
- Exporting and importing license related data
- Importing all of the data
- Importing the data specific to a tenant and optionally mapping it to a new tenant either with or without the core data. This allows you to build a base configuration for new tenants. When specifying tenant tags, the imported data is used to find the old tenant id and the current configuration is used to find the new tenant id.

Some of the advantages that come with the use of multi-tenant architecture are that you can move configurations for a tenant from one cluster to another to export a tenant template data and then import that data as another tenant.



### Note

A user constrained to a specific tenant can only export or import data for that tenant.

The `cnr_exim` tool also serves to export unprotected resource record information. However, `cnr_exim` simply overwrites existing data and does not try to resolve conflicts.



### Note

You cannot use `cnr_exim` tool for import or export of data from one version of Cisco Prime Network Registrar to another. It can be used only for import or export of data from or to the same versions of Cisco Prime Network Registrar.

Before using the `cnr_exim` tool, exit from the CLI, then find the tool on:

- **Windows**—...\`bin\cnr_exim.exe`
- **Solaris and Linux**—.../`usrbin/cnr_exim`

You must reload the server for the imported data to become active.

Note that text exports are for reading purposes only. You cannot reimport them.

The text export prompts for the username and password (the cluster defaults to the local cluster). The syntax is:

```
> cnr_exim -e exportfile [-N username -P password -C cluster]
```

To export (importable) raw data, use the `-x` option:

```
> cnr_exim -e exportfile -x
```

To export DNS server and zone components as binary data in raw format, use the `-x` and `-c` options:

```
> cnr_exim -e exportfile -x -c "dnsserver,zone"
```

The data import syntax is (the import file must be in raw format):

```
> cnr_exim -i importfile [-N username -P password -C cluster]
```

You can also overwrite existing data with the `-o` option:

```
> cnr_exim -i importfile -o
```

Table 8-2 describes all the qualifying options for the `cnr_exim` tool.

**Table 8-2** *cnr\_exim Options*

| Option                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-a value</code>        | Allows exporting and importing of protected or unprotected RRs. Valid <i>values</i> are:<br><b>protectedRR</b><br><b>unprotectedRR</b><br>On export or import, all RRs are exported by default, so you must use a value to export or import just the protected or unprotected RRs.                                                                                                                                                                                                                                                                                              |
| <code>-c "components"</code> | Imports or exports Cisco Prime Network Registrar components, as a quoted, comma-delimited string. Use <code>-c help</code> to view the supported components. User are not exported by default; you must explicitly export them using this option, and they are always grouped with their defined groups and roles. Secrets are never exported.<br><b>Note</b> After you import administrator names, you must set new passwords for them. If you export groups and roles separately from usernames (which are not exported by default), their relationship to usernames is lost. |
| <code>-C cluster</code>      | Imports from or exports to the specified cluster. Preset to <b>localhost</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <code>-e exportfile</code>   | Exports the configuration to the specified file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <code>-h</code>              | Displays help text for the supported options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <code>-i importfile</code>   | Imports the configuration to the specified file. The import file must be in raw format.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <code>-N username</code>     | Imports or exports using the specified username.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <code>-o</code>              | When used with the <code>-i</code> (import) option, overwrites existing data.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <code>-p port</code>         | Port used to connect to the SCP server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <code>-P password</code>     | Imports or exports using the specified password.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <code>-t exportfile</code>   | Specifies a file name to export to, exports data in s-expression format.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <code>-v</code>              | Displays version information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <code>-x</code>              | When used with the <code>-e</code> (export) option, exports binary data in (importable) raw format.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <code>-d</code>              | Specifies the directory path of <code>cnr_exim</code> log file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <code>-f taglid</code>       | Specifies the source tenant. Valid for export and import.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <code>-g taglid</code>       | Specifies the destination tenant. Valid for import only. The tenant-id can not be changed when exporting data, only when the data is imported.)                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <code>-b</code>              | Specifies that the core (base) objects are to be included in the import/export. This includes all objects either with an explicit tenant-id of 0 and those that have no tenant-id attribute                                                                                                                                                                                                                                                                                                                                                                                     |
| <code>-w</code>              | Specifies the view tag to export. This option allows the user to export zone and RRs data which has the same view tag as mentioned in "w" option. All other objects will not take this option into consideration and will be exported as earlier if it is used.                                                                                                                                                                                                                                                                                                                 |

## Using the `cnrdb_recover` Utility

The `cnrdb_recover` utility is useful in restoring the Cisco Prime Network Registrar databases to a consistent state after a system failure. You would typically use the `-c` and `-v` options with this command (Table 8-3 describes all of the qualifying options). The utility is located in the installation bin directory.

**Table 8-3** *cnrdb\_recover Options*

| Option              | Description                                                                                                                                                                                                                                                                    |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-c</code>     | Performs a catastrophic recovery instead of a normal recovery. It not only examines all the log files present, but also recreates the <code>.ndb</code> (or <code>.db</code> ) file in the current or specified directory if the file is missing, or updates it if is present. |
| <code>-e</code>     | Retains the environment after running recovery, rarely used unless there is a <code>DB_CONFIG</code> file in the home directory.                                                                                                                                               |
| <code>-h dir</code> | Specifies a home directory for the database environment. By default, the current working directory is used.                                                                                                                                                                    |
| <code>-t</code>     | Recovers to the time specified rather than to the most current possible date. The time format is <code>[[CC]YY]MMDDhhmm[.ss]</code> (the brackets indicating optional entries, with the omitted year defaulting to the current year).                                          |
| <code>-v</code>     | Runs in verbose mode.                                                                                                                                                                                                                                                          |
| <code>-V</code>     | Writes the library version number to the standard output, and exits.                                                                                                                                                                                                           |

In the case of a catastrophic failure, restore a snapshot of all database files, along with all log files written since the snapshot. If not catastrophic, all you need are the system files at the time of failure. If any log files are missing, `cnrdb_recover -c` identifies the missing ones and fails, in which case you need to restore them and perform the recovery again.

Use of the catastrophic recovery option is highly recommended. In this way, the recovery utility plays back all the available database log files in sequential order. If, for some reason, there are missing log files, the recovery utility will report errors. For example, the following gap in the log files listed:

```
log.0000000001
log.0000000053
```

results in the following error that might require you to open a TAC case:

```
db_recover: Finding last valid log LSN:file:1 offset 2411756
db_recover: log_get: log.0000000002: No such file or directory
db_recover: DBENV->open: No such for or directory
```

## Using the `cnrdb_verify` Utility

The `cnrdb_verify` utility is useful for verifying the structure of the Cisco Prime Network Registrar databases. The command requires a file parameter. Use this utility only if you are certain that there are no programs running that are modifying the file. Table 8-4 describes all its qualifying options. The utility is located in the installation bin directory. The syntax is described in the usage information when you run the command:

```
C:\Program Files\Network Registrar\Local\bin>cnrdb_verify
usage: db_verify [-NoqV] [-h dir] [-P password] file
```

Table 8-4 *cnrdb\_verify Options*

| Option                   | Description                                                                                                                                     |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-h dir</code>      | Specifies a home directory for the database environment. By default, the current working directory is used.                                     |
| <code>-N</code>          | Prevents acquiring shared region locks while running, intended for debugging errors only, and should not be used under any other circumstances. |
| <code>-o</code>          | Ignores database sort or hash ordering and allows <code>cnrdb_verify</code> to be used on nondefault comparison or hashing configurations.      |
| <code>-P password</code> | User password, if the file is protected.                                                                                                        |
| <code>-q</code>          | Suppresses printing any error descriptions other than exit success or failure.                                                                  |
| <code>-V</code>          | Writes the library version number to the standard output, and exits.                                                                            |

## Using the `cnrdb_checkpoint` Utility

The `cnrdb_checkpoint` utility is useful in setting a checkpoint for the database files so as to keep them current. The utility is located in the installation bin directory. The syntax is described in the usage information when you run the command:

```
C:\Program Files\Network Registrar\Local\bin>cnrdb_checkpoint ?
usage: db_checkpoint [-lVv] [-h home] [-k kbytes] [-L file] [-P password] [-p min]
```

## Restoring DHCP Data from a Failover Server

You can restore DHCP data from a failover server that is more current than the result of a shadow backup. Be sure that the failover partner configurations are synchronized, then, on the failover partner:

### On Windows

1. Set the default path; for example:
 

```
SET PATH=%PATH%;.;C:\PROGRA~1\NETWOR~1\LOCAL\BIN
```
2. Stop the server agent:
 

```
net stop "Network Registrar Local Server Agent"
```
3. Delete the eventstore, ndb, and logs directories:
 

```
del C:\NetworkRegistrar\Local\data\dhcpeventstore*. *
del C:\NetworkRegistrar\Local\data\dhcp\ndb\dhcp.ndb
del C:\NetworkRegistrar\Local\data\dhcp\ndb\logs*. *
```
4. Restart the server agent:
 

```
net start "Network Registrar Local Server Agent"
```

### On Solaris and Linux

1. Stop the server agent:
 

```
/etc/init.d/nwreglocal stop
```

2. Determine the processes running:

```
/opt/nwreg2/local/usrbin/cnr_status
```

3. Kill the remaining processes:

```
kill -9 pid
```

4. Delete the eventstore, ndb, and logs directories:

```
rm /var/nwreg2/data/dhcpeventstore/*.*
rm -r /var/nwreg2/data/dhcp/ndb/*
```

5. Restart the server agent:

```
/etc/init.d/nwreglocal start
```