Backup and Recovery

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

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.

Recommendation

When upgrading to 11.1 (or later) from a pre-11.1 version of Cisco Prime Network Registrar and when there are significant number of DHCPv6 leases (and/or DHCPv6 lease history records), customers SHOULD schedule a DHCP database dump and load (see Using the cnrdb_util Utility , on page 16) to reduce the size of the DHCPv4 database after the upgrade. The upgrade does NOT reduce the size of the original dhcp.ndb database when the DHCPv6 leases (active + history) are moved to the new dhcp6.ndb and the only way to reduce the size of the original database is to do a dump and load. Viewing the size of the dhcp6.ndb file (using the ls command) will give you an estimate as to the size by which the database can be reduced.

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. ".../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. ".../bin/*program*" means the program file in the bin directory, which by default is /opt/nwreg2/{local | regional}/usrbin/*program*.

Note Use only the approved utilities for each type of database.

Backup Strategy

The backup strategy involves either:

Making CCM perform a nightly shadow backup for you (See the Setting Automatic Backup Time, on page 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.

Manual Backup (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
- **Smart License databases**—...data/sanosync.data, ...data/sapiidsync.data, and ...data/satimeflagsync.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.

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
```

Note You must restart Cisco Prime Network Registrar for a change to **cnr.backup-time** to take effect.

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 Restoring DHCP Data from a Failover Server, on page 19.

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 Setting Automatic Backup Time, on page 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 Backup Strategy, on page 2. 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.

Backing Up CNRDB Data

In case of 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 databases are in the .../data/dhcp/ndb, .../data/dhcp/ndb6, and .../data/dhcp/clientdb directories, with the database log files in the logs subdirectories of these directories. The shadow copies are in the .../data.bak/dhcp/ndb, .../data.bak/dhcp/ndb6, and.../data.bak/dhcp/clientdb directories.

- **DNS**—The operational database is in the .../data/dns/rrdb directory, with the database log files in the logs subdirectory. The important operational components are the High-Availability (HA) DNS is in the .../data/dns/hadb directory, with the log files in the .../data/dns/hadb/logs directory. The shadow copies are in the .../data.bak/dns directory.
- **CCM**—The operational databases are in the .../data/ccm/ndb, .../data/ccm/rrdb, and .../data/ccm/clientdb directories, with the database log files in the logs subdirectories of these directories. The shadow copies are in the .../data.bak/ccm directory.
- **MCD change log**—The operational database and log files are in the .../data/mcd/ndb directory, with the database log files in the logs subdirectory. The shadow copies are in the .../data.bak/mcd 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, with the database log files in the logs subdirectories of these directories. The shadow copies are in the .../data.bak/leasehist and .../data.bak/lease6hist directories.
- **DHCP utilization**—The operational database and log files are in the .../data/subnetutil directory, with the database log files in the logs subdirectory. The shadow copies are in the .../data.bak/subnetutil directory.
- **Replica**—The operational database and log files are in the .../data/replica directory, with the database log files in the logs subdirectory.

Following table lists the database files in Cisco Prime Network Registrar.

Table 1: Database Files

Directory	Subdirectory	File Name
dhcp	.../data/dhcp/ndb	dhcp.ndb
	.../data/dhcp/ndb6	dhcp6.ndb
	.../data/dhcp/clientdb	*.db
dns	.../data/dns/csetdb	dnscset.db
	.../data/dns/hadb	dnsha.db
	.../data/dns/rrdb	*.db
ccm	.../data/ccm/clientdb	changelog.db config.db
	.../data/ccm/ndb	*.db
	.../data/ccm/rrdb	changelog.db config.db

The log files are listed as 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
> /var/nwreg2/local/conf or /var/nwreg2/regional/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.

Database Recovery Strategy

Cisco Prime Network Registrar uses the CNRDB database. The following table lists the types of CNRDB database that must be backed up and recovered.

Table 2: 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.

Subdirectory	Cluster	Type	Description
cdns	local		Caching DNS database. Stores the initial DNSSEC root trust anchor and root hints.
dhcp[1]	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	DHCP Utilization database which includes databases for subnets and prefixes separately.

[1] Restoring the DHCP databases (.../data/dhcp/ndb and .../data/dhcp/ndb6) from a backup is NOT RECOMMENDED. This is because, this data is constantly changing as the DHCP server is running (because of client activity and lease expirations either on this server or its partner). Therefore, restoring the DHCP ndb/ndb6 databases from a backup would set the clock back in time for the server, but not for clients. Hence, it is best to retain the DHCP server databases rather than recovering from a backup, or if recovery is needed, delete the databases and recover the current leases from the partner via failover (see Restoring DHCP Data from a Failover Server, on page 19).

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.

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 Performing Manual Backups, on page 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 **systemctl stop nwreglocal** 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 **systemctl stop nwreglocal** 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.

Recovering from Regional Cluster Database Issues

There is no high availability solution for the regional cluster. The regional cluster is not critical to the operation of the local clusters - except for licensing. If the worst happens and restoring from a backup (such as a nightly shadow backup) fails, the regional cluster can be rebuilt.

While the regional cluster databases are very reliable (as they are transaction based), there are some situations (for example, running out of disk space or physical disk issues such as bad blocks) that can result in database problems, where CCM is unable to start or unable to perform certain functions.

There are four main databases used by the regional cluster:

- The CCM database (ccm directory) which contains the configuration objects.
- The lease history databases (lease6hist and leasehist) which contain the lease history collected from local clusters (if enabled).
- The subnet utilization database (subnetutil) which contains the scope and prefix utilization history collected over time (if enabled).
- The replica database (replica) which contains the configuration periodically pulled from local clusters.

The following sections describe the steps used if one or more of these databases develop issues (this can be determined from the config_ccm_1_log file and errors reported there – possibly including the inability of the regional to start).

Note Before proceeding with any of these steps, you should first see if the Troubleshooting Databases, on page 12 section can help correct the database, and if not, confirm whether a recent backup is available that might be restored.

Handling Lease History Database Issues

The lease history databases can potentially grow very large depending on the period for which data is saved and the rate of client activity. If this database is corrupted and cannot be restored, one way to recover the regional cluster operation is to delete this database (this will cause loss of lease history).

Use the following steps:

Step 1 Stop the regional cluster.

Step 2 Delete (or rename) the lease6hist and/or leasehist database directories. Delete (or rename) only the database that has issues.

Note If you were able to restore one or both of these databases from a recent backup, you can copy the backup lease6hist and/or leasehist directories (and all files and directories below them) to replace the deleted (or renamed) databases.

Step 3 Start the regional cluster.

Note These steps may also be used if you decide to no longer want to collect lease history and wish to delete all history. Before performing Step 1, be sure to disable all lease history collection.

Handling Subnet Utilization Database Issues

The subnet/prefix utilization databases can potentially grow very large depending on the period for which data is saved, the frequency of polling, and the number of subnets/prefixes. If this database is corrupted and cannot be restored, one way to recover the regional cluster operation is to delete this database (this will cause loss of utilization history).

Use the following steps:

Step 1 Stop the regional cluster.

Step 2 Delete (or rename) the subnetutil database directory.

Note If you were able to restore the subnetutil database from a recent backup, you can copy the backup subnetutil directory (and all files and directories below it) to replace the deleted (or renamed) database directory.

Step 3 Start the regional cluster.

Note These steps may also be used if you decide to no longer want to collect utilization data and wish to delete all collected data. Before performing Step 1, be sure to disable all utilization history collection.

Handling Replica Utilization Database Issues

The replica database can easily be recreated from the local clusters (since it stores a copy of each local cluster's configuration). If this database is corrupted, the best way to deal with it is to delete this database.

Use the following steps:

Step 1 Stop the regional cluster.

Step 2 Delete (or rename) the replica database directory.

Note It is best not to restore just this database from a backup as it is easily rebuilt from the local clusters.

Step 3 Start the regional cluster.

Step 4 Initiate a pull of replica data from each local cluster (this will occur automatically for each local cluster within several hours, so you can also wait for it to occur).

It is usually a good idea to pull the (IPv4 and IPv6) address space (if using DHCP) and the zone data once the replica database has been updated to assure that the regional cluster is consistent with the local clusters.

Rebuilding the Regional Cluster

If the ccm database is corrupt, and recovery from a backup is not possible or rebuilding the indexes (for more details on the rebuild_indexes tool, contact the Cisco Technical Assistance Center (TAC)) does not resolve the issue, it may be necessary to completely rebuild the regional. In some cases, it may be necessary to rebuild the regional cluster on a new system.

If the existing regional cluster is operating, it may be possible to extract the configuration data. However, this is problematic as it may also extract old or corrupt data (and for some database corruptions, it may loop exporting the same data over and over). To do this, you can run the cnr_exim tool to export the configuration in binary mode (use the **-x** option). If successful, this can later be imported. However, not all data is imported and therefore, it is important to follow the steps below.

If this is a new system:

Step 1 Install the Cisco Prime Network Registrar regional cluster.

Step 2 Set up the admin account and add the licenses.

Step 3 Register all of the local clusters with the regional. This requires issuing the **license register** command. If the address and port of the regional have not changed, then there is no need to specify the regional server’s address and port.

Step 4 If you used cnr_exim to export data from the old regional cluster, you can import it now using cnr_exim.

Step 5 Skip the "existing regional cluster" steps and proceed with the "common steps" below.

If this is an existing regional cluster:

Step 1 Stop the regional cluster, if running.

Step 2 Delete the /var/nwreg2/regional/data directory (itself and all files and directories under it).

Note You can retain the lease6hist, leasehist, and/or subnetutil directories (and all files in or below these directories) if these databases have not been corrupted and you prefer to retain this historical information. If deleted, this historical data will be lost.

Note You MUST NEVER retain the replica database as its data will not be usable if the ccm database is deleted. Failure to delete the replica database can cause significant issues.

Step 3 Create an empty /var/nwreg2/regional/data directory (if entirely deleted or moved).

Step 4 Start the regional cluster.

Step 5 Set up the admin account and add the licenses.

Step 6 If you used cnr_exim to export data from the old regional cluster, you can import it now using cnr_exim.

Step 7 Restart the regional cluster (this is required to assure all services are running).

Step 8 Re-register all of the local clusters with the regional. This requires issuing the **license register** command (note that additional parameters are not needed as this will re-register with the existing regional information at the local - servers, IP address, and port).

Step 9 Continue with the common steps below.

Common steps (either for new or existing regional cluster):

Step 1 Assure that all of the replica data is up-to-date - this can be done by pulling the replicas for each local cluster (either in web UI or using the **cluster** *name* **updateReplicaData** command).

Step 2 Pull the v4 and v6 address space if using DHCP (either in web UI or using the **ccm pullAddressSpace** and **ccm pullIPv6AddressSpace** commands).

Step 3 Pull the zone data if using DNS (either in web UI or using the **ccm pullZoneData** command).

Step 4 Pull the administrators or other objects (policies, templates, and so on) as appropriate from one of the local clusters that has this information (either in web UI or using the **pull** subcommand).

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.

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 them 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 in the *install-path*/usrbin directory.

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
```

Starting Cisco Prime Network Registrar 11.2, the nrcmd text export/import feature allows the customer the ability to export config files in nrcmd text format and import a text version as well. For this feature, the option -n can be used to export a configuration in nrcmd format. This allows the configuration to be both human-readable, editable and be able to be imported. The importing happens using nrcmd.

The cnr_exim command can be run to export configuration in nrcmd format.

```
> cnr_exim –e exportfile –n –o –N username –P password –C cluster –c component –p CCM-Port
```

The nrcmd import command is:

```
> nrcmd –N username –P password –C cluster –b exportfile
```

The following table describes all the qualifying options for the **cnr_exim** tool.

Table 3: cnr_exim Options

Option	Description
–a	Allows exporting and importing of protected or unprotected RRs. Valid *values* are: **protectedRR**, **unprotectedRR**, and **none** **Export**: All RRs are exported by default, so you must explicitly specify the export of protected or unprotected RRs using the option "-a protectedRR", "-a unprotectedRR", or "-a none". If this option is not specified, all RRs are exported. **Import:** All RRs are imported by default, so you must explicitly specify the import of protected or unprotected RRs using the option "-a protectedRR" or " -a unprotectedRR". If this option is not specified, all RRs are imported.
–b	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.
–c	Imports or exports Cisco Prime Network Registrar components, as a quoted, comma-delimited string. Use **–c help** 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. **Note** 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.
–C *cluster*	Imports from or exports to the specified cluster. Preset to **localhost**.
–d	Specifies the directory path of cnr_exim log file.
–e *exportfile*	Exports the configuration to the specified file.

Option	Description
–f	Specifies the source tenant. Valid for export and import.
–g	Specifies the destination tenant. Valid for import only. The *tenant-id* can not be changed when exporting data, only when the data is imported.)
–h	Displays help text for the supported options.
–i *importfile*	Imports the configuration to the specified file. The import file must be in raw format.
–n	When used with the -e (export) option, exports data in nrcmd command format. To import the exported data, the user must use "nrcmd -b < exportfile". All the exporting limiting options (-c, etc) should apply to the nrcmd format.
–N *username*	Imports or exports using the specified username.
–o	When used with the **–i** (import) option, overwrites existing data.
–p *port*	Port used to connect to the SCP server.
–P *password*	Imports or exports using the specified password.
–t *exportfile*	Specifies a file name to export to, exports data in s-expression format.
–v	Displays version information
–w	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.
–x	When used with the **–e** (export) option, exports binary data in (importable) raw format.

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. The following table describes all of the qualifying options. The utility is located in the *install-path*/bin directory.

Table 4: cnrdb_recover Options

Option	Description
–c	Performs a catastrophic recovery instead of a normal recovery. It not only examines all the log files present, but also recreates the .ndb (or .db) file in the current or specified directory if the file is missing, or updates it if is present.
–e	Retains the environment after running recovery, rarely used unless there is a DB_CONFIG file in the home directory.
–h *dir*	Specifies a home directory for the database environment. By default, the current working directory is used.

Option	Description
–t	Recovers to the time specified rather than to the most current possible date. The time format is *[[CC]YY]MMDDhhmm[.ss]* (the brackets indicating optional entries, with the omitted year defaulting to the current year).
–v	Runs in verbose mode.
–V	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. The following table describes all its qualifying options. The utility is located in the *install-path*/bin directory.

The syntax is described in the usage information when you run the command:

```
./cnrdb_verify

usage: cnrdb_verify [-mNoqV] [-b blob_dir] [-h home] [-P password] db_file ...
```

Table 5: cnrdb_verify Options

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

Option	Description
–V	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 *install-path*/bin directory.

The syntax is described in the usage information when you run the command:

```
./cnrdb_checkpoint
```

```
usage: cnrdb_checkpoint [-1Vv] [-h home] [-k kbytes] [-L file] [-m msg_pfx] [-P password][-p
 min]
```

Using the cnrdb_util Utility

The **cnrdb_util** utility is useful for dumping and loading the Cisco Prime Network Registrar databases. In addition, you can use this utility to shadow backup and recover the Cisco Prime Network Registrar databases, to clear the log files, as well as to change the database page size.

Stop Cisco Prime Network Registrar on the local cluster while using cnrdb_util utility.

The utility is located in the *install-path*/usrbin directory:

Important It is strongly recommended that a backup be done before performing any operation on the Cisco Prime Network Registrar databases. If existing backup files are to be retained, they must be backed up as well.

The **cnrdb_util** utility runs in two modes.

- **Interactive mode**—Prompts the user for operations and options.
- **Batch mode**—Requires information (both operation and options) as arguments while executing this utility.

The syntax is described in the usage information when you run the command:

```
./cnrdb_util -h
```

The following tables describe all of the qualifying operations and options.

Table 6: cnrdb_util Operations

Operation	Description
–d	Dump one or all Cisco Prime Network Registrar databases.
–l	Load one or all Cisco Prime Network Registrar databases.
–b	Create shadow backup of all Cisco Prime Network Registrar databases.
–r	Recover one or all Cisco Prime Network Registrar databases from shadow backup.

Operation	Description
–c	Cleanup sleepycat log files in one or all Cisco Prime Network Registrar databases.
–h	Display help text for the supported options.

Important You can perform only one operation at a time.

Table 7: cnrdb_util Options

Option	Description
-m { local | regional }	Specifies the Cisco Prime Network Registrar installation mode. If not specified, this inf file. If the file is not found, local mode is used by default.
-prog *path*	Specifies the path to the dump, load, or shadow backup executable. If not specified, this Prime Network Registrar installation path.
-db *db-path*	Specifies the path to the dump, load, or shadow backup executable. If not specified, this Prime Network Registrar installation path.
-db_pagesize *number*	Specifies the size of database pages (in bytes) to be used when creating new databases. The minimum page size is 512 bytes and the maximum page size is 64K bytes, and mus size is specified, a page size is selected based on the underlying filesystem I/O block size has a lower limit of 512 bytes and an upper limit of 16K bytes.) Usually the default is appropriate. However, large page sizes may not have good perform good sizes. You can determine the page size of the database by using the **cnrdb_stat** ut
-n { ccm | dhcp | dns | mcd | leasehist | lease6hist | replica | subnetutil | all }	Specifies the name of the source database for the '-d' dump, '-l' load, or '-r' recover operat will be performed on all databases present in database path. This option is not applicabl • Valid database names for local mode are { ccm | dhcp | dns | mcd | all } • Valid database names for regional mode are { ccm | dns | leasehist | lease6hist | rep
-s	Specifies that this program should attempt to stop the Cisco Prime Network Registrar S
-out *path*	Specifies the destination path for output files. If not specified, the source db path is used the '-b' backup and '-c' cleanup operations.

Important If the source and target directories are the same, the Dump and Load operations will delete the source files when the target files are created. This is done to minimize the disk space requirements when a dump/load operation is run to recapture the unused space in large database files.

Note The Dump operation will dump each database to a file in the specified location using the database file name appended by '.dbdump'. The Load operation will only load database files if a *.dbdump file is found; the name of the database file is the name without '.dbdump'.

Using the cnr_rpz_zone Utility

The cnr_rpz_zone utility allows the end user to take a list of domains and generate an RPZ zone along with the corresponding RRs. This utility will be included with a Cisco Prime Network Registrar install (that is, /opt/nwreg2/local/usrbin/cnr_rpz_zone). Below shown are the sample commands for the cnr_rpz_zone utility.

To create rpz specific zone:

```
./cnr_rpz_zone -f <txt file specify zone to be restricted> -Z <rpz zone name>
./cnr_rpz_zone -f zone.txt -Z rpz.zone.com
./cnr_rpz_zone -f zone.txt -Z <rpz zone name> -a <rpz-action>
./cnr_rpz_zone -f zone.txt -Z rpz.zone.com -a NXDOMAIN
```

To overwrite existing action in rpz zone use -o option:

```
./cnr_rpz_zone -f zone.txt -Z <rpz zone name> -o -a <rpz-action>
./cnr_rpz_zone -f zone.txt -Z rpz.zone.com -o -a NOERROR
```

To reload DNS server:

```
./cnr_rpz_zone -f zone.txt -Z <rpz zone name> -r
./cnr_rpz_zone -f zone.txt -Z rpz.zone.com -r
```

The following table describes all the qualifying options for the **cnr_rpz_zone** tool.

Table 8: cnr_rpz_zone Options

Option	Description
–a *RPZ_ACTION*	Specifies RPZ action to be performed. Valid options: NXDOMAIN, NODATA, NOOP, REDIRECT or DROP. Defaults to NXDOMAIN.
-d *RPZ_DATA*	Specifies the RR data and is required for the REDIRECT action. The RR data can contain A, AAAA or CNAME data.
–f *FILENAME*	Specifies the rpz filename to use to generate RPZ zone. This is a mandatory flag to use this utility.
–Z *ZONENAME*	Specifies the rpz zone name to use to generate RPZ zone. This is a mandatory flag to use this utility.
–o	Specifies if existing zone should be overwritten. Defaults to false.
–r	Specifies if the DNS server should be reloaded. Defaults to false.

Option	Description
–p *CLUSTER_PORT*	Specifies the CPNR cluster port. Defaults to 1234.
–C *CLUSTER*	Specifies the CPNR cluster to connect to other cluster. Defaults to localhost.
–N *USERNAME*	Specifies the CPNR admin username.
–P *PASSWORD*	Specifies the CPNR admin password.
–h	Prints a help screen.

The input file (-f FILENAME) should be a list of domains that you would want blocked. For example (i.e. rpz-blocked-domains.txt):

wrongdomain.com

baddomain.com

Running the cnr_rpz_zone command, generates an RPZ zone and the RPZ RRs.

If the zone exists, we will use the existing zone and only add name sets that are new. If the zone exists and the -o (overwrite) option is given, we will recreate the zone and drop all existing RRs and replace them with the generated RPZ RRs based on the input file.

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. Also, ensure that the following steps are run on the bad failover partner (that is, the one whose database is bad) and that you want to restore to.

1. Stop the server agent:

   ```
   systemctl stop nwreglocal
   ```

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/
   rm -r /var/nwreg2/data/dhcp/ndb6/
   ```

Warning When removing either DHCP databases, BOTH MUST be removed - the DHCPv4 (data/dhcp/ndb) or DHCPv6 (data/dhcp/ndb6) lease databases. Removing only one (and leaving the other) is unsupported and may produce unpredictable results.

5. Restart the server agent:

```
systemctl start nwreglocal
```