



## Managing Zones

---

The Domain Name System (DNS) is a distributed database for objects in a computer network. By using a nameserver approach, the network consists of a hierarchy of autonomous domains and zones. The namespace is organized as a tree that often resembles the organizations that are responsible for the administration boundaries. For an introduction to the protocol, see [Chapter 13, “Introduction to the Domain Name System.”](#)

The basic function of DNS nameservers is to provide data about network objects by answering queries. You can configure the Cisco CNS Network Registrar DNS server and zones by accepting the system defaults or changing them.

This chapter describes the basics of configuring the Network Registrar DNS servers, and their primary and secondary zones. [Chapter 15, “Managing Resource Records,”](#) describes how to manage DNS resource records (RRs) and hosts, and [Chapter 16, “Managing DNS Server Properties,”](#) describes how to set some of the more advanced zone and DNS server properties.

## Staged and Synchronous Modes

You can perform additions or edits to DNS zones, RRs, and hosts in one of two modes—staged or synchronous:

- **Staged (or CCM)**—Changes to zones (and their hosts and protected server RRs) are written to the CCM database, but not immediately propagated to the DNS server until a synchronization is requested. This mode is reflected on the List/Add CCM Server Protected RRs for Zone page.
- **Synchronous (or DNS)**—After committing changes to CCM, hosts and protected RRs are immediately propagated to the DNS server. If propagation cannot occur because of an unreachable server, RRs are propagated at the next synchronization. This mode is reflected on the List/Add DNS Server RRs for Zone page (see [Figure 15-1](#) on [page 15-2](#)).

Synchronizations can occur on a zone basis or by creating a zone distribution. In synchronous mode, changes are written to the DNS server right away, even though a server reload is necessary for the zone to be published on the network. Staged or synchronous zone modes are defaulted for the session based on settings on the Main Menu page. The regional cluster defaults to staged, while the local cluster defaults to synchronous.

In the CLI, you can change the zone edit mode by setting the session *zone-edit-mode* attribute to staged or synchronous. For example:


```
nrcmd> session set zone-edit-mode=sync
```

# Creating and Applying Zone Templates

A zone template is a convenient way to create a boilerplate for primary zones that share many of the same attributes. You can apply a zone template to any zone, and override the zone's attributes with those of the template. You can create zone templates in the local and regional cluster Web UIs and in the CLI.

**Step 1** In the Web UI, click **DNS**, then **Zone Templates** (see [Figure 14-1](#) for the regional cluster version).

**Figure 14-1 List Zone Templates Page (Regional)**



**Step 2** You can add a zone template at the local and regional clusters, and you can also pull and push zone templates at the regional cluster in the Web UI:

- To add a zone template at the local cluster or explicitly add one at the regional cluster, click **Add Zone Template**. This opens the Add Zone Template page, which is almost identical to the Add Zone page for the local cluster.

To make the zone template meaningful, you would enter, in addition to its name, at least the suggested serial number, nameserver, contact e-mail address, and list of nameservers, because they are required for the zone itself. You might also want to specify any zone owners or zone distributions. You do not necessarily need to add these values for the zone template, because you can do so for the zone once it is created from the template. However, the template name and zone default TTL are required. (For a description of the minimally required zone attributes, see the “[Creating Primary Zones](#)” section on page 14-5.)

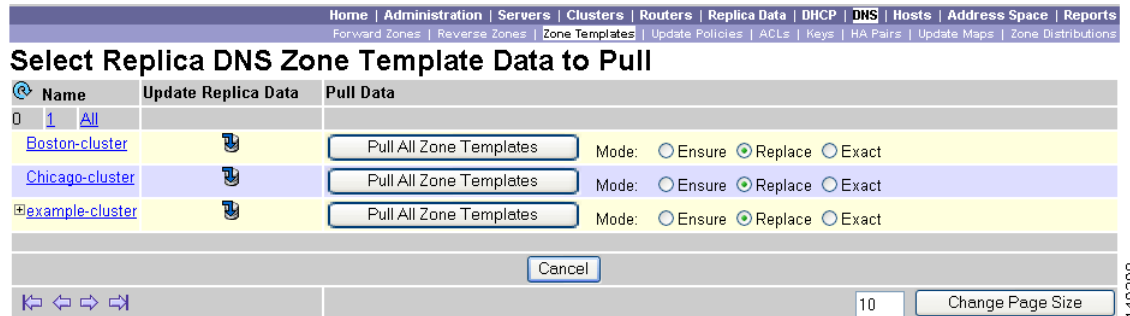
Once you are done entering these value, click **Add Zone Template** at the bottom of the page.

In the CLI, use `zone-template name create` to create the zone template. (See the “[Adding Primary Forward Zones](#)” section on page 14-5 for how to apply the template to a zone.) For example:


```
nrcmd> zone-template zone-template-1 create serial=1
```

- At the regional cluster, to pull a zone template from one or more local clusters, click **Pull Replica Zone Templates** on the List Zone Templates page. This opens the Select Replica DNS Zone Template Data to Pull page (see [Figure 14-2](#)).

Figure 14-2 Select Replica DNS Zone Template Data to Pull Page (Regional)



This page shows a tree view of the regional server's replica data for the local clusters' zone templates. The tree has two levels, one for the local clusters and one for the templates in each cluster. You can pull individual templates from the clusters, or you can pull all of their templates:

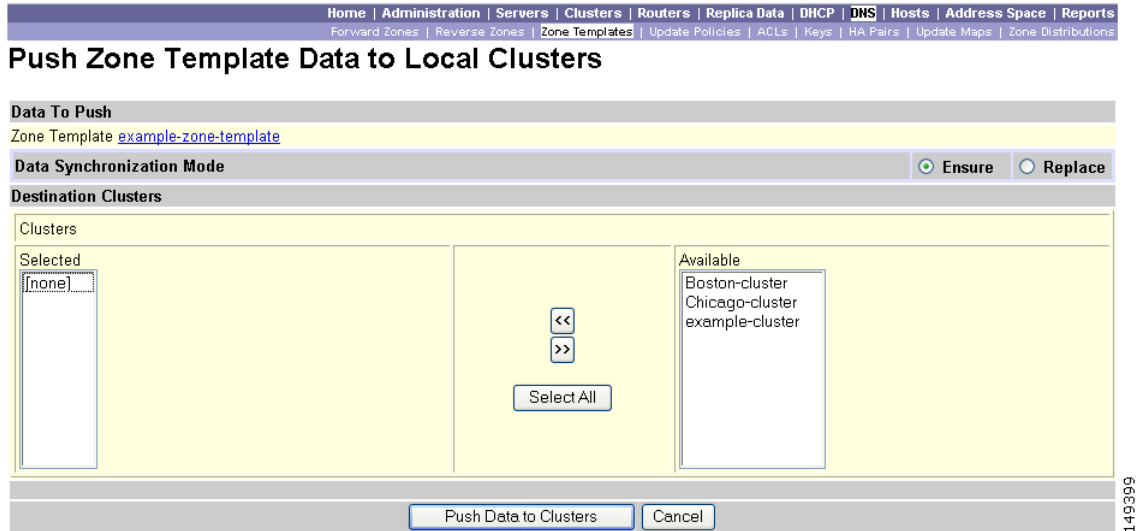
- To pull individual zone templates, expand the tree for the cluster, choose a pull criterion next to its name, then click **Pull Zone Template**.
- To pull all the templates from a cluster, choose a pull criterion, then click **Pull All Zone Templates from Cluster**.
- To update all the replica data for a cluster, click the Replica icon () next to its name.

The pull selection criteria are:

- Ensure—Pulls each template, except if an existing template by that name already exists at the regional cluster, in which case it does not overwrite the regional cluster data.
- Replace—Pulls each template and overwrites the data for it if it already exists at the regional cluster, without affecting any additional templates at the regional cluster. This is the default and recommended setting.
- Exact—Pulls each template, overwrites the data for it if it already exists at the regional cluster, and removes any additional templates at the regional cluster.
- At the regional cluster, to push a zone template to one or more local clusters:
  - To push all the zone templates on the page List Zone Templates page—Click **Push All Zone Templates**.
  - To push individual zone templates on the page List Zone Templates page—Click **Push Zone Template** next to the template name.

Both of these actions open a version of the Push Zone Template Data to Local Clusters page (see [Figure 14-3](#)).

Figure 14-3 Push Zone Template Data to Local Clusters Page (Regional)



This page provides a choice of the synchronization mode and the destination clusters. Move the desired cluster or clusters from the Available field to the Selected field, then click one of the data synchronization mode radio buttons:

- Ensure—Pushes each template, except if an existing template by that name already exists at the local cluster, in which case it does not overwrite the local cluster data. This is the default and recommended setting.
- Replace—Pushes each template and overwrites the data for it if it already exists at the local cluster, without affecting any additional templates at the local cluster.
- Exact—Available for “push all” operations only, it pushes each template, overwrites the data for it if it already exists at the local cluster, and removes any additional templates at the local cluster.

After making these choices, click **Push Data to Clusters**. This opens the View Push Zone Template Data Report page, where you can view the intended results of the push operation. Click **OK** to implement the push operation.

**Step 3** You can apply the template to a new or existing zone.

In the Web UI:

- a. New zone—Select the template from the Template drop-down list when you create the zone, as described in the “Adding Primary Forward Zones” section on page 14-5.
- b. Existing zone—After you create a zone (see the “Adding Primary Forward Zones” section on page 14-5), you can apply the template when you edit the zone on the Edit Zone page. Click the template name in the Template drop-down list, then click **Apply Template**.

In the CLI, use `zone-template name apply-to zone`. Note that the syntax permits one or more comma-separated zones and also the `all` keyword for all zones. You can also clone a template from an existing template by using `zone-template clone-name create clone=template`, and then make adjustments to the clone. For example:

```
nrcmd> zone-template zone-template-1 apply-to example.com,boston.example.com
nrcmd> zone-template cloned-template create clone=zone-template-1 owner=owner-1
```

**Caution**

Be careful in applying a template to an existing zone. The template overwrites all the explicitly set attributes for the zone (other than its name). This can result in severe consequences if the zone is already configured in a network. If you want to make a limited attribute change to multiple zones using a template, be sure to change only that attribute (or attributes), leaving the others unset, before you apply the template to the zones.

## Managing Primary DNS Servers

Adding a zone involves creating a domain name. You can also define an owner and use a zone template. If you do not use a template, you must also define the Start of Authority (SOA) and Name Server (NS) properties for the zone.

You do not need to create a loopback zone for the local host, because Network Registrar automatically creates one. A loopback zone is a reverse zone that a host uses to resolve its loopback address, 127.0.0.1, to localhost so that it can direct network traffic to itself. The loopback zone is 127.in-addr.arpa, which appears on the list of reverse zones. (If you inadvertently delete the loopback zone, see the Usage Guidelines for the **zone** command in the *Network Registrar CLI Reference Guide*.)

## Adding Primary Forward Zones

This section explains how to configure a primary nameserver with a primary forward zone. When you are done with this procedure, follow the procedure in the “[Adding Primary Reverse Zones](#)” section on [page 14-11](#) to configure a reverse zone for each network that you use.

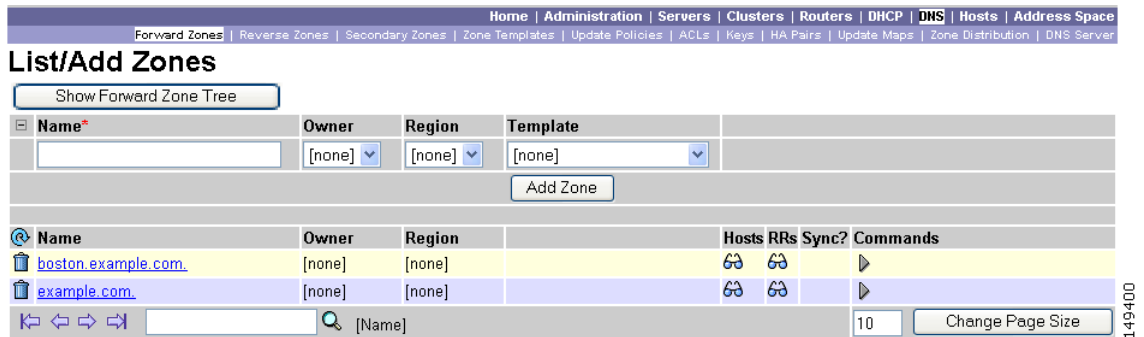
## Creating Primary Zones

The first thing to do when creating a forward zone is to give the zone a name and set its Start of Authority (SOA) RRs. The SOA record designates the top of the zone in the DNS inverted-tree namespace. A zone can have only one SOA record, which sets these primary zone properties:

- SOA time to live (TTL)—*soattl*
- Primary server name—*ns*
- Hostmaster (person in charge) name—*person*
- Serial number—*serial*
- Secondary refresh time—*refresh*
- Secondary retry time—*retry*
- Secondary expire time—*expire*
- Minimum TTL—*minttl*

**Step 1** In the Web UI, click **DNS**, then **Forward Zones** to open the List/Add Zones page (see [Figure 14-4](#)). (At the regional cluster, this is called the List Forward Zones page.)

Figure 14-4 List/Add Zones Page (Local)



**Step 2** Enter the zone name (in domain name format).

In the CLI, use **zone name create primary nameserver contact**. If you want to apply an existing template at this point, use the *template* attribute. For example:

```
nrcmd> zone example.com create primary ns1 hostmaster template=zone-template-1
```


**Note** In the previous CLI example that applies a template to the zone, even though you need to specify the nameserver and contact as part of the syntax, the template definition (if any) overwrites them.

**Step 3** Optionally choose a predefined owner or region.

**Step 4** To apply an existing zone template in the Web UI (see the “Creating and Applying Zone Templates” section on page 14-2), click the template name in the Template drop-down list.

In the CLI, if you want to apply a template after creating the zone, use **zone name applyTemplate template**; for example:

```
nrcmd> zone example.com applyTemplate zone-template-1
```

 **Caution** Be careful applying a template to a zone that is already live. Explicitly defined attributes on the template replace the existing ones defined for the zone.

**Step 5** Click **Add Zone** to open the Add Zone page (see Figure 4-8 on page 4-18).

**Step 6** For now, set only the serial number, primary server, and hostmaster data for the zone’s SOA record:

a. Set the *serial number*.

A primary DNS server uses a serial number to indicate when its database changes and uses any incrementing of this number to trigger a zone transfer to a secondary server. The serial number you can enter here is the *suggested* one only, and the DNS server does not always accept it. If you edit the serial number to be less than the actual serial number that the server maintains, the server logs a warning message and ignores the suggested serial number. You must reload the server for your change to take effect. The actual serial number always equals or is higher than the suggested one. You can get the actual serial number by using **zone name get serial** (if the DNS server is running; if the server is not running, or listing or showing the zone attributes, it always returns the suggested serial number), or by refreshing the DNS Server Value for the zone Serial Number attribute in the Web UI. In the Web UI, you must explicitly enter this suggested serial number when creating a zone. In the CLI, the serial number defaults to 1.

- b. Set the *primary DNS server* name.

Enter either just the hostname (such as *exampleDNSServ1*) or its fully qualified name (such as *exampleDNSServ1.example.com.*, ending with a trailing dot). Use the fully qualified name if the primary nameserver is in a different zone. The primary DNS server becomes the *ns* value in the zone's SOA record. In the Web UI, you must also specify one or more authoritative nameservers for the zone—these become the Name Server (NS) records for the zone. In the CLI, the primary DNS server automatically becomes the first NS record and also appears as the first entry in the *nameservers* attribute list.

- c. Set the *hostmaster* (person in charge) name and address as a slightly altered form of the e-mail address.

Substitute a dot (.) for the at symbol (@), and end the address with a trailing dot (for example, enter *hostmaster@example.com* as *hostmaster.example.com.*). Escape any dot before the @ in the original address with a backslash (\) (for example, enter *hostmaster.marketing@example.com* as *hostmaster\marketing.example.com.*).

- Step 7** In the Web UI, do not click **Add Zone** yet. You must still enter the authoritative nameservers for the zone, unless you specified this in the zone template you applied to the zone.
- 

## Adding Authoritative Nameservers for Zones

Authoritative nameservers validate the data in their zones. Both primary and secondary servers can be authoritative. The crucial difference is where they get their zone data. A primary server obtains its data from an administrator, as stored in the server's configuration database, and from DNS updates, typically from a DHCP server. A secondary server obtains the zone data from its designated master servers by way of a zone transfer.

You must add at least one nameserver for a zone—Network Registrar does not consider the zone data complete unless you do so. The nameservers you list should be those that you want people outside your domain to query when trying to resolve names in your zone. In the CLI, creating a primary zone requires specifying a primary DNS server—this server becomes the first entry in the nameserver list. In the Web UI, you must add the authoritative nameservers in addition to the primary server for the zone. If the primary DNS server for the zone is in the zone, you must create a host address for it—see the [“Adding Host Addresses for Nameservers”](#) section on page 14-8.

- Step 1** In the Web UI, on the Add Zone page, enter the name of an authoritative nameserver (either as hostname or fully qualified, if in another zone) in the field next to the Add Nameserver button.

In the CLI, enter a comma-separated list of fully qualified domain names using **zone name set nameservers=list**. Note that only the first server entered is confirmed by the command. Use **zone name show** to show all the server names. Then reload the server.

- Step 2** In the Web UI, click **Add Nameserver**.

- Step 3** Repeat the previous steps for each additional nameserver you add.

- Step 4** In the Web UI, unless you want to add additional attributes for the zone, you can now click **Add Zone**.
-

## Adding Host Addresses for Nameservers

For every DNS internal-to-zone nameserver, you must create an Address (A) RR to associate the server's domain name with an IP address.

- 
- Step 1** Create the zone, as described in the [“Adding Primary Forward Zones” section on page 14-5](#).
  - Step 2** In the Web UI, click **Host** to open the List Zones page.
  - Step 3** Click the zone name to open the List/Add Hosts for Zone page.
  - Step 4** Enter the hostname.
  - Step 5** Enter the IP address of the authoritative server.
  - Step 6** Click **Add Host**. The server's hostname and address appear in the list.

In the CLI, use `zone name addRR hostname A address` to add the authoritative server's hostname and address. To list the host, use `zone name listHosts`. To remove the host, use `zone name removeRR hostname A`.

---


## Creating Zone Templates from Zones

You can save zone information as a template so that you can re-use the data for other zones. You do this from the Edit Zone page in the local cluster Web UI.

- 
- Step 1** After you modify the zone information, click **Modify Zone and Save Template**.
  - Step 2** On the Save New Zone Template page, give the template a name in the Value field.
  - Step 3** Click **Save Zone Template**. You return to the List/Add Zones page.
- 

## Confirming Zone Nameservers

Confirm your zone's NS RR configuration by looking at the RRs that you created.

In the Web UI, click the View icon () in the Configuration RRs column of the zone name on the List/Add Zones page to open the List/Add CCM Protected Server RRs for Zone page or List/Add DNS Server RRs for Zone page. There should be an A record for each nameserver host in the zone. Edit these records or add more on this page.

In the CLI, use `zone name listRR` to check the RRs you added. To activate the records, you must reload the DNS server.

## Importing and Exporting Zone Data

The easiest and quickest way to create a primary zone is to import an existing BIND format zone file, defined in RFC 1035. You can also export these same kinds of files to another server. BIND 4.x.x uses a boot file, called named.boot, to point the server to its database files. You can import your entire BIND 4.x.x configuration using the **import** command. BIND 8 and BIND 9 use a configuration file, called named.config, with a different syntax.

You can import and export zone data only by using the CLI.

When a BIND file contains an `$INCLUDE` directive, BIND searches for the include file relative to the directory that the `directory` directive in the `named.boot` file specifies. In contrast, the `nrcmd` program searches for the include file relative to the directory containing the zone file being processed.

To avoid this problem, ensure that the BIND configuration uses absolute paths whenever specifying an include file in a zone file. If your zone files contain relative paths when specifying include files, and the directory containing the zone file is not the same as the directory that the `directory` directive in the `named.boot` file specifies, your configuration cannot load properly. You need to convert the relative paths in your zone files to absolute paths so that you can import your BIND configuration into Network Registrar. Here is an example of a configuration and how to fix paths in directory hierarchy, configuration files, and zone files:

- Directory hierarchy:

```
/etc/named.conf
/etc/named.boot
/usr/local/domain/primary/db.example
/usr/local/domain/primary/db.include
/usr/local/domain/secondary
```

- Configuration file (`/etc/named.conf`):

```
#BIND searches for zone files and include files relative to /usr/local/domain
option directory /usr/local/domain
#BIND finds zone file in /usr/local/domain/primary
zone example.com {
    type master ;
    file primary/db.example ;
#end of /etc/named.conf
```

- Configuration file (`/etc/named.boot`):

```
#BIND searches for zone files and include files relative to /usr/local/domain
directory /usr/local/domain
#BIND finds zone file in /usr/local/domain/primary
primary example.com primary/db.example
#end of /etc/named.boot
```

- Incorrect zone file (`/usr/local/domain/primary/db.example`):

```
#BIND searches for include file relative to /usr/local/domain
$INCLUDE primary/db.include
#end of /usr/local/domain/primary/db.example
```

To make the configuration loadable, change the relative path (`$INCLUDE primary/db.include`) in the file `db.example` to an absolute path (`$INCLUDE /usr/local/domain/primary/db.include`).

Table 14-1 describes the `named.boot` and `named.conf` file directives that BIND 4 and BIND 9 support, and the corresponding Network Registrar user interface location or syntax, if any.

**Table 14-1 BIND-to-CLI Command Mappings**

BIND 4 Command	BIND 9 Command	Mapping to User Interface
—	<code>acl name { addr-match-list };</code>	Web UI: List/Add Access Control Lists page fields (see the “Access Control Lists” section on page 27-5). CLI: <code>acl name create value match-list=addr-match-list</code>
<code>forwarders addrlist</code>	<code>options { forwarders { addr; addr;... } };</code>	Web UI: Edit DNS Server page, set Forwarders: IP Address field. CLI: <code>dns addForwarder addr[,addr...]</code>

Table 14-1 BIND-to-CLI Command Mappings (continued)

BIND 4 Command	BIND 9 Command	Mapping to User Interface
—	key <i>id</i> { algorithm <i>string</i> ; secret <i>string</i> ; }	Web UI: List/Add Encryption Keys page fields. CLI: <b>key name create secret algorithm=alg</b>
limit transfers-in <i>num</i>	options { transfers-in <i>num</i> ;};	Web UI: Edit DNS Server page, set <i>xfer-client-concurrent-limit</i> . CLI: <b>session set visibility=3</b> <b>dns set xfer-client-concurrent-limit=number</b>
—	options { allow-query <i>addr-match-list</i> ;};	Web UI: Edit DNS Server page, enable <i>restrict-query-acl</i> CLI: <b>dns set restrict-query-acl</b>
options allow-recursion <i>addr-match-list</i>	options { allow-recursion <i>addr-match-list</i> ;};	Web UI: Edit DNS Server page, enable <i>restrict-recursion-acl</i> . CLI: <b>dns set restrict-recursion-acl</b>
options forward-only	options { forward only ;};	Web UI: Edit DNS Server page, enable <i>Slave mode</i> . CLI: <b>dns enable slave-mode</b>
options listen-on <i>port</i>	options { listen-on <i>port</i> { <i>addr-match-list</i> } ;};	Web UI: Edit DNS Server page, set <i>Listening port</i> . CLI: <b>dns set local-port-number=port</b>
options max-cache-ttl <i>num</i>	options { max-cache-ttl <i>num</i> };	Web UI: Edit DNS Server, set <i>Max. RR caching TTL</i> . CLI: <b>dns set max-cache-ttl=num</b>
options no-fetch-glue	options { fetch-glue no ;};	Web UI: Edit DNS Server page, enable <i>Don't fetch missing glue records</i> . CLI: <b>dns enable no-fetch-glue</b>
options no-recursion	options { recursion no ;};	Web UI: Edit DNS Server page, enable <i>Recursive queries</i> . CLI: <b>dns enable no-recurse</b>
options notify yes	options { notify yes ;};	Web UI: Edit DNS Server page, enable <i>Send zone change notification (NOTIFY)</i> . CLI: <b>dns enable notify</b>
options rrset-order <i>order order ...</i>	options { rrset-order <i>order</i> ; <i>order</i> ; ... ;};	Web UI: Edit DNS Server page, enable <i>Enable round-robin</i> . CLI: <b>dns enable round-robin</b>
options support-ixfr yes	options { request-ixfr yes ;};	Web UI: Edit DNS Server page, enable <i>Request incremental transfers (IXFR)</i> . CLI: <b>dns enable ixfr-enable</b>
options transfer-format many-answers	options { transfer-format many-answers ;};	Web UI: Edit DNS Server page, enable <i>Use multirec format for zone transfers</i> . CLI: <b>dns enable axfr-multirec-default</b>
primary <i>zonename file</i>	zone " <i>name</i> " { type master; };	Web UI: Add Zone page fields. CLI: <b>zone name create primary file=file</b>
secondary <i>zonename</i> <i>addr list [backupfile]</i>	zone " <i>name</i> " { type slave; };	Web UI: Add Secondary Zone page fields. CLI: <b>zone name create secondary ip-addr [ip-addr...]</b>

**Table 14-1 BIND-to-CLI Command Mappings (continued)**

BIND 4 Command	BIND 9 Command	Mapping to User Interface
slave	zone "name" { type slave; };	Web UI: Edit DNS Server page, enable <i>Slave mode</i> . CLI: <b>dns enable slave-mode</b>
—	zone "name" { allow-query { addr; ... } };	Web UI: Edit Zone page, set <i>restrict-query-acl</i> . CLI: <b>zone name set restrict-query-acl=addr[,addr...]</b>
tcp <sub>list</sub> <i>addrlist</i> xfer <sub>nets</sub> <i>addrlist</i>	zone "name" { allow-transfer { addr; ... } };	Web UI: Edit Zone page, enable <i>restrict-xfer</i> and set <i>restrict-xfer-acl</i> . CLI: <b>zone name enable restrict-xfer</b> <b>zone name set restrict-xfer-acl=addr[,addr...]</b>

For details on the following topics, see the Usage Guidelines for the **zone** command in the *Network Registrar CLI Reference Guide*:

- Importing zone data
- Exporting zone data
- Exporting Unix hosts files

## Adding Primary Reverse Zones

For a correct DNS configuration, you must create a reverse zone for each network that you use. A reverse zone is a primary zone that DNS clients use to convert IP addresses back to hostnames, and resides in a special in-addr.arpa domain. You can create a reverse zone manually or import it from BIND. You can also create reverse zones from subnets (see the “[Adding Reverse Zones from Subnets](#)” section on [page 14-11](#)).

### Adding Reverse Zones as Zones

To add a reverse zone manually as a zone:

- 
- Step 1** In the local cluster Web UI, click **Zone**, then **Reverse Zones** to open the List/Add Reverse Zones page. This page is almost identical to the List/Add Zones page (see [Figure 14-4 on page 14-6](#)).
- Step 2** In the Web UI, add a reverse zone the same way you would add a forward zone, as described in the “[Adding Primary Forward Zones](#)” section on [page 14-5](#), except use the reverse of the forward zone’s network number added to the special in-addr.arpa domain as the zone name. Use the same template or SOA and nameserver values as you used for the related forward zone.

In the CLI, use **zone name create primary** and **zone name addRR PTR** to add the primary reverse zone and pointer records for the server. You can also apply a zone template.

---

### Adding Reverse Zones from Subnets

An alternative to creating reverse zones manually is to create them from existing subnets. You can do this in the Web UI only.

**Step 1** In the Web UI, click **Address Space**, then **Subnets** to open the List/Add Subnets page (see [Figure 14-5](#)).

**Figure 14-5** List/Add Subnets Page (Local)

- Step 2** Create a subnet for the reverse zone, or use one of the existing subnets. If the subnet already has a reverse zone created from it, the Reverse Zone column shows the View icon (🔍), which opens the List/Add Reverse Zones page. If the subnet has the Create icon (🔧) in the Reverse Zone column, click the icon to open the Create Reverse Zone(s) for Subnet page.
- Step 3** On the Create Reverse Zone(s) for Subnet page, you must choose an existing zone template.
- Step 4** Click **Report** to show the change sets for the creation.
- Step 5** Click **Run** to run the creation.
- Step 6** Click **Return** to return to the List/Add Subnets page.
- Step 7** Confirm the creation by clicking **DNS**, then **Reverse Zones** to see the newly created zone on the List/Add Reverse Zones page.

## Managing Secondary Servers

When you configure a zone, choose at least one secondary server. If you have only one nameserver and it becomes unavailable, there is nothing that can look up names. A secondary server splits the load with the primary or handles the whole load if the primary is unavailable. When a secondary server starts up, it contacts the primary and pulls the zone data over. This is known as a zone transfer.



### Tip

If the authoritative server for your secondary zones is also running Network Registrar 6.0 or later, see the [“Managing Zone Distributions”](#) section on page 14-17 for how to avoid entering these zones manually. If you have only one secondary server, remove it geographically from the primary. They should not even be on the same network segment, switch, or router, but on a different cluster entirely.

You can configure a secondary DNS server to be responsible for a secondary zone, which makes the server a secondary for that zone. You also need to give the address of the master server from which to perform zone transfers. Network Registrar must know about this master server.

## Adding Secondary Forward Zones

To add a secondary forward zone in the local cluster Web UI:

- Step 1** Click **DNS**, then **Secondary Zones** to open the List Secondary Zones page.
- Step 2** Click **Add Secondary Zone** to open the Add Secondary Zone page (see [Figure 14-6](#)).

**Figure 14-6 Add Secondary Zone Page (Local)**

Attribute	Value	Data Type	Default
Name* (origin)	<input type="text"/>	DNS name	
master-servers*	<input type="text"/>	list	
restrict-xfer	<input type="radio"/> true <input type="radio"/> false	boolean	false
restrict-xfer-acl	<input type="text"/>	address match list	
IXFR and NOTIFY Settings			
Subzone Forwarding Settings			
Checkpoint Settings			
Query Settings			

149402

In the CLI, use `zone name create secondary`. The IP address you include is that of the nameserver from which data is expected, typically a primary nameserver. You cannot apply a template to a secondary zone.

## Adding Secondary Reverse Zones

You should add a secondary reverse zone, just as you added a secondary forward zone.

- Step 1** Add the secondary reverse zone the same way you did the primary reverse zone, except set the zone type to Secondary. See the [“Adding Primary Reverse Zones”](#) section on page 14-11.
- Step 2** Make the secondary zone’s domain name an in-addr.arpa reverse domain.
- Step 3** Add the nameserver address for the secondary forward zone and set any zone transfer address restrictions, as in the [“Adding Secondary Forward Zones”](#) section on page 14-13.
- Step 4** Reload the DNS server and confirm its status.



## Enabling Zone Transfers

A secondary server periodically contacts its master server for changes, called a zone transfer. The interval is defined in the server’s SOA record as the secondary refresh time. You can restrict zone transfers by setting the `restrict-xfer` attribute to true (the default is false) on the master server.

**Note**

If you restrict zone transfers, the **nslookup** utility **ls** command may fail because it tries to do a full zone transfer, unless you include the IP address that **ls** runs from in the zone's *restrict-xfer-acl* list.

In the local cluster Web UI:

- 
- Step 1** On the List/Add Zones page, click the name of the primary zone to open the Edit Zone page.
- Step 2** In the zone attributes, you can set the *restrict-xfer* attribute to false (the default). If you set the attribute to *true*, you can also specify a list of servers to which to restrict the zone transfers by using the *restrict-xfer-acl* attribute, separating the IP addresses with commas.
- Step 3** Click **Modify Zone**.
- Step 4** You can force zone transfers for the DNS server in two ways:
- On the List Secondary Zones page, click the Run icon () in the Force Zone Transfer column.
  - To force all zone transfers from the primary server, on the DNS Server Commands page (see [Figure 6-3 on page 6-2](#)), click the Run icon () next to Force all zone transfers.
- 

Secondary zones can also restrict zone transfers from other secondary zones, so that the *restrict-xfer* and *restrict-xfer-acl* attributes are also available for secondary zone configurations.

In the CLI, zone transfers are enabled by default, unless you restrict them using **zone name enable restrict-xfer**. If you want to force a zone transfer, use **zone name forceXfer secondary**.

## Adding Subzones

As the zone grows, you might want to divide it into smaller pieces called subzones. You can delegate administrative authority for these subzones, and have them managed there or served by separate servers. This partitioning is called subzone delegation. Establish subzone delegation by performing these tasks:

1. Choose a subzone name
2. Specify a nameserver name
3. Specify a nameserver address

## Choosing Subzone Names and Servers

After you decide to divide the zone into subzones, you must create names for them. Involve the people responsible for the subzones in deciding their names, and try to maintain a consistent naming scheme.

These suggestions can help you avoid subzone naming problems:

- Consider not naming a subzone by its organizational name. In a changing business environment, organizations merge and are renamed. Naming a subzone after an organization could result in a name that is no longer meaningful over time.
- Consider not using geographical names that indicate the subzone location. Geographical names are meaningless to people outside your organization.
- Do not use cryptic names; make them obvious.

- Do not use existing or reserved top-level domain names as subzones. Using existing names can result in routing problems.

After you choose a subzone name, specify its nameservers, the ones the parent domain's nameservers use when queried about the subzone. To ensure that the subzone is always reachable, you should specify two nameservers. They must be authoritative for this zone as either primary or secondary, or lame delegation will result (see the “[Reporting Lame Delegation](#)” section on page 16-11).

Whenever a subzone's nameserver changes its name or address, the subzone administrator must inform its parent zone so that the latter's administrator can change the subzone's nameserver and *glue records*. A glue record is an A record with the address of a subzone's authoritative nameserver. If the subzone's administrator fails to inform its parent, the glue records are invalid. The common symptom is that a host cannot reach a host in another domain by its name, only by its address.

**Note**

Network Registrar detects lame delegation by reporting missing subzone NS records in the parent zone, if NS record addresses do not match, and if glue A records are required.

## Creating and Delegating Subzones

You delegate a subzone by creating it in the parent zone. There should be one NS record for each nameserver to which the subzone is delegated. Each NS record requires a corresponding A record describing the address of the nameserver, unless the nameserver is outside the parent zone or subzone. This A record is called a *glue* record.

- Step 1** In the Web UI, create a zone as a subdomain of the parent domain on the List/Add Zones (or List Forward Zones) page:
- If you apply a zone template, you go right to Step 2.
  - If not applying a zone template, on the Add Zone page, add the SOA records and the nameserver with its address, then click **Add Zone**.

In the CLI, on the subzone's primary nameserver machine, create the subdomain:

```
nrcmd> zone boston.example.com. create primary bostonDNSserv1 hostmaster
```


- Step 2** In the Web UI, if Network Registrar detects a parent zone based on the subzone name, the Create Subzone in Parent Zone page appears. Click **Create as Subzone** (or **Create as Unparented Zone** if you do not want it to be a subzone) on this page.

- Step 3** If you configured a nameserver in the subzone, you need to create a glue Address (A) record for it. In the field provided, enter the IP address of the nameserver, then click **Specify Glue Records**. (If there are multiple subzone nameservers, there are multiple fields for the glue records.)

In the CLI, create an A record for the subzone's nameserver:

```
nrcmd> zone boston.example.com. addRR bostonDNSserv1 A 192.168.40.1
```

- Step 4** In the Web UI:
- a. Click **Report** to show the intended change sets for the added records, then click **Run**.
  - b. Click **Return** after viewing the actual change sets implemented.
  - c. To confirm the added records for the subzone, click the View icon (🔍) in the RRs column for the subzone. The glue A record or records for the subzone nameserver should appear. Click **Return to Zone List**.

- d. To confirm the added records for the parent zone, click the View icon () in the RRs column for the parent zone. The subzone nameserver (NS) record or records plus the glue A record or records for them should appear. Click **Return to Zone List**.

In the CLI, on the parent zone's nameserver machine:

- a. Add an NS record for the subzone's nameserver:

```
nrcmd> zone example.com. addRR boston NS bostonDNSserv1.boston.example.com.
```

- b. Create a glue A record for the subzone's nameserver:

```
nrcmd> zone example.com. addRR bostonDNSserv1.boston.example.com. A 192.168.40.1
```

**Step 5** Reload the DNS server.

---

## Undelegating Subzones

If you undelegate a subzone, you need to remove any associated NS and glue A records from the parent zone.



### Note

If you delete the subzone, Network Registrar cleans up the delegation records automatically.

---


In the Web UI, on the List/Add CCM Server Protected RRs for Zone page or List/Add DNS Server RRs for Zone page, delete the NS record for the subzone, then delete the glue A record for the subzone's server host.

In the CLI, use **zone name removeRR NS** and **zone name removeRR A** to remove the subzone's NS and glue A records.

## Editing Subzone Delegation

You can edit the subzone's RRs:

---

- Step 1** In the Web UI, on the List/Add CCM Protected Server RRs for Zone page or List/Add DNS Server RRs for Zone page, edit the NS RR for the subzone by clicking the Edit icon () next to the record to open the Edit RR in Zone page.
  - Step 2** Edit the NS record data.  
In the CLI, use **zone name removeRR** to delete the NS and glue A records, then use **zone name addRR** to replace them. Reload the DNS server.
  - Step 3** In the Web UI, click **Modify Resource Record**.
  - Step 4** Edit the glue A RR for the subzone's server in the same way as in the previous steps.
  - Step 5** Reload the DNS server.
-

## Enabling DNS Updates

DNS Update (RFC 2136) integrates DNS and DHCP so that they can work together. DNS update automatically records the association between the hosts and their DHCP-assigned addresses. Using DHCP and DNS update, you can configure a host automatically for network access whenever it attaches to the network. You can locate and access the host using its unique DNS hostname.

DNS update is described more fully in [Chapter 27, “Configuring DNS Update.”](#) The chapter includes sections on the following:

- Update policy (the **Update Policies** tab)—Determines what kind of RRs you want updated when name-to-address associations change through DHCP.
- Update map (the **Update Maps** tab)—Defines an update relationship between a DNS server or HA DNS pair and a DHCP failover pair, DHCP policies, client-class, or access control list. (See the [“Configuring DNS Update Policies”](#) section on page 27-8.)

## Managing Zone Distributions

Creating a zone distribution simplifies creating multiple zones that share the same secondary zone attributes. The zone distribution requires adding one or more predefined secondary servers. Running a zone distribution synchronization adds secondary zones managed by secondary (slave) servers for each primary zone managed by a primary server. You can also use zone distributions to synchronize zone data from the CCM database to the local DNS server and regional and local cluster zone data.

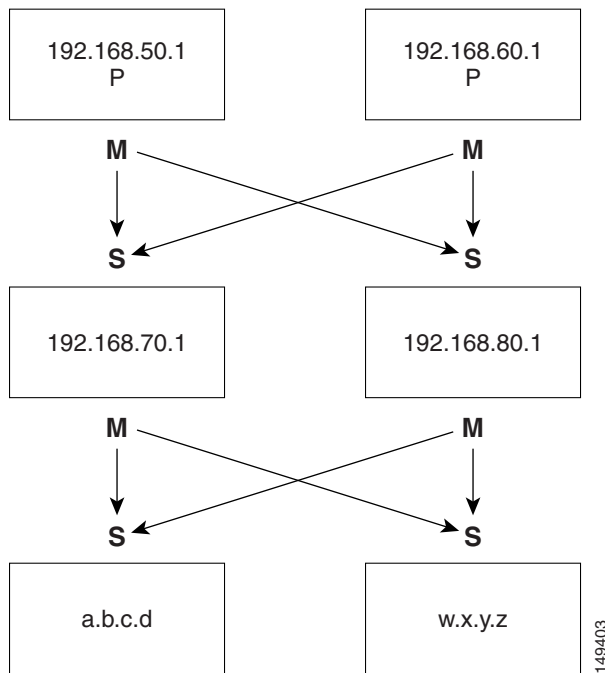
The distribution must be in a star topology, that is, one primary server and multiple secondary servers. The authoritative (master) server can only be the local primary server where the zone distribution default is defined. Starting with Network Registrar 6.2, you can manage one zone distribution at the local cluster and multiple distributions at the regional clusters.

## Preparing the Zone Distribution Map

To prepare for creating a zone distribution, draw a zone distribution map diagram on paper.

- 
- Step 1** Start by identifying the HA DNS pair that is primary (or the primary server if HA is not involved) for all the zones that you include in the map:
- a. Create a box for each server in the HA DNS pair. For example, the server pair for the Chicago-cluster consists of the servers at 192.168.50.1 and 192.168.60.1.
  - b. Write the IP addresses of each server in each box.
  - c. Write a **P** (for Primary) inside each box (see [Figure 14-7](#)).

Figure 14-7 Diagramming a Zone Distribution Map



- Step 2** Identify the role as master for each server by writing an **M** below the box. In the example, both primary servers are, by definition, also masters that will send copies of their zones to other servers over zone transfers. Even so, write the **M** below the box to make later steps easier.
- Step 3** Identify all slave servers that will receive zone transfers directly from these masters. Below the master server boxes on the page, add a box for each slave, and write its IP address inside the box. For example, the slave servers at 192.168.70.1 and 192.168.80.1 get zone transfers from the Chicago-cluster masters.
- Step 4** Write an **S** above each slave server box.
- Step 5** Draw arrows from the **M** to each **S** representing the zone transfer flow (see the diagram). In this HA DNS example, the arrows go from each master to both slaves.
- Step 6** As you can see from the diagram, you can extend the boxes further so that the original slaves can become masters to another set of servers (a.b.c.d and w.x.y.z).
- Step 7** In the Web UI, enter the IP address in each box with an **M** below it in the Master Servers list when creating the zone distribution (see Figure 14-9 on page 14-19). In the CLI, set the master-servers attribute to the list of IP addresses; for example:

```
nrcmd> zone-dist dist-1 create Chicago-cluster master-servers=192.168.50.1,192.168.60.1
```

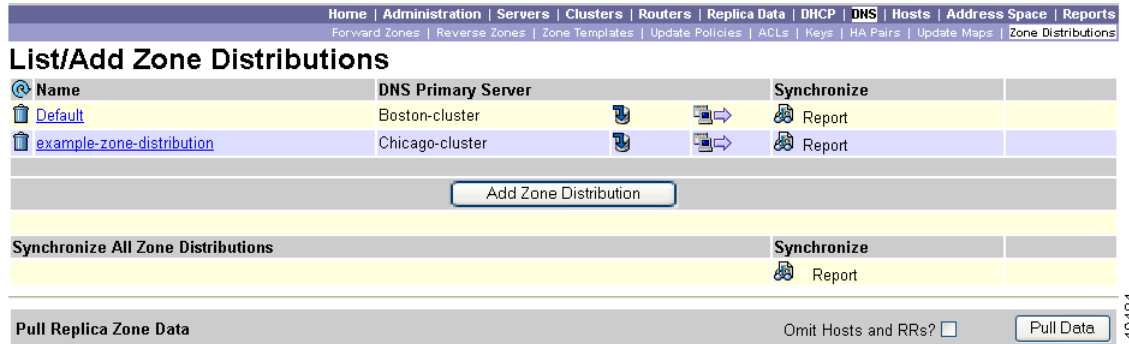
- Step 8** In the Web UI, from the Secondary Servers drop-down list on the Add or Edit Zone Distribution Secondary Server page, choose the cluster associated with the slave server IP addresses in the boxes that have an **S** above them. In the CLI, use **zone-dist name addSecondary cluster**; for example:

```
nrcmd> zone-dist dist-1 addSecondary Boston-cluster
```

# Creating a Zone Distribution

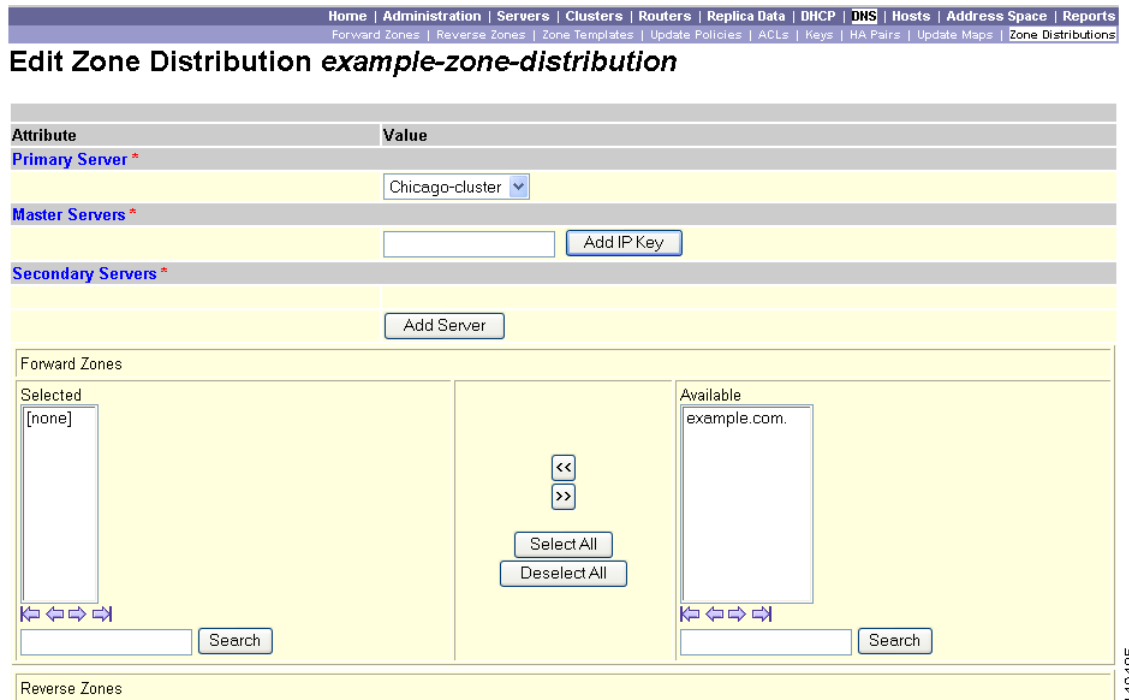
- Step 1** In the Web UI, click **DNS**, then **Zone Distributions** (for the regional cluster) or **Zone Distribution** (for the local cluster). This opens the List/Add Zone Distributions page (see [Figure 14-8](#)) or the List Zone Distributions page at the local cluster. Note that the Default zone distribution is predefined at both clusters; however, the Default cluster is the only one available at the local cluster.

**Figure 14-8 List/Add Zone Distributions Page (Regional)**



- Step 2** At the regional cluster, to add a new zone distribution, click **Add Zone Distribution** to open the Add Zone Distribution page. To edit an existing zone distribution, click its name to open the Edit Zone Distribution page (see [Figure 14-9](#) for a partial view). The Add Zone Distribution page and Edit Zone Distribution page are functionally equivalent.

**Figure 14-9 Edit Zone Distribution Page (Regional)**



In the regional CLI, use **zone-dist name create primary-cluster**. (The primary cluster can also be the HA DNS pair.) For example:

```
nrcmd> zone-dist dist-2 create Chicago-cluster
```

**Step 3** In the Primary Server field, enter the cluster (or configured HA DNS pair) that has the primary server. (In the CLI, this already occurred in the previous step.) This primary server is authoritative for the zones that you will determine further down the page. This selection is subtractive: the next zone distribution you create will no longer have the cluster that you set here as one of the choices.

**Step 4** In the Master Servers list, add the IP address (and optional key) for each master server. The master server is generally the primary server. However, you might want to set up a hierarchy of primaries and secondaries where you need to define the master servers for each of the secondary relationships. You might also want to determine the HA DNS server pairs from the master server list. You can also add an optional TSIG key (see the “[Transaction Security](#)” section on page 27-6) to the master server address by hyphenating the entry in the format *address-key*. For each entry, click **Add IP Key**

In the CLI, use **zone-dist name set master-servers=addresses**, separating the addresses with commas. For example:

```
nrcmd> zone-dist zone-dist-2 set master-servers=192.168.50.1,192.168.60.1
```

**Step 5** For a zone distribution, you need to add at least one secondary server. Click **Add Server** on the Edit Zone Distribution page to open the Add Zone Distribution Secondary Server page. Here, choose the cluster of the secondary server. Optionally, if the master servers are other than the primary servers indicated for the zone distribution, add the master server addresses, separating multiple addresses with commas. After clicking **Add Secondary Server** returns you to the Add or Edit page, you can connect to the secondary server cluster, delete it, or edit it to change the master servers.

In the CLI, use **zone-dist name addSecondary secondary-cluster**. For example:

```
nrcmd> zone-dist zone-dist-2 AddSecondary Boston-cluster
```

**Step 6** In the Web UI, choose the forward and reverse zones for the zone distribution. The Default zone distribution includes all the created forward and reverse zones. For all other created zone distributions, you must move the zone or zones into the Selected column.

In the CLI, you must associate the zone distribution directly with the zone or zone template. Use **zone name set dist-map=zone-dist-list** or **zone-template name set dist-map=zone-dist-list**, separating the zone distribution entries with commas. For example:

```
nrcmd> zone example.com set dist-map=zone-dist-2
nrcmd> zone-template zone-template-1 set dist-map=zone-dist-2
```


**Step 7** In the Web UI, click **Add Zone Distribution** or **Modify Zone Distribution**.

**Step 8** Synchronize the zone distribution with the local cluster DNS servers. A synchronization:

- Pushes staged zone, RR, or host edits to the primary server cluster or HA DNS pair for the regional cluster in Ensure, Replace, or Exact modes, or from the local cluster in Exact mode.
- Creates secondary zones for secondary servers, in Exact mode.

**Step 9** Choose a synchronization mode:

- Update—Adds new zones, RR sets, and hosts; replaces existing hosts if there are conflicts; and creates new secondary zones.
- Complete—Like Ensure mode, except that it always replaces existing RR sets and hosts, and modifies the master server list on existing secondary zones.
- Exact—Like Complete mode, except that it deletes extra zones, RR sets, hosts, and secondary zones no longer on the primary.

**Step 10** In the Web UI, click the Report icon () in the Synchronize column (or the same icon in the Synchronize All Zone Distributions area of the page at the regional cluster). This opens the Sync Zone Distribution page that shows a preview of the data synchronized. Click **Run**.

In the CLI, use **zone-dist name sync**. You can do a synchronization in update, complete, or exact mode, and you can exclude RRs and secondary zones:

- At the local cluster, this synchronizes staged edits to the DNS server and primary zones to secondaries. Regardless of the synchronization mode, this always synchronizes the exact list of authoritative zones.
- At the regional cluster, this synchronizes primary zones with the local clusters, and primaries to secondaries. This replaces primary zones at the local cluster in Update and Complete modes, and deletes extra primary zones at the local cluster in Exact mode.
- For secondary zones, the same synchronization logic occurs at the local and regional clusters. In Update mode, this ensures that corresponding secondary zones exist on the server. In Complete mode, existing zones are updated to use the master server list specified by the zone distribution map. In Exact mode, any zones not matching the distribution map are deleted. For example:

```
nrcmd> zone-dist zone-dist-1 sync exact no-rrs no-secondaries
```

**Note**

If you move a zone from one zone distribution to another, synchronize the first zone distribution, move the zone, then synchronize the second zone distribution.

## Pulling Zone Distributions from Replica Data

To pull zone distributions from the replica data of the local clusters instead of explicitly creating them:

- Step 1** On the List/Add Zone Distributions page, click **Pull Replica Zone Data**.
- Step 2** Choose the data synchronization mode (**Update**, **Complete**, or **Exact**) on the Select Pull Replica Zone Data page. These modes are described in the table on that page.
- Step 3** Click **Report** at the bottom of the page.
- Step 4** Click **Run** on the Report Pull Replica Zone Data page.
- Step 5** Click **OK** on the Run Pull Replica Zone Data page.

