Managing High Availability DNS

A second primary server can be made available as a hot standby that shadows the main primary server. This configuration is called High-Availability (HA) DNS. The Cisco Prime Network Registrar web UI and CLI have features with which you can duplicate the primary setup required for HA DNS for the server pair. The server pair is responsible for detecting communication failures and the like. After the HA DNS is configured, the shadowing and error detection is done automatically. In a Cisco Prime Network Registrar deployment where Cisco Prime Network Registrar DHCP is updating Cisco Prime Network Registrar DNS, the failure detection and failover also happens automatically.

Introduction to HA DNS Processing

In normal state, both the main and backup primary servers are up and running. The main server processes all DNS updates from clients and sends all accepted updates to the hot standby backup. The main server will forward RR updates to the backup server and the backup server only accepts updates from the main in normal state. In normal states, updates from DDNS clients are ignored or dropped by a backup server. Both servers can respond to queries and zone transfer requests. The main and the backup partners always stay in communication to detect availability of the other.

If the main goes down, the backup waits a short time, then begins servicing the DNS updates from clients that the main would normally service and records the updates. When the main returns, the backup sends it the updates, and the main synchronizes with the backup any updates that were not sent and which it had before it went down.

Whenever you add a new zone, both the primary and the secondary servers must be reloaded to automatically synchronize with the HA backup.

The synchronization is done on a per-zone basis. This allows updates to all other zones while a given zone is in the process of getting synchronized.

If the hot standby backup goes down, the main waits a short time, then records the updates that the partner did not acknowledge. When the backup server comes back up, the main sends the recorded updates to the backup.

Both the main and backup can traverse the following states:

- **Startup**—The servers establish communication and agree on the HA version to use. In this state, the servers do not accept DNS updates or RR edits, and they defer scavenging, if enabled.
- **Negotiating**—Each server is waiting for the other to get ready to synchronize. In this state, DNS Updates and RR edits are not allowed.
- **Normal**—Both servers are up and healthy, exchanging DNS updates and heartbeat messages. The main accepts DNS updates and RR edits, sends RR Update messages to the backup, and performs history trimming and scavenging, if enabled. The backup ignores DNS updates, refuses RR edits, but processes RR Update messages from the main server. The backup also performs history trimming, but defers scavenging, if enabled. In this state, the synchronization takes place.
- **Communication-Interrupted**—The server goes into this state after not getting a response or request from the partner during the communication timeout (*ha-dns-comm-timeout*) period (preset to 30 seconds). The server continues listening for communication from the partner (they both send heartbeat messages every 12 seconds) and tries to connect, meanwhile accepting DNS updates and RR edits and disabling scavenging.
- **Partner-Down**—The server administrator notifies the partner that it will be down for an extended time. This manual intervention is possible only in Communication-Interrupted state. Either server continues listening for communication from the partner and tries to connect, accepts DNS updates and RR edits, and performs scavenging.

When a DNS server starts up, it:

1. Tries to establish a connection with its partner.
2. Transitions to Negotiating state.
3. Transitions to Normal state, after it receives a Negotiating response.

Once the server is in Normal state, the zone level synchronization begins. Zone synchronization is always managed by the Main HA server. The zones traverse through the following states:

- **Sync-Pending State**—A zone enters this state when the HA DNS server transitions to the normal state or if a manual sync is requested. In this state RR updates for the zone will be accepted on the main server, and forwarded to the backup server.
- **Synchronizing State**—The RR synchronization for the zone takes place in the synchronizing state. RR updates are not accepted, and notifies are disabled.
- **Sync-Complete State**—A zone transitions to this state from the synchronizing state once it has successfully synchronized resource record changes with its corresponding zone on the HA DNS backup. In this state, the zone on the HA DNS main server accepts all dynamic DNS update requests, allow resource record configuration changes, and re-enables notifies. Resource record modifications will be forwarded to the backup server.
- **Sync-Failed State**—A zone transitions to the sync-failed state from the synchronizing state if it fails to sync. The zone will accept resource record updates on the main server, and changes will be forwarded to the backup. The server will retry synchronizing the zone after ha-dns-zonesync-failed-timeout. A manual sync request or server restart will also restart zone synchronization.

HA DNS is fully integrated with DHCP servers, and the partners are updated when hosts get added to the network (see the *"Managing DNS Update" chapter in Cisco Prime Network Registrar 8.3 DHCP User Guide*). From the DHCP side of HA DNS, the DHCP server sends DNS updates to a single DNS server at a time.

DHCP autodetects the main being down and start sending updates to the backup. The DHCP server tries to contact the main DNS server, twice. It tries the backup partner if both of the attempts are unsuccessful.

The backup detects the main server down and starts accepting updates from DDNS clients. When the servers come up again, HA communication will be automatically established and the servers will get into Normal state where they carry out zone synchronization and make sure that both have the same RRs, etc.

If both DNS partners are communicating, the backup server drops the update, whereby the DHCP server times out and retries the main DNS server. If both servers are unreachable or unresponsive, the DHCP server continually retries each DNS partner every 4 seconds until it gets a response.

For zone level sync, an **Advanced** mode command is added in the local cluster Zone Commands page, if the local cluster is configured as the main HA server. The sync is run using the HA server algorithms by default. In **Expert** mode, the following three options are provided:

- Sync All RRs using the HA server synchronization algorithms
- Sync All RRs from Main to Backup
- Sync All RRs from Backup to Main

HA DNS status is modified to include the zone synchronization status. Status includes count and percentage of synchronized zones, zones pending synchronization, and zones that have failed synchronization.

Zone status has been modified to also include the HA synchronization status (ha-server-pending, sync-pending, sync-complete, synchronizing, or sync-failed), if HA is configured.

The *ha-dns-comm-timeout* attribute managed through the HA pair indicates the time required to determine if a partner is unreachable, after network communication is not acknowledged, which triggers the Communication-Interrupted state (see the description of this state in Introduction to HA DNS Processing, on page 1). The preset value is 30s. The server tries to communicate and then back off at multiples of the *ha-dns-comm-timeout* interval.

Creating High Availability DNS Pairs

The attributes needed to set up an HA DNS server pair from the main server are:

- ***ha-dns*** —Enabled or disabled. The preset value is enabled.
- ***main*** —cluster for the main primary DNS server.
- ***backup*** —cluster for the backup primary DNS server.

The specific IP addresses for the main or backup is specified only when the cluster IP is only used for management and DNS works on a different interface

Local Basic or Advanced and Regional Web UI

Step 1 Create a cluster for the backup server.

Step 2 From the **Deploy** menu, choose **HA** under the**DNS** submenu to open the View/Add HA DNS Server Pair page.

Step 3 Click the **Add HA Pair** icon in the **HA Pairs** pane to open the Add HA DNS Server dialog box.

Step 4 Enter the name of the server pair in the name field. This can be any identifying text string.

Step 5 Click the cluster name of the main DNS server in the Main Server drop-down list.

Note If you change the IP address of your local host machine, you must modify the localhost cluster (on the Edit Cluster page) to change the address in the IP Address field. Do not set the value to 127.0.0.1.

Step 6 Click the cluster name of the backup DNS server in the Backup Server drop-down list. This cannot be the same as the main server cluster. Set the *ha-dns-main-server* and *ha-dns-backup-server* attributes only if the server is configured with different interfaces for configuration management and update requests. (Configure the HA DNS protocol only with the interface used to service updates.)

Step 7 Click **Add HA DNS Server**.

Step 8 Once the server pair appears on the List/Add HA DNS Server Pair page, synchronize the servers:

a) Select the HA in the HA Pairs pane and click the Sync HA DNS Server Pair tab.
b) Choose the direction of synchronization (Main to Backup or Backup to Main).
c) Choose the operation type (Update, Complete, or Exact). See the table on the page for details on the operations for each operation type.
d) Click the **Report** button to display the prospective synchronization changes on the View HA DNS Sync Report page.
e) Click Run Complete to complete the synchronization.
f) Click **Return** to return to the List HA DNS Server Pairs page.

Step 9 Reload both DNS servers to begin HA communication.

CLI Commands

Create the HA DNS server pair (**ha-dns-pair** *name* **create** *mainaddr backupaddr*). Then synchronize the servers using **ha-dns-pair** *name* **sync**, specifying the synchronization operation (update, complete, or exact) and direction (main-to-backup or backup-to-main). Be sure to reload both DNS servers. For example:

```
nrcmd> ha-dns-pair examplehadnspair create localhost test-cluster

nrcmd> ha-dns-pair examplehadnspair sync exact main-to-backup

nrcmd> dns reload
```

See the **ha-dns-pair** command in the **CLIGuide.html** file in the /docs directory for syntax and attribute descriptions . The CLI provides an additional command for the DNS server to set the HA DNS partner down, if necessary, which is possible only while in Communication-Interrupted state:

```
nrcmd> dns setPartnerDown
```

The partner down is useful because it limits the bookkeeping data a server maintains, thus optimizing its performance. When both servers start communicating again, the sync sends all the zone RRs rather than trying to determine individual changes.

HA DNS Configuration Synchronization

This section describes the migration procedure used to migrate Cisco Prime Network Registrar product databases from the HA DNS main server to the HA DNS backup server. Throughout this procedure the source system is referred as DNS HA main server and destination as DNS HA backup server. When you enable the HA DNS with large DNS configuration, you will notice that the process takes long time to complete. This section provides a workaround, which you can use until the defect is addressed.

Danger To perform this process, you must have HA main server and HA backup server running on the same OS, Cisco Prime Network Registrar version, and DNS configuration.

Pre-install Cisco Prime Network Registrar on the HA DNS backup server

You need to pre-install Cisco Prime Network Registrar on the HA DNS backup system before migrating the database directory from the HA DNS main system, to reduce the time required during the Cisco Prime Network Registrar software installation process. During the installation process, the installer will verify whether any previous configuration is up to date with the Cisco Prime Network Registrar data schema for the version being installed. Even if the versions are identical, the time required to perform this verification can be avoided by pre-installing Cisco Prime Network Registrar on the HA DNS backup system.

Pre-migration Steps for HA DNS Main Server

You must ensure that the service of DHCP and TFTP servers are available and running on different systems, especially when there is a large DNS configuration. If the servers are found on the same system, the migration from HA DNS main server to backup server may cause DHCP or TFTP conflicts, and DHCP clients may be destabilized.

Follow the pre-migration steps as below:

Step 1 Disable the automatic start-on-reboot setting for the DHCP and TFTP server.

Note The default setting of start-on-reboot for the TFTP server is disabled.

Example:

```
nrcmd> server dhcp disable start-on-reboot
nrcmd> server tftp disable start-on-reboot
```

Step 2 Stop the Cisco Prime Network Registrar on the HA DNS main server using the Windows Service Control manager (Windows) or nwreglocal script in /etc/init.d (Linux and Solaris).

Step 3 Once the Cisco Prime Network Registrar is stopped by using Windows Process Manager (Windows) or ps command line utility (Linux/Solaris), navigate to the parent directory of the Cisco Prime Network Registrar data directory, InstallDir\Network Registrar\Local\ (Windows) or /var/nwreg2/local/ on (Linux/Solaris).

Step 4 Using tar or an equivalent compression utility, bundle up the contents of the data subdirectory. InstallDir is the directory where you have installed your Cisco Prime Network Registrar: tar -cvf cnrdatadir.tar data.

Tip Replace all the .bak database backup directories temporarily from HA DNS main server. The HA backup server does not need these backup directories and replacing them reduces the overall archive size. Be sure that you do not replace any other database files other than .bak; otherwise, the HA DNS backup cluster may not function properly.

Restart Cisco Prime Network Registrar on the HA DNS Main Server

Step 1 Restart the Cisco Prime Network Registrar servers on the HA DNS main system using the Windows Service Control manager (Windows) or nwreglocal script in /etc/init.d (Linux and Solaris).

Step 2 Restore the DHCP and TFTP server start-on-reboot attribute values to their pre-migration values:

Example:

```
nrcmd> server dhcp enable start-on-reboot
nrcmd> server dhcp start
nrcmd> server tftp enable start-on-reboot
nrcmd> server tftp start
```

Copy Cisco Prime Network Registrar Database Files to HA DNS Backup Server

Step 1 Use FTP or an equivalent network file copy mechanism to transfer the Cisco Prime Network Registrar database archive that was generated in the previous step to the parent directory of the Cisco Prime Network Registrar data directory (typically C:\NetworkRegistrar\Local\ on Windows, and /var/nwreg2/local/ on Linux/Solaris) on the HA DNS backup server.

Step 2 Ensure that the mechanism used to transfer the database archive preserves binary file data. If FTP sessions default to ASCII mode, change it to binary mode in order to produce a usable database on the HA DNS backup server.

Step 3 Stop the Cisco Prime Network Registrar product on the HA DNS backup server completely using the Windows Service Control manager (Windows) or nwreglocal script in /etc/init.d (Linux and Solaris). Ensure that the product is completely stopped, either by using the Windows Process Manager or the ps command line utility on Linux/Solaris, navigate to the parent directory of the Cisco Prime Network Registrar data directory (typically C:\NetworkRegistrar\Local\ on Windows, and /var/nwreg2/local/ on Linux/Solaris).

Step 4 Ensure to recursively remove all contents of the existing data directory, to prevent any conflicts with the database archive that is about to be extracted. Using tar or an equivalent utility, extract the contents of the database archive file: tar -xvf cnrdatadir.tar.

Reconfigure Cisco Prime Network Registrar on the HA DNS Backup Server

Step 1 Start the Cisco Prime Network Registrar servers on the HA DNS backup system using the Windows Service Control manager (Windows) or nwreglocal script in /etc/init.d (Linux and Solaris).

Step 2 Rectify the conflicts, if any, between HA DNS main system and any DHCP or TFTP server configuration settings.

Step 3 The DHCP integrity will be compromised if the DHCP server has a configuration similar to that of HA DNS main system. To know more on increasing the DHCP service availability, refer to the Cisco Prime Network Registrar product documentation. Cisco recommends that you completely remove any DHCP and/or TFTP related configuration on the HA DNS backup system using either the web UI or nrcmd CLI. You can restore the original DHCP and TFTP server-start-on-reboot attribute values, only after you confirm that the configuration values do not conflict with that of the HA DNS main system.

Example:

```
nrcmd> server dhcp enable start-on-reboot
nrcmd> server tftp enable start-on-reboot (only if it had be previously enabled)
```

Step 4 Edit the localhost Cluster object in the HA DNS backup server to reflect the values in use on the local server.

Configure Cisco Prime Network Registrar HA DNS on the HA DNS Main Server

Step 1 In HA DNS main server, define appropriate Cluster objects for both the HA DNS main and HA DNS backup servers.

Step 2 Create an HA Pair object by specifying appropriate Cluster names for the main and backup DNS server roles, and enable HA DNS for the HA Pair.

Step 3 Generate the report of changesets and exchange them between the two servers using the default report generation settings (Main-to-backup, Complete).

Step 4 Perform the changeset synchronization while the list of changesets is displayed.

Reload the DNS Servers

Step 1 Reload the DNS servers on both HA DNS systems to initiate the DNS RR synchronization process. Do it either through the Manage Servers page on the HA DNS main cluster when the HA DNS main server's DNS server has finished reloading, or to save a little time, initiate through separate connections to both clusters to perform the reloads in parallel instead of series.

Step 2 When the DNS servers are synchronizing, Cisco Prime Network Registrar does not allow DNS configuration updates (such as DDNS), but provides DNS queries and zone transfer. You can monitor the DNS server log files on the main and backup clusters to follow the progress of the DNS server synchronization process. The servers are fully operational when HA DNS enters Normal state.

Synchronizing HA DNS Zones

Local Advanced Web UI

To manually synchronize an HA DNS zone:

Step 1 From the **Design** menu, choose **Forward Zones** or **Reverse Zones** under the **Auth DNS** submenu to open the **List/Add Forward Zones** or List/Add Reverse Zones page.

Step 2 Click the **Commands** button for the zone which you want to synchronize on the Edit Zone page.

Step 3 Click the **Command** icon next to **Synchronize HA Zone** to synchronize the HA DNS zone.

Synchronizing the HA DNS zone will always sync the associated views and named ACLs for primary zones.

Note In the Expert mode, you have the option to choose the type of synchronization. The **Use Server Algorithms** option is checked by default. If you click the command icon next to the **Synchronize HA Zone** without choosing another option, server algorithms will be used to synchronize the zone. You can override this by checking either **Push Full Zone From Main to Backup** check box or **Pull Full Zone From Backup to Main** check box.

CLI Commands

Use **zone** *name* **ha-sync-all-rrs** to manually schedule HA zone synchronization for the zone, or to raise its priority, if the zone is already in the sync-pending state (see the **zone** command in the CLIGuide.html file in the /docs directory for syntax and attribute descriptions).

Enable Logging of HA DNS Information

The log settings, *ha-details* and *ha-messages* , enable logging of HA DNS-related information.

Note The HA communications with versions earlier to 8.0 are not supported. So, you have to upgrade both the main and the backup servers in the same maintenance window.

Viewing HA DNS Statistics

You can view HA DNS statistics.

Local Basic or Advanced Web UI

Click the **Statistics** tab on the Manage DNS Authoritative Server page to open the DNS Server Statistics page. The statistics appear under the HA Statistics and Max Counter Statistics subcategories of both the Total Statistics and Sample Statistics categories.

CLI Commands

Use **dns getStats ha** [**total**] to view the HA DNS Total counters statistics, and **dns getStats ha sample** to view the Sampled counters statistics.