



Cisco Prime Network Registrar 11.0 Caching and Authoritative DNS User Guide

First Published: 2021-04-23

Last Modified: 2021-10-20

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021 Cisco Systems, Inc. All rights reserved.



CONTENTS

PART I

Introduction 15

CHAPTER 1

Introduction to the Domain Name System 1

How DNS Works 1

Overview of Concepts in DNS 2

Domains 2

Nameservers 4

Reverse Nameservers 5

Authoritative and Caching DNS Servers 6

High-Availability DNS 6

EDNS 6

DNS Views 7

CHAPTER 2

DNS Server Status Dashboard 9

Opening the Dashboard 9

Display Types 9

General Status Indicators 10

Graphic Indicators for Levels of Alert 10

Magnifying and Converting Charts 11

Legends 11

Tables 11

Line Charts 11

Area Charts 13

Other Chart Types 13

Getting Help for the Dashboard Elements 14

Customizing the Display 14

- Refreshing Displays 14
- Setting the Polling Interval 15
- Displaying Charts as Tables 15
- Exporting to CSV Format 15
- Selecting Dashboard Elements to Include 15
- Configuring Server Chart Types 16

PART II

Caching DNS Server 19

CHAPTER 3

Managing Caching DNS Server 21

- Setting DNS Caching Server Properties 21
 - Setting General Caching DNS Server Properties 22
 - Specifying Log Settings 22
 - Enabling Packet Logging 23
 - Specifying Activity Summary Settings 24
 - Activity Summary Statistics 25
 - Specifying Top Names Settings 33
 - Top Names Statistics 34
 - Specifying TLS Settings 34
 - TLS Statistics 37
 - Setting Prefetch Timing 37
 - Setting Cache TTLs 37
 - Local Basic or Advanced Web UI 37
 - CLI Commands 38
 - Enabling Smart Caching 38
 - Defining Root Nameservers 39
 - Dynamic Allocation of UDP Ports 40
 - Setting Maximum Memory Cache Sizes 40
 - Specifying Resolver Settings 40
 - Configuring Case Randomization Exclusions 41
 - Specifying Network Settings 42
 - Specifying Advanced Settings 42
 - Enabling Round-Robin 42
 - Flushing Caching DNS Cache 43

Detecting and Preventing DNS Cache Poisoning	43
Handling Unresponsive Nameservers	45
Running DNS Caching Server Commands	45
Configuring Caching DNS Server Network Interfaces	46
Local Advanced Web UI	46

CHAPTER 4**Advanced Caching DNS Server 47**

Using Forwarders	47
Local and Regional Web UI	48
CLI Commands	48
Using Exceptions	49
Local and Regional Web UI	50
CLI Commands	50
Managing DNS64	51
Local and Regional Advanced Web UI	51
CLI Commands	51
Managing DNSSEC	52
Local Advanced Web UI	52
CLI Commands	52
Managing Caching Rate Limiting	53
Client Rate Limiting	53
Domain Rate Limiting	53
Managing Rate Limiting	54
Per Domain Limit	55
CLI Commands	55
Managing DNS Views	56
Setting up Caching DNS and Authoritative DNS Servers on the Same Operating System	56
Managing DNS Firewall	56
Configuring Caching DNS to Use Umbrella	57

CHAPTER 5**Caching DNS Metrics 59**

Caching DNS General Indicators	59
How to Interpret the Data	59
Troubleshooting Based on the Results	59

- DNS Caching Activity 60
 - How to Interpret the Data 60
 - Troubleshooting Based on the Results 60
- DNS Caching Server Queries Per Second 60
- DNS Caching Server Recursion Rate Limit 60
- DNS Incoming Queries 60
 - How to Interpret the Data 61
- DNS Queries Responses 61
 - How to Interpret the Data 61
 - Troubleshooting Based on the Results 61
- DNS Queries Type 62
 - How to Interpret the Data 62
- DNS Recursive Query Time 62
 - How to Interpret the Data 62
 - Troubleshooting Based on the Results 62

PART III

Authoritative DNS Server 63

CHAPTER 6

Managing Authoritative DNS Server 65

- Setting DNS Server Properties 65
 - Setting General DNS Server Properties 66
 - Specifying Log Settings 66
 - Enabling Packet Logging 67
 - Specifying Activity Summary Settings 69
 - Activity Summary Statistics 70
 - Specifying Top Names Settings 92
 - Top Names Statistics 92
 - Specifying TLS Settings 92
 - TLS Statistics 94
 - Enabling Round-Robin 94
 - Enabling Weighted Round-Robin 95
 - Enabling Incremental Zone Transfers (IXFR) 96
 - Restricting Zone Queries 96
 - Enabling NOTIFY 96

Blocking Recursive Queries from Authoritative Server	97
Drop Recursive Queries Statistics	98
Running DNS Authoritative Server Commands	98
Configuring DNS Server Network Interfaces	99
Local Advanced Web UI	99
Managing Authoritative DNSSEC	99
Enabling Authoritative DNSSEC	100
Local Advanced Web UI	101
CLI Commands	101
Managing Authoritative DNSSEC Keys	102
Local and Regional Advanced Web UI	103
CLI Commands	103
Exporting DS Record	103
Setting Advanced Authoritative DNS Server Properties	104
Setting SOA Time to Live	104
Setting Secondary Refresh Times	105
Setting Secondary Retry Times	105
Setting Secondary Expiration Times	105
Setting Local and External Port Numbers	106
Handling Malicious DNS Clients	106
Tuning DNS Properties	106
Running Caching DNS and Authoritative DNS on the Same Server	107
Local Advanced Web UI	108
CLI Commands	109
Troubleshooting DNS Servers	109

CHAPTER 7

DNS Host Health Check	111
DNS Host Health Check Configuration Settings	111
Enabling Host Health Check	112
Local Advanced Web UI	112
CLI Commands	112
Host Health Check RR Set Settings	113
Local Advanced Web UI	113
CLI Commands	113

Viewing DNS Host Health Check Statistics 113

Local Advanced Web UI 113

CLI Commands 114

CHAPTER 8

Managing DNS Firewall 115

Managing DNS Firewall 115

Setting Up RPZ Primary Zones on the Authoritative DNS Server 119

Setting Up DNS Firewall Rules 120

Changing Priority of DNS Firewall Rules 120

Enabling TLS for RPZ 121

CHAPTER 9

Managing High Availability DNS 123

Introduction to HA DNS Processing 123

Creating High Availability DNS Pairs 125

Local Basic or Advanced and Regional Web UI 125

CLI Commands 126

Synchronizing HA DNS Zones 126

Local Advanced Web UI 126

CLI Commands 127

Enable Logging of HA DNS Information 127

Local Basic or Advanced Web UI 127

CLI Command 127

Viewing HA DNS Statistics 127

Local Basic or Advanced Web UI 127

CLI Commands 127

CHAPTER 10

Managing Zones 129

Managing Primary DNS Servers 130

Related Topics 130

Creating and Applying Zone Templates 130

Local and Regional Web UI 130

CLI Commands 132

Staged and Synchronous Modes 132

Local and Regional Web UI 132

CLI Commands	133
Configuring Primary Forward Zones	133
Creating Primary Zones	133
Editing Primary Zones	135
Confirming Zone Nameserver Configuration	136
Synchronizing Zones	136
Zone Commands	137
Importing and Exporting Zone Data	137
Configuring Primary Reverse Zones	139
Related Topics	140
Adding Reverse Zones as Zones	140
Adding Reverse Zones from Subnets	141
Getting Zone Counts on the Server	141
Enabling DNS Updates	142
Managing Secondary Servers	142
Adding Secondary Forward Zones	142
Enabling Zone Transfers	143
Configuring Subzones	144
Related Topics	144
Choosing Subzone Names and Servers	144
Creating and Delegating Subzones	145
Editing Subzone Delegation	146
Undelegating Subzones	146
Managing Zone Distributions	147
Related Topics	147
Preparing the Zone Distribution Map	147
Creating a Zone Distribution	148
Pulling Zone Distributions from Replica Data	150
Managing DNS ENUM Domain	151
Managing DNS ENUM Defaults	151
Adding DNS ENUM Domains	151
Adding DNS ENUM Numbers	152
Pulling and Pushing ENUM Domains	153
Pulling and Pushing ENUM Numbers	154

CHAPTER 11	Managing DNS Views	157
	DNS Views Processing	157
	Key Points to Remember While Working on DNS Views	158
	Managing DNS Views	159
	Local and Regional Web UI	159
	Reorder DNS Views	160
	CLI Commands	160
	Synchronizing DNS Views	160
	Pushing and Pulling DNS Views	160
	Pushing DNS Views to Local Clusters	160
	Regional Web UI	161
	CLI Commands	161
	Pulling DNS Views from Local Clusters	161
	Regional Web UI	161
	CLI Commands	161

CHAPTER 12	Managing Resource Records	163
	Managing Resource Records for Zone	163
	Related Topics	164
	Adding Resource Record to Zone	164
	Local and Regional Web UI	164
	CLI Commands	165
	Editing Resource Records	165
	Removing Resource Records from Zone	165
	Local and Regional Web UI	165
	CLI Commands	165
	Managing Resource Records for Host	166
	Protecting Resource Record Sets	166
	Local and Regional Web UI	166
	Unprotecting Resource Record Sets	166
	CLI Commands	167
	Searching Server-Wide for Records and Addresses	167
	Local Advanced Web UI	167

Local Advanced Web UI	168
CLI Commands	168
Filtering Resource Records	169
Local Basic or Advanced and Regional Web UI	169
CLI Commands	169
Advertising Services to Network Using Service Location (SRV) Records	169
Name Resolution in a Namespace Using NAPTR Resource Records	170
Local Basic or Advanced and Regional Web UI	170
DNS Certification Authority Authorization (CAA) Resource Record	171
Local and Regional Web UI	172
CLI Commands	172
Uniform Resource Identifier (URI) Resource Records	172
Local and Regional Web UI	173
CLI Commands	174

CHAPTER 13
Managing Hosts 175

Adding Hosts in Zones	175
Local Basic or Advanced Web UI	175
CLI Commands	176
Adding Additional RRs for the Host	176
Local Basic or Advanced Web UI	176
CLI Commands	176
Editing Hosts	176
Local Basic or Advanced Web UI	176
CLI Commands	177
Removing Hosts	177
Local Basic or Advanced Web UI	177
CLI Commands	177

CHAPTER 14
Authoritative DNS Metrics 179

DNS General Indicators	179
How to Interpret the Data	179
Troubleshooting Based on the Results	179
DNS Inbound Zone Transfers	180

- How to Interpret the Data 180
- Troubleshooting Based on the Results 180
- DNS Network Errors 180
 - How to Interpret the Data 181
 - Troubleshooting Based on the Results 181
- DNS Outbound Zone Transfers 181
 - How to Interpret the Data 181
 - Troubleshooting Based on the Results 181
- DNS Queries Per Second 181
- DNS Related Servers Errors 181
 - How to Interpret the Data 182
 - Troubleshooting Based on the Results 182

APPENDIX A **Resource Records 183**

- Resource Records 183

APPENDIX B **DNS Anycast with Cisco Prime Network Registrar 193**

- Basic Requirements for DNS Anycast 193
- Anycast Routing 194
 - FRRouting 194
 - Quagga 194
- Script 195
- Router Configuration 195
- Sample Anycast Configuration Using BGP 195
- Network Router Configuration 196
- Configure FRRouting on DNS Servers 197
 - Enable zebra and bgpd in Daemons File 198
 - FRR Zebra Configuration 198
 - FRR BGP Configuration 198
 - Start FRR Service 199
 - Restart FRR Service 199
- Configure Quagga on DNS Servers 199
 - Quagga Zebra Configuration 199
 - Quagga BGP Configuration 200

Start BGP daemon	200
Run Diagnostics on Router	200
Monitor BGP Traffic Logs	201
Configure DNS Zones	202

APPENDIX C

DNS Security and Attack Prevention	203
Prevention of DNS Attacks in Cisco Prime Network Registrar	203



PART I

Introduction

- [Introduction to the Domain Name System, on page 1](#)
- [DNS Server Status Dashboard, on page 9](#)



CHAPTER 1

Introduction to the Domain Name System

The Domain Name System (DNS) handles the growing number of Internet users. DNS translates names, such as `www.cisco.com`, into IP addresses, such as `192.168.40.0` (or the more extended IPv6 addresses), so that computers can communicate with each other. DNS makes using Internet applications, such as the World Wide Web, easy. The process is as if, when phoning your friends and relatives, you could autodial them based on their names instead of having to remember their phone numbers.


- [How DNS Works, on page 1](#)
- [Overview of Concepts in DNS, on page 2](#)

How DNS Works

To understand how DNS works, imagine a typical user, John, logging in to his computer. He launches his web browser so that he can view the website at a company, ExampleCo (see the image below). He enters the name of their website—`http://www.example.com`. Then:

1. John's workstation sends a request to the DNS server about the IP address of `www.example.com`.
2. The DNS server checks its database to find that `www.example.com` corresponds to `192.168.1.4`.
3. The server returns this address to John's browser.
4. The browser uses the address to locate the website.
5. The browser displays the website on John's monitor.

Figure 1: Domain Names and Addresses




Overview of Concepts in DNS

This section provides an overview of the concepts in DNS.

Domains

John can access the ExampleCo website because his DNS server knows the `www.example.com` IP address. The server learned the address by searching through the domain namespace. DNS was designed as a tree structure, where each named domain is a node in the tree. The top-most node of the tree is the DNS root domain (`.`), under which there are subdomains, such as `.com`, `.edu`, `.gov`, and `.mil` (see the image below).

Figure 2: DNS Hierarchy




The fully qualified domain name (FQDN) is a dot-separated string of all the network domains leading back to the root. This name is unique for each host on the Internet. The FQDN for the sample domain is `example.com.`, with its domain `example`, parent domain `.com`, and root domain `“.”` (`dot`).

Learning ExampleCo Address

When John’s workstation requests the IP address of the website `www.example.com` (see the image below):

Figure 3: DNS Hierarchical Name Search



1. The local DNS server looks for the `www.example.com` domain in its database, but cannot find it, indicating that the server is not authoritative for this domain.
2. The server asks the authoritative root nameserver for the top-level (root) domain `“.”` (`dot`).
3. The root nameserver directs the query to a nameserver for the `.com` domain that knows about its subdomains.

4. The .com nameserver determines that example.com is one of its subdomains and responds with its server address.
5. The local server asks the example.com nameserver for the www.example.com location.
6. The example.com nameserver replies that its address is 192.168.1.4.
7. The local server sends this address to John's web browser.

Establishing a Domain

ExampleCo has a website that John could reach because it registered its domain with an accredited domain registry. ExampleCo also entered its domain name in the .com server database, and requested a network number, which defines a range of IP addresses.

In this case, the network number is 192.168.1.0, which includes all assignable hosts in the range 192.168.1.1 through 192.168.1.254. You can only have numbers 0 through 255 (28) in each of the address fields, known as octets. However, the numbers 0 and 255 are reserved for network and broadcast addresses, respectively, and are not used for hosts.

Difference Between Domains and Zones


The domain namespace is divided into areas called zones that are points of delegation in the DNS tree. A zone contains all domains from a certain point downward, except those for which other zones are authoritative.

A zone usually has an authoritative nameserver, often more than one. In an organization, you can have many nameservers, but Internet clients can query only those that the root nameservers know. The other nameservers answer internal queries only.

The ExampleCo company registered its domain, example.com. It established three zones—example.com, marketing.example.com, and finance.example.com. ExampleCo delegated authority for marketing.example.com and finance.example.com to the DNS servers in the marketing and finance groups in the company. If someone queries example.com about hosts in marketing.example.com, example.com directs the query to the marketing.example.com nameserver.

In the image below, the domain example.com includes three zones, with the example.com zone being authoritative only for itself.

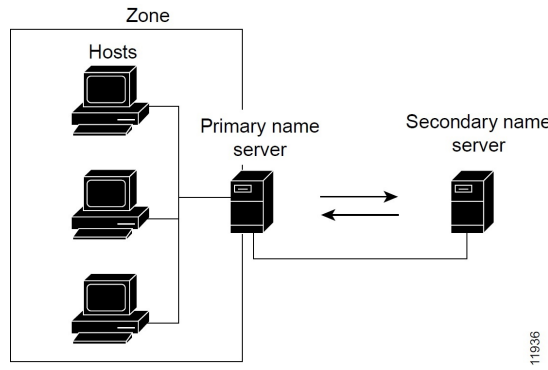
Figure 4: Example.com With Delegated Subdomains



ExampleCo could choose not to delegate authority to its subdomains. In that situation, the example.com domain is a zone that is authoritative for the subdomains for marketing and finance. The example.com server answers all outside queries about marketing and finance.

As you begin to configure zones by using Cisco Prime Network Registrar, you must configure a nameserver for each zone. Each zone has one primary server, which loads the zone contents from a local configuration database. Each zone can also have any number of secondary servers, which load the zone contents by fetching the data from the primary server. The image below shows a configuration with one secondary server.


Figure 5: Primary and Secondary Servers for Zones



Nameservers

DNS is based on a client/server model. In this model, nameservers store data about a portion of the DNS database and provide it to clients that query the nameserver across the network. Nameservers are programs that run on a physical host and store zone data. As administrator for a domain, you set up a nameserver with the database of all the Resource Records (RRs) describing the hosts in your zone or zones (see the image below).

Figure 6: Client/Server Name Resolution



The DNS servers provide name-to-address translation, or name resolution. They interpret the information in an FQDN to find its address.


Each zone must have one primary nameserver that loads the zone contents from a local database, and a number of secondary servers, which load a copy of the data from the primary server (see the image below). This process of updating the secondary server from the primary server is called a zone transfer.

Even though a secondary nameserver acts as a kind of backup to a primary server, both types of servers are authoritative for the zone. They both learn about hostnames in the zone from the zone authoritative database, not from information learned while answering queries. Clients can query both servers for name resolution.

The DNS server functionality is enhanced to provide separate DNS servers for authorization and caching.

As you configure the Cisco Prime Network Registrar DNS nameserver, you specify what role you want the server to perform for a zone—primary or secondary. The type of server is meaningful only in context to its role. An Authoritative DNS server can only be a primary or a secondary server for a zone, it does not specify zones for caching servers.

Figure 7: DNS Zone Transfer



To configure the:

- Primary nameserver, see [Managing Primary DNS Servers, on page 130](#).
- Secondary nameserver, see [Managing Secondary Servers, on page 142](#).

Reverse Nameservers


The DNS servers described so far perform name-to-address resolution. They can do this easily by searching through their database for the correct address, because they index all the data by name. However, there are times when you need address-to-name resolution so that you can interpret certain output, such as computer log files.

Finding a domain name when you only know the address, however, would require searching the entire namespace. DNS solves this problem by supporting a domain namespace that uses addresses as names, known as the `in-addr.arpa` or `.arpa` domain. This reverse zone contains subdomains for each network based on the network number. For consistency and natural grouping, the four octets of a host number are reversed.

The IP address as a domain name appears backward, because the name is in leaf-to-root order. For example, the ExampleCo example domain network number is 192.168.1.0. Its reverse zone is 1.168.192.in-addr.arpa. If you only know the DNS server address (192.168.1.1), the query to the reverse domain would find the host entry 1.1.168.192.in-addr.arpa that maps back to example.com.

Reverse domains are handled through Pointer (PTR) RRs, as indicated in the image below.

Figure 8: Reverse Domains



Authoritative and Caching DNS Servers

In Cisco Prime Network Registrar, the authoritative and caching services are separated, and are handled by two separate servers. The authoritative server holds authoritative zone data and responds only to queries for which it is authoritative. The caching server is the recursive/caching server and does not contain any authoritative zone data.

High-Availability DNS

Because there can be only one primary DNS server per zone, failure of this server makes it impossible to update the zone data. These updates can occur on the primary DNS server only; software such as DHCP servers, that update DNS resource records must send the updates directly to the primary server. A second primary server can become a hot standby that shadows the main primary server. This is called High-Availability (HA) DNS.

EDNS

To send a DNS message above 512 bytes over User Datagram Protocol (UDP), you need to use an extension of the DNS protocol known as Extended DNS (EDNS). The EDNS protocol expands the number of flags, label types, and return codes available to the DNS protocol. A version of EDNS specified by RFC 6891 is known as EDNS0. EDNS uses a pseudo resource record known as OPT Resource Record (OPT RR). OPT RR differentiates conventional DNS from EDNS. OPT RRs appear only in the route transmission between DNS clients and servers, they are not cached or persisted to disk. The DNS client is responsible for an EDNS0 OPT RR to indicate it accepts larger UDP responses.

The Authoritative and Caching DNS servers support the EDNS0 extension. You can modify the UDP payload size of the DNS server. The minimum UDP payload size of the DNS server is 512 bytes. The maximum UDP packet size is 64 KB, the default size for the DNS server is 1232 bytes. Also, the DNS servers will cap the UDP replies to 1232.



Note

The DNS Server can handle requests from clients that do not support EDNS0, however, the DNS server is not permitted to use any extended capabilities, when it handles requests from clients that do not support EDNS0. The response to client requests are inserted into a default 512 byte message. Clients may indicate that they support EDNS by including an OPT RR in the query. If a server does not support EDNS (or the support is disabled), the server will return FORMERR and the client retries without EDNS. If an answer is larger than the size that the client has reported (either with EDNS or the default 512 bytes), the server will mark the result as truncated and the client may retry using TCP.



Note IP fragmentation is a problem on the Internet today, especially when it comes to large DNS messages. Even if fragmentation works, it might not be secure enough for DNS. These issues can be fixed by a) setting the EDNS buffer size lower to limit the risk of IP fragmentation and b) allowing DNS to switch from UDP to TCP when a DNS response is too big to fit in this limited buffer size. The default EDNS buffer size for both the Caching and Authoritative DNS servers is 1232 bytes.

Use the following commands to set the EDNS buffer size:

Authoritative DNS servers:

```
nrcmd> session set visibility=3
nrcmd> dns set edns-max-payload=2000
nrcmd> dns reload
```

Caching DNS servers:

```
nrcmd> session set visibility=3
nrcmd> cdns set edns-buffer-size=2000
nrcmd> cdns set max-udp-size=2000
nrcmd> cdns reload
```

DNS Views

DNS Views allow you to present alternate versions of zone data to different communities of clients using a single name server.

For example, a DNS server for example.com could maintain two views of the zone, where the view of example.com that can be queried internally includes many hosts that do not exist in the external view. Each zone view is treated as an independent copy of the zone. The DNS server, when answering queries on the zone, uses the match criteria defined in each view to determine the matching zone for the client. The query is then answered based on that zone contents.

In Cisco Prime Network Registrar 11.0, zones can be referenced by multiple views without the need to make copies of the zone. Use the *alternate-view-ids* attribute for this purpose. For more information, see [Key Points to Remember While Working on DNS Views, on page 158](#).



CHAPTER 2

DNS Server Status Dashboard

The Cisco Prime Network Registrar server status dashboard in the web user interface (web UI) presents a graphical view of the system status, using graphs, charts, and tables, to help in tracking and diagnosis. These dashboard elements are designed to convey system information in an organized and consolidated way, and include:

- Significant protocol server and other metrics
- Alarms and alerts
- Database inventories
- Server health trends

The dashboard is best used in a troubleshooting desk context, where the system displaying the dashboard is dedicated for that purpose and might be distinct from the systems running the protocol servers. The dashboard system should point its browser to the system running the protocol servers.

You should interpret dashboard indicators in terms of deviations from your expected normal usage pattern. If you notice unusual spikes or drops in activity, there could be communication failures or power outages on the network that you need to investigate.

- [Opening the Dashboard, on page 9](#)
- [Display Types, on page 9](#)
- [Customizing the Display, on page 14](#)
- [Selecting Dashboard Elements to Include, on page 15](#)

Opening the Dashboard

The Dashboard feature is available on the regional cluster also. It provides System Metrics chart by default. It allows you to display the server specific (DHCP, DNS, and CDNS) charts for various clusters. This can be configured in the Chart Selections page.

To open the dashboard in the web UI, from the **Operate** menu, choose **Dashboard**.

Display Types


Provided you have DNS and Caching DNS privileges through administrator roles assigned to you, the preset display of the dashboard consists of the following tables (See the table below for an example):

- **System Metrics**—See the "System Metrics" section in *Cisco Prime Network Registrar 11.0 Administration Guide*.
- **DNS General Indicators**—See [Caching DNS Metrics](#), on page 59 and [Authoritative DNS Metrics](#), on page 179.



Tip These are just the preset selections. See [Selecting Dashboard Elements to Include](#), on page 15 for other dashboard elements you can select. The dashboard retains your selections from session to session.

Figure 9: Preset Dashboard Elements



Each dashboard element initially appears as a table or a specific panel chart, depending on the element:

- **Table**—See [Tables](#), on page 11.
- **Line chart**—See [Line Charts](#), on page 11.
- **Area chart**—See [Area Charts](#), on page 13.

General Status Indicators

Note the green indicator in the Server State description in the above image. This indicates that the server sourcing the information is functioning normally. A yellow indicator indicates that server operation is less than optimum. A red indicator indicates that the server is down. These indicators are the same as for the server health on the Manage Servers page in the regular web UI.

Graphic Indicators for Levels of Alert

Graphed lines and stacked areas in the charts follow a standard color and visual coding so that you can immediately determine key diagnostic indicators at a glance. The charts use the following color and textural indicators:

- **High alerts or warnings**—Lines or areas in red, with a hatched texture.
- **All other indicators**—Lines or areas in various other colors distinguish the data elements. The charts do not use green or yellow.

Magnifying and Converting Charts

You can magnify a chart in a separate window by clicking the **Chart Link** icon at the bottom of the panel chart and then by clicking the **Magnified Chart** option (see the image below). In magnified chart view, you can choose an alternative chart type from the one that comes up initially (see [Other Chart Types, on page 13](#)).

Figure 10: Magnifying Charts



Note Automatic refresh is turned off for magnified charts. To get the most recent data, click the **Refresh** icon next to the word Dashboard at the top left of the page.

To convert a chart to a table, see the *Displaying Charts as Tables* section. You cannot convert tables to a graphic chart format.

Legends

Each chart includes a color-coded legend by default.

Tables

Dashboard elements rendered as tables have data displayed in rows and columns. The following dashboard elements are preset to consist of (or include) tables:

- DHCP DNS Updates
- DHCP Address Current Utilization
- DHCP General Indicators
- DNS General Indicators
- Caching DNS General Indicators



Note If you view a table in Expert mode, additional data might appear.

Line Charts

Dashboard elements rendered as line charts can include one or more lines plotted against the x and y axes. The three types of line charts are described in the following table.

Table 1: Line Chart Types

Type of Line Chart	Description	Dashboard Elements Rendered
Raw data line chart	Lines plotted against raw data.	<ul style="list-style-type: none"> • Java Virtual Machine (JVM) Memory Utilization (Expert mode only) • DHCP Buffer Capacity • DHCP Failover Status (two charts) • DNS Network Errors • DNS Related Servers Errors
Delta line chart	Lines plotted against the difference between two sequential raw data.	<ul style="list-style-type: none"> • DNS Inbound Zone Transfers • DNS Outbound Zone Transfers
Rate line chart	Lines plotted against the difference between two sequential raw data divided by the sample time between them.	<ul style="list-style-type: none"> • DHCP Server Request Activity (see the image below) • DHCP Server Response Activity • DHCP Response Latency • DNS Query Responses • DNS Forwarding Errors



Tip To get the raw data for a chart that shows delta or rate data, enter Expert mode, go to the required chart, click the **Chart Link** icon at the bottom of the panel chart, and then click **Data Table** . The Raw Data table is below the Chart Data table.

Figure 11: Line Chart Example



Area Charts

Dashboard elements rendered as area charts have multiple related metrics plotted as trend charts, but stacked one on top of the other, so that the highest point represents a cumulative value. The values are independently shaded in contrasting colors. (See the image below for an example of the DHCP Server Request Activity chart shown in [Figure 11: Line Chart Example, on page 12](#) rendered as an area chart.)

Figure 12: Area Chart Example



They are stacked in the order listed in the legend, the left-most legend item at the bottom of the stack and the right-most legend item at the top of the stack. The dashboard elements that are pre-set to area chart are:

- DHCP Buffer Capacity
- DHCP Failover Status
- DHCP Response Latency
- DHCP Server Leases Per Second
- DHCP Server Request Activity
- DHCP Server Response Activity
- DNS Inbound Zone Transfers
- DNS Network Errors
- DNS Outbound Zone Transfers
- DNS Queries Per Second
- DNS Related Server Errors

Other Chart Types

The other chart types available for you to choose are:

- **Line**—One of the line charts described in [Line Charts, on page 11](#).
- **Area**—Charts described in the [Area Charts, on page 13](#).
- **Column**—Displays vertical bars going across the chart horizontally, with the values axis being displayed on the left side of the chart.

- **Scatter**—A scatter plot is a type of plot or mathematical diagram using Cartesian coordinates to display values for typically two variables for a set of data.



Tip Each chart type shows the data in distinct ways and in different interpretations. You can decide which type best suits your needs.

Getting Help for the Dashboard Elements

You can open a help window for each dashboard element by clicking the help icon on the table/chart window.

Customizing the Display

To customize the dashboard display, you can:

- Refresh the data and set an automatic refresh interval.
- Expand a chart and render it in a different format.
- Convert a graphic chart to a table.
- Download data to comma-separated value (CSV) output.
- Display or hide chart legends.
- Configure server chart types.
- Reset to default display

Each chart supports:

- Resizing
- Drag and drop to new cell position
- Minimizing
- Closing

Each chart has a help icon with a description of the chart and a detailed help if you click the link (more...) at the bottom of the description.



Note The changes made to the dashboard/chart will persist only if you click **Save** in the Dashboard window.


Refreshing Displays

Refresh each display so that it picks up the most recent polling by clicking the **Refresh** icon.

Setting the Polling Interval

You can set how often to poll for data. Click the **Dashboard Settings** icon in the upper-right corner of the dashboard display. There are four options to set the polling interval of the cached data, which polls the protocol servers for updates (See the image below).

Figure 13: Setting the Chart Polling Interval



You can set the cached data polling (hence, automatic refresh) interval to:

- **Disabled**—Does not poll, therefore does not automatically refresh the data.
- **Slow**—Refreshes the data every 30 seconds.
- **Medium**—Refreshes the data every 20 seconds.
- **Fast** (the preset value)—Refreshes the data every 10 seconds.

Displaying Charts as Tables

Use the **Chart Link** icon at the bottom of the panel chart to view the chart link options (see the image below). You can choose to display a graphic chart as a table by clicking the **Data Table** option.

Figure 14: Specifying Chart Conversion to Table Format



Exporting to CSV Format

You can dump the chart data to a comma-separated value (CSV) file (such as a spreadsheet). In the Chart Link controls at the bottom of the panel charts (see the above image), click the **CSV Export** option. A Save As window appears, where you can specify the name and location of the CSV file.

Selecting Dashboard Elements to Include

You can decide how many dashboard elements you want to display on the page. At times, you might want to focus on one server activity only, such as for the DHCP server or the DNS server, and exclude all other metrics for the other servers. In this way, the dashboard becomes less crowded, the elements are larger and more readable. At other times, you might want an overview of all server activities, with a resulting smaller element display.

You can select the dashboard elements to display from the main Dashboard page by clicking the Dashboard Settings icon and then clicking **Chart Selections** in the Dashboard Settings dialog. Clicking the link opens the Chart Selection page (see [Figure 15: Selecting Dashboard Elements, on page 16](#)).

Configuring Server Chart Types

You can set the default chart types on the main dashboard view. You can customize the server charts in the dashboard to display only the specific chart types as default.

To set up default chart type, check the check box corresponding to the Metrics chart that you want to display and choose a chart type from the **Type** drop-down list. The default chart types are consistent and shared across different user sessions (see the image below).




Note You can see either the CDNS or DNS Metrics in the **Dashboard Settings > Chart Selection** page based on the service configured on the server.



Tip The order in which the dashboard elements appear in the Chart Selection list does not necessarily determine the order in which the elements will appear on the page. An algorithm that considers the available space determines the order and size in a grid layout. The layout might be different each time you submit the dashboard element selections. To change selections, check the check box next to the dashboard element that you want to display.

Figure 15: Selecting Dashboard Elements



The above image displays the Charts Selection table in the regional web UI. The **Clusters** column is available only in regional dashboard and it displays the list of local clusters configured. You can add the local cluster by clicking the Edit icon and then by selecting the local cluster name from the Local Cluster List dialog box.

To change selections, check the check box next to the dashboard element that you want to display.

Specific group controls are available in the **Change Chart Selection** drop-down list, at the top of the page (see the image above). To:


- Uncheck all check boxes, choose **None**.

- Revert to the preset selections, choose **Default**. The preset dashboard elements for administrator roles supporting DHCP and DNS are:
 - Host Metrics: System Metrics
 - DHCP Metrics: General Indicators
 - DNS Metrics: General Indicators
- Select the DHCP metrics only, choose **DHCP** (see the "*DHCP Metrics*" section in *Cisco Prime Network Registrar 11.0 DHCP User Guide*).
- Select the DNS metrics only, choose **DNS** (see the "*Authoritative DNS Metrics*" section in *Cisco Prime Network Registrar 11.0 Authoritative and Caching DNS User Guide*).
- Select the DNS metrics only, choose **CDNS** (see the "*Caching DNS Metrics*" section in *Cisco Prime Network Registrar 11.0 Authoritative and Caching DNS User Guide*)
- Select all the dashboard elements, choose **All**.

Click **OK** at the bottom of the page to save your choices, or **Cancel** to cancel the changes.

You can change the chart type by clicking the **Chart Type** icon at the bottom of the panel chart and then by selecting the required chart type (see the image below). The different types of chart available are: Line Chart, Column Chart, Area Chart, and Scatter Chart.

Figure 16: Selecting the Chart Type





PART II

Caching DNS Server

- [Managing Caching DNS Server, on page 21](#)
- [Advanced Caching DNS Server, on page 47](#)
- [Caching DNS Metrics, on page 59](#)



CHAPTER 3

Managing Caching DNS Server

In Cisco Prime Network Registrar, the authoritative and caching services are separated, and are handled by two separate servers. This chapter explains how to set the Caching DNS server parameters. Before you proceed with the tasks in this chapter, see [Introduction to the Domain Name System, on page 1](#), which explains the basics of DNS.

- [Setting DNS Caching Server Properties, on page 21](#)
- [Running DNS Caching Server Commands, on page 45](#)
- [Configuring Caching DNS Server Network Interfaces, on page 46](#)

Setting DNS Caching Server Properties

You can set properties for the Caching DNS server. These include:

- **General server properties**—See [Setting General Caching DNS Server Properties, on page 22](#)
- **Log settings**—See [Specifying Log Settings, on page 22](#)
- **Packet logging**—See [Enabling Packet Logging, on page 23](#)
- **Activity summary settings**—See [Specifying Activity Summary Settings, on page 24](#)
- **Top names settings**—See [Specifying Top Names Settings, on page 33](#)
- **TLS settings**—See [Specifying TLS Settings, on page 34](#)
- **Caching settings**—See [Setting Prefetch Timing, on page 37](#)
- **Cache TTLs**—See [Setting Cache TTLs, on page 37](#)
- **Smart caching**—See [Enabling Smart Caching, on page 38](#)
- **Root name servers**—See [Defining Root Nameservers, on page 39](#)
- **UDP ports**—See [Dynamic Allocation of UDP Ports, on page 40](#)
- **Maximum memory cache sizes**—See [Setting Maximum Memory Cache Sizes, on page 40](#)
- **Resolver settings**—See [Specifying Resolver Settings, on page 40](#)
- **Network settings**—See [Specifying Network Settings, on page 42](#)
- **Advanced settings**—See [Specifying Advanced Settings, on page 42](#)

- **Flush cache**—See [Flushing Caching DNS Cache, on page 43](#)
- **Prevent DNS cache poisoning**—See [Detecting and Preventing DNS Cache Poisoning, on page 43](#)
- **Handle unresponsive nameservers**—See [Handling Unresponsive Nameservers, on page 45](#)

Setting General Caching DNS Server Properties

You can view general Caching DNS server properties, such as log settings, basic cache settings, SNMP traps, and root nameservers.

The following subsections describe some of the most common property settings. They are listed in [Setting DNS Caching Server Properties, on page 21](#).

Local Basic or Advanced Web UI

-
- Step 1** To access the server properties, from the **Deploy** menu, choose **CDNS Server** under the **DNS** submenu to open the Manage DNS Caching Server page.
 - Step 2** The local CDNS Server page is automatically selected when you choose the **CDNS Server** tab, either from the Deploy menu or by clicking the **CDNS Server** tab in the left pane. The page displays all the Caching DNS server attributes.
 - Step 3** Click **Save** to save the Caching DNS server attribute modifications.
-

CLI Commands

Use **cdns show** to display the Caching DNS server properties (see the **cdns** command in the CLIGuide.html file in the /docs directory for syntax and attribute descriptions).

Specifying Log Settings

The *log-settings* attribute determines which detailed events the Caching DNS server logs. Logging these additional details can help analyze a problem. However, leaving detailed logging enabled for a long period, can fill the log files and cause the loss of important information.

The possible options are:

- **activity-summary**—Causes logging of a server statistics summary at a regular interval.
- **config**—Controls logging pertaining to server configuration and server de-initialization.
- **query**—Causes logging of all DNS queries to the server.
- **scp**—Controls logging pertaining to SCP message processing.
- **server-detailed-ops**—Controls detailed logging of server operations.
- **server-ops**—Controls high level logging of server operations.
- **name-servers**—Enables logging when name servers for exceptions and forwarders become unresponsive or again become responsive.

The *immediate-response-stats* attribute (available in Advanced mode) enables collecting response times statistics when queries are answered immediately. If this feature is disabled, the related statistics (*immediate-response-count*, *immediate-response-average*, and *immediate-response-median*) will show zero.

Enabling Packet Logging

Cisco Prime Network Registrar supports packet logging for Caching DNS server to help analyze and debug the Caching DNS server activity. The packet logging settings determine the type of packet logging (summary or detail), the type of packets logged, and to which log file the messages are logged. By default, the Caching DNS server does not log any packet log messages.

Use the following server level attributes to enable packet logging for the Caching DNS server:

Table 2: Caching DNS Server Packet Logging Attributes

Attribute	Description
Packet Logging (<i>packet-logging</i>)	<p>Determines the type of packet logging that is logged to the CDNS logs. The type of packets logged can be controlled with the <i>packet-log-settings</i> attribute.</p> <ul style="list-style-type: none"> • disabled—This settings disables packet logging. • summary—This setting enables one line summary packet logging. • detail—This setting enables detailed packet tracing. <p>Note: While packet logging can be helpful for debugging and troubleshooting, it does have an impact on DNS server performance. Therefore, Cisco does not recommend leaving packet logging enabled in production environments.</p>
Packet Logging File (<i>packet-logging-file</i>)	<p>Determines the destination log of packet log messages when packet logging is enabled.</p> <ul style="list-style-type: none"> • cdns—Packet logging messages are logged to the standard CDNS log file (<i>cdns_log*</i>). • packet—Packet logging messages are logged to a separate CDNS packet log file (<i>cdns_query_log*</i>).
Packet Log Settings (<i>packet-log-settings</i>)	<p>Determines the type of packets to log when packet logging is enabled. Packet logging can be enabled by configuring the <i>packet-logging</i> attribute.</p> <ul style="list-style-type: none"> • query-in—This setting enables logging of incoming query packets. These are packets coming in from DNS clients. • query-out—This setting enables logging of outgoing query packets. These are queries going to upstream DNS servers. • response-in—This setting enables logging of incoming query response packets. These are responses coming from upstream DNS servers. • response-out—This setting enables logging of outgoing query response packets. These are responses going to DNS clients.

Local Advanced Web UI

Step 1 On the Manage DNS Caching Server page, under the **Packet Logging** section, select the value for **packet-logging** from the drop-down list. The value can be **summary** or **detail**.

Step 2 For the *packet-log-settings* attribute, check the desired check boxes.

Step 3 Click **Save** to save the changes.

CLI Commands

Use **cdns set packet-logging=summary** to enable one line summary packet logging.

Use **cdns set packet-logging=detail** to enable detailed packet tracing.

Use **cdns set packet-log-settings=value** to set the type of packets to log when packet logging is enabled.



Note Reloading of Caching DNS server is not required for the *packet-logging* and *packet-log-settings* attributes to take effect immediately (similar to log settings). However, the *packet-logging-file* attribute requires a Caching DNS server reload.

Specifying Activity Summary Settings



Note To specify the activity summary settings, you have to check *activity-summary* under Log Settings.

You can specify the interval at which to log activity summary information using the Statistics Interval (*activity-summary-interval*) attribute. It has a default value of 60 seconds.

The Caching DNS server logs sample and/or total statistics based on the option you check for the Statistics Type (*activity-summary-type*) attribute. The default value is "sample".

The option checked for the Statistics Settings (*activity-summary-settings*) attribute determines the category of statistics that is logged as part of activity summary. The possible settings are:

- **cache**—Logs statistics on the RR cache.

For the list of activity summary statistics that are displayed in the logs for the **cache** setting, see [Cache Statistics, on page 25](#).

- **firewall**— Logs statistics on DNS firewall usage.

For the list of activity summary statistics that are displayed in the logs for the **firewall** setting, see [Firewall Statistics, on page 26](#).

- **memory**—Logs statistics on memory usage.

For the list of activity summary statistics that are displayed in the logs for the **memory** setting, see [Memory Statistics, on page 27](#).

- **query**—Logs statistics related to incoming queries.

For the list of activity summary statistics that are displayed in the logs for the **query** setting, see [Query Statistics, on page 27](#).

- **query-type**—Logs statistics on the RR types that are being queried.

For the list of activity summary statistics that are displayed in the logs for the **query-type** setting, see [Query by Type Statistics, on page 29](#).

- **rate-limiting**—Logs the number of rate limiting events.

For the list of activity summary statistics that are displayed in the logs for the **rate-limiting** setting, see [Rate Limiting Statistics, on page 29](#).

- **resol-queue**—Logs statistics on the resolution queue.

For the list of activity summary statistics that are displayed in the logs for the **resol-queue** setting, see [Resolution Queue Statistics, on page 30](#).

- **responses**—Logs statistics about query responses.

For the list of activity summary statistics that are displayed in the logs for the **responses** setting, see [Responses Statistics, on page 31](#).

- **system**—Logs statistics on system usage.

For the list of activity summary statistics that are displayed in the logs for the **system** setting, see [System Statistics, on page 33](#).

- **top-names**—Logs the top names queried and hit count.

For the list of activity summary statistics that are displayed in the logs for the **top-names** setting, see [Top Names Statistics, on page 33](#).

Activity Summary Statistics

Following sections describe the list of activity summary statistics that are displayed in the logs under each of the *activity-summary-settings* category.

Cache Statistics

The **cache** activity-summary-settings logs statistics on the RR cache.

Sample log message:

```
10/06/2021 10:22:44 cdns Activity Stats 0 22173 [Cache] Sample since Wed Oct 6 10:21:44
2021: hits=number, misses=number, prefetches=number, message-overflow=number,
rrset-overflow=number, remote-ns-overflow=number, key-overflow=number, smart-cache=number
```

Table 3: Cache Statistics

Activity Summary Name	Statistic ¹	Description
hits	cache-hits	Total number of queries that were answered from cache.
misses	cache-misses	Total number of queries that were not found in the cache.
prefetches	cache-prefetches	Number of prefetches performed.

Activity Summary Name	Statistic ¹	Description
rrset-overflow	mem-cache-exceeded	Number of times the RRSet cache has gone over the configured limit. This indicates that the configured limit may be undersized for its environment.
message-overflow	mem-query-cache-exceeded	Number of times the message cache has gone over the configured limit. This indicates that the configured limit may be undersized for its environment.
remote-ns-overflow	remote-ns-cache-exceeded	Number of times the remote name server cache has gone over the configured limit. This indicates that the configured limit may be undersized for its environment.
key-overflow	key-cache-exceeded	Number of times the key cache has gone over the configured limit. This indicates that the configured limit may be undersized for its environment.
smart-cache	smart-cache	Total number of times the CDNS Server employed a smart-cache response, when <i>smart-cache</i> is enabled.

¹ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Caching DNS server statistics, see the "CDNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Firewall Statistics

The **firewall** activity-summary-settings logs statistics on DNS Firewall usage.

Sample log message:

```
11/18/2021 12:39:20 cdns Activity Stats 0 22322 [Firewall] Sample since Thu Nov 18 12:38:20
2021: redirected=number, dropped=number, refused=number, redirect-nxdomain=number, rpz=number
```

Table 4: Firewall Statistics

Activity Summary Name	Statistic ²	Description
dropped	firewall-dropped	Number of times DNS Firewall dropped a query.
redirected	firewall-redirected	Number of times DNS Firewall redirected a query.
refused	firewall-refused	Number of times DNS Firewall refused a query.
redirect-nxdomain	firewall-redirect-nxdomain	Number of times DNS Firewall redirected a query with an NXDOMAIN answer.

Activity Summary Name	Statistic ²	Description
rpz	firewall-rpz	Number of times DNS Firewall RPZ rules matched an incoming query.

² The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Caching DNS server statistics, see the "CDNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Memory Statistics

The **memory** activity-summary-settings logs statistics on memory usage.

Sample log message:

```
10/06/2021 10:22:44 cdns Activity Stats 0 22303 [Memory] Current: mem-cache-process=number,
mem-cache-rrset=number, mem-cache-message=number, mem-mod-iterator=number,
mem-mod-validator=number
```

Table 5: Memory Statistics

Activity Summary Name	Statistic ³	Description
mem-cache-process	mem-process	An estimate of the memory in bytes of the CDNS process.
mem-cache-rrset	mem-cache	Memory in bytes allocated to the RRset cache. Note that the allocated memory will be maintained across server reloads, unless the <i>rrset-cache-size</i> configuration has changed.
mem-cache-message	mem-query-cache	Memory in bytes allocated to the message cache. Note that the allocated memory will be maintained across server reloads, unless the <i>msg-cache-size</i> configuration has changed.
mem-mod-iterator	mem-iterator	Memory in bytes used by the CDNS iterator module.
mem-mod-validator	mem-validator	Memory in bytes used by the CDNS validator module.

³ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Caching DNS server statistics, see the "CDNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Query Statistics

The **query** activity-summary-settings logs statistics related to incoming queries.

Sample log message:

```
10/06/2021 10:22:44 cdns Activity Stats 0 22171 [Query] Sample since Wed Oct 6 10:21:44
2021: total=number, acl-failures=number, udp=number, tcp=number, ipv4=number, ipv6=number,
  tls=number, tls-errors-in=number, tls-errors-out=number, edns=number, dnssec=number,
dns64-aaaa=number, dns64-ptr=number, dns64-ns=number, unwanted-class=number
```

Table 6: Query Statistics

Activity Summary Name	Statistic ⁴	Description
total	queries-total	Total number of queries received by the CDNS Server.
acl-failures	queries-failing-acl	Number of queries being dropped or refused due to ACL failures.
tcp	queries-over-tcp	Total number of queries received over TCP by the CDNS Server.
udp	N/A	Total number of queries received over UDP by the CDNS Server.
ipv4	N/A	Total number of IPv4 queries received by the CDNS Server.
ipv6	queries-over-ipv6	Total number of IPv6 queries received by the CDNS Server.
tls	queries-over-tls	Total number of queries received over TLS by the CDNS Server.
tls-errors-in	tls-errors-in	Total number of TLS related errors on inbound DNS query attempts.
tls-errors-out	tls-errors-out	Total number of TLS related errors on outbound DNS query attempts.
edns	queries-with-edns	Number of queries with EDNS OPT RR present.
dnssec	queries-with-edns-do	Number of queries with EDNS OPT RR with DO (DNSSEC OK) bit set.
dns64-aaaa	dns64-a2aaaa-conversions	Number of times dns64 has converted a type A RR to a type AAAA RR.
dns64-ptr	dns64-ptr-conversions	Number of times dns64 has converted an IPv4 PTR RR to an IPv6 PTR RR.
unwanted-class	queries-unwanted-class	Total number of queries with an unwanted classes.

⁴ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Caching DNS server statistics, see the "CDNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Query by Type Statistics

The **query-type** activity-summary-settings logs statistics on the RR types that are being queried.

Sample log message:

```
10/06/2021 10:22:44 cdns Activity Stats 0 22172 [Query-by-Type] Sample since Wed Oct 6
10:21:44 2021: A=number, AAAA=number, ANY=number, CNAME=number, PTR=number, MX=number,
NS=number, SOA=number, DS=number, DNSKEY=number, RRSIG=number, NSEC=number, NSEC3=number,
Other=number
```

Table 7: Query by Type Statistics

Activity Summary Name	Statistic ⁵	Description
A	queries-type-A	Number of A queries received.
AAAA	queries-type-AAAA	Number of AAAA queries received.
CNAME	queries-type-CNAME	Number of CNAME queries received.
PTR	queries-type-PTR	Number of PTR queries received.
NS	queries-type-NS	Number of NS queries received.
SOA	queries-type-SOA	Number of SOA queries received.
MX	queries-type-MX	Number of MX queries received.
DS	queries-type-DS	Number of DS queries received.
DNSKEY	queries-type-DNSKEY	Number of DNSKEY queries received.
RRSIG	queries-type-RRSIG	Number of RRSIG queries received.
NSEC	queries-type-NSEC	Number of NSEC queries received.
NSEC3	queries-type-NSEC3	Number of NSEC3 queries received.
Other	queries-type-other	All other queries received.
ANY	queries-type-ANY	Number of ANY queries received.

⁵ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Caching DNS server statistics, see the "CDNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Rate Limiting Statistics

The **rate-limiting** activity-summary-settings logs the number of rate limiting events.

Sample log message:

```
11/30/2021 16:20:37 cdns tid: 0 Activity Stats 0 22388 [Ratelimit] Sample since Tue Nov 30
16:19:37 2021: client-ratelimited=number, domain-ratelimited=number
```

```
11/30/2021 16:20:37 cdns tid: 0 Activity Stats 0 22390 [Ratelimit-Domain] from 16:19:37 to
16:20:33; interval=number, num-ratelimited=number, total-counted=number, not-counted=number
```

```
11/30/2021 16:20:37 cdns tid: 0 Activity Stats 0 22390 [Ratelimit-Client] from 08:29:43 to
08:30:43; interval=number, num-ratelimited=number, total-counted=number, not-counted=number
```

Table 8: Rate Limiting Statistics

Activity Summary Name	Logging Sub Category	Statistic ⁶	Description
client-ratelimited	Ratelimit	client-rate-limit	Number of times a client was rate limited.
domain-ratelimited	Ratelimit	domain-rate-limit	Number of times a domain was rate limited.
interval	Ratelimit-Domain	N/A	Length of data collection period.
num-ratelimited	Ratelimit-Domain	N/A	Total number of domains that were rate limited.
total-counted	Ratelimit-Domain	N/A	Total number of times a domain was rate limited.
not-counted	Ratelimit-Domain	N/A	Number of times the domain rate limiting table overflowed.
interval	Ratelimit-Client	N/A	Length of data collection period.
num-ratelimited	Ratelimit-Client	N/A	Total number of clients that were rate limited.
total-counted	Ratelimit-Client	N/A	Total number of times a client was rate limited.
not-counted	Ratelimit-Client	N/A	Number of times the client rate limiting table overflowed.

⁶ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Caching DNS server statistics, see the "CDNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Resolution Queue Statistics

The **resol-queue** activity-summary-settings logs statistics on the resolution queue.

Sample log message:

```
10/06/2021 10:22:44 cdns Activity Stats 0 22174 [Resolution-Queue] Sample since Wed Oct 6
10:21:44 2021: num-entries=number, user-queries=number, system-queries=number,
```

average-num-entries=*number*, max-num-entries=*number*, entries-overwritten=*number*, exceeded-limit=*number*, replies-sent=*number* exceeded-max-target-count=*number*

Table 9: Resolution Queue Statistics

Activity Summary Name	Statistic ⁷	Description
num-entries	requestlist-total	Total number of queued requests waiting for recursive replies.
user-queries	requestlist-total-user	Total number of queued user requests waiting for recursive replies.
system-queries	requestlist-total-system	Total number of queued system requests waiting for recursive replies.
average-num-entries	requestlist-total-average	Average number of requests on the request list.
max-num-entries	requestlist-total-max	Maximum number of requests on the request list.
entries-overwritten	requestlist-total-overwritten	Number of requests on the request list that were overwritten by newer entries.
exceeded-limit	requestlist-total-exceeded	Number of requests dropped because the request list was full.
replies-sent	recursive-replies-total	Total number of query replies that were not found in the cache and required external resolution.
exceeded-max-target-count	exceeded-max-target-count	Number of queries that exceeded the maximum number of name servers glue lookups allowed.

⁷ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Caching DNS server statistics, see the "CDNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Responses Statistics

The **responses** activity-summary-settings logs statistics about query responses.

Sample log message:

```
10/06/2021 10:22:44 cdns Activity Stats 0 22175 [Responses] Sample since Wed Oct 6 10:21:44
2021: no-error=number, no-data=number, formerr=number, servfail=number, nxdomain=number,
notimp=number, refused=number, notauth=number, other-errors=number, secure=number,
unsecure=number, rrsset-unsecure=number, unwanted=number
```

Table 10: Responses Statistics

Activity Summary Name	Statistic ⁸	Description
no-error	answers-with-NOERROR	Number of answers from cache or recursion that result in rcode of NOERROR being returned to client.
nxdomain	answers-with- NXDOMAIN	Number of answers from cache or recursion that result in rcode of NXDOMAIN being returned to client.
no-data	answers-with-NODATA	Number of answers that result in pseudo rcode of NODATA being returned to client.
other-errors	answers-with-other-errors	Number of answers that result in pseudo rcode of NODATA being returned to client.
secure	answers-secure	Number of answers that correctly validated.
unsecure	answers-unsecure	Number of answers that did not correctly validate.
rrset-unsecure	answers-rrset-unsecure	Number of RRSets marked as bogus by the validator.
unwanted	answers-unwanted	Number of replies that were unwanted or unsolicited. High values could indicate spoofing threat.
refused	answers-with-REFUSED	Number of answers from cache or recursion that result in rcode of REFUSED being returned to client.
servfail	answers-with-SERVFAIL	Number of answers from cache or recursion that result in rcode of SERVFAIL being returned to client.
formerr	answers-with-FORMERR	Number of answers from cache or recursion that result in rcode of FORMERR being returned to client.
notauth	answers-with-NOTAUTH	Number of answers from cache or recursion that result in rcode of NOTAUTH being returned to client.
notimp	answers-with-NOTIMP	Number of answers from cache or recursion that result in rcode of NOTIMP being returned to client.

⁸ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Caching DNS server statistics, see the "CDNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

System Statistics

The **system** activity-summary-settings logs statistics on system usage.

Sample log message:

```
10/26/2021 6:04:44 cdns tid: 0 Activity Stats 0 22375 [System] Current: contrack-max=number,
contrack-count=number, contrack-usage=number
```

Table 11: System Statistics

Activity Summary Name	Description
contrack-max	Maximum number of connection tracking entries allowed.
contrack-count	Number of connection tracking entries currently in use.
contrack-usage	Percentage of connection tracking entries in use.

Top Names Statistics

The **top-names** activity-summary-settings logs the top names queried and hit count.

Sample log message:

```
10/26/2021 12:07:08 cdns Activity Stats 0 22371 [Top-Names] from 12:06:48 to 12:06:58;
interval=number, total-counted=number
```

Table 12: Top Names Statistics

Activity Summary Name	Statistic ⁹	Description
interval	N/A	Length of data collection period. It corresponds to the CDNS <i>top-names-max-age</i> setting, which controls how long it has to collect the top names for each log entry. It then lists a configurable number of top names (default 10) and the number of queries for those names.
total-counted	total-counted	Total number of queries counted in this collection period.

⁹ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Caching DNS server statistics, see the "CDNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Specifying Top Names Settings

The *top-names* attribute specifies if top names data should be collected. When enabled, a snapshot of the cache hits for the top names that are queried is collected for each interval set by the *top-names-max-age* value. The list of top names that is reported with activity summary statistics is the most current snapshot.

You can specify the maximum age (based on last access time) of a queried name allowed in the list of top names by using the *top-names-max-age* attribute.



Note The *top-names-max-age* attribute has a default value of 60 seconds.

You can specify the maximum number of entries in the list of top names queried by using the *top-names-max-count* attribute. This limit is applied to the lists of top names that are logged or returned as part of activity summary. The default value is 10.

Local Basic or Advanced Web UI

To enable Top Names, on the Edit Local CDNS Server tab, under the **Top Names Settings** section, enable the *top-names* attribute by selecting the **enabled** option, and then click **Save** to save the changes.

Top Names Statistics

The Top Names tab displays the relevant information with respect to top N domains and other important statistics attributes.

Local Basic or Advanced Web UI

-
- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page.
 - Step 2** Select **CDNS** in the Manage Servers pane.
 - Step 3** Click the **Top Names** tab available in the Local CDNS Server page.
-

CLI Commands

Use `cdns getStats top-names` to view the Top Names statistics.

Specifying TLS Settings

DNS queries without encryption are vulnerable to spoofing and other attacks that threaten privacy. To address these issues, Cisco Prime Network Registrar 11.0 supports DNS over TLS (DoT) as specified by RFC 7858 for both Authoritative DNS server and Caching DNS server.

DNS over TLS is a security protocol for encrypting and wrapping DNS queries and answers via the Transport Layer Security (TLS) protocol. It improves privacy and security between clients and resolvers. It uses TCP as the basic connection protocol and layers over TLS encryption and authentication.

TLS Keys

TLS key pair consists of a private key and a public key. These two keys are related to one another by means of a cryptographic algorithm. The private key is “private” to the server which receives the incoming TLS connection and must be kept secret. The server introduces itself to the client by handing over its certificate. The certificate is a signed (“certified”) container that includes the server’s public key.

In Cisco Prime Network Registrar 11.0, the DNS server listens on configurable port 853 for TLS. On port 853, only TCP TLS connections are accepted and other connections are dropped. The DNS server has

configurable parameters to enable or disable TLS, and to add TLS private and public key files, and TLS certificate bundle for upstream.

Caching DNS exceptions and forwarders have configuration parameters to enable or disable TLS for upstream.



Note

- Cisco Prime Network Registrar does not support a command for generating self-signed certificates. However, they can be generated using readily available command line tool like openssl. For example:

```
# openssl req -new -x509 -days 365 -nodes -out public.pem -keyout private.pem
```

- TLS is not supported in hybrid mode and in zone transfers.
- TLS keys are not supported with password phrase.

Adding Public Key to the Certificate Authority Bundle

For upstream queries, copy the public.pem of forwarder/exception servers to the Caching DNS server and update the same in *tls-upstream-cert-bundle* using the following commands:

```
scp -r public.pem @client-ip:/etc/pki/ca-trust/source/anchors/
# update-ca-trust
```

The above command will update the */etc/pki/tls/certs/ca-bundle.crt* file.

Copy the updated */etc/pki/tls/certs/ca-bundle.crt* file in the *<cnr.datadir>/cdns/tls* and set this filename in *tls-upstream-cert-bundle*.

Table 13: TLS Attributes in the Caching DNS Server

Attribute	Description
TLS (<i>tls</i>)	<p>Enables or disables TLS support for Caching DNS. Before enabling TLS, the private key files must be placed in the CDNS data directory under <i>cdns/tls</i>, and the <i>tls-service-key</i> attribute be set.</p> <p>If using managed CDNS certificates, the certificate settings will be automatically set. Otherwise, the public certificate file must be placed in the CDNS data directory under <i>cdns/tls</i> and the <i>tls-service-pem</i> attribute be set.</p> <p>Enabling or disabling TLS service requires a Cisco Prime Network Registrar service restart for the change to take effect.</p>
TLS Port (<i>tls-port</i>)	The port number on which to provide TCP TLS service. The Caching DNS server will not serve non-TLS queries on this port.
TLS Private Key File (<i>tls-service-key</i>)	Defines the file name which contains the private key to be used by DNS for TLS sessions. The file must be in the CDNS data directory under the tls subdirectory (that is, <i><cnr.datadir>/cdns/tls</i>). The openssl tool can be used to create TLS private and public key files.

Attribute	Description
TLS Public Key File (<i>tls-service-pem</i>)	Defines the pem file name which contains the public key certificate to be used by CDNS for TLS sessions. The file must be in the CDNS data directory under the tls subdirectory (that is, <cnr.datadir>/cdns/tls). Note that if using managed CDNS certificates, this attribute will be ignored and should be left unset.
TLS Certificate Bundle File (<i>tls-upstream-cert-bundle</i>)	Defines the file name which contains the certificate bundle. These certificates are used for TLS connections made to outside peers. These certificates are used to authenticate connections made to upstream DNS servers. The file must be in the CDNS data directory under the tls subdirectory (that is, <cnr.datadir>/cdns/tls). You can copy the /etc/pki/tls/certs/ca-bundle.crt file or create a soft link for it.

You can also enable TLS at the forwarder (see [Using Forwarders, on page 47](#)), exception (see [Using Exceptions, on page 49](#)), and at the firewall (see [Enabling TLS for RPZ, on page 121](#)) level.

Local Advanced Web UI

To enable TLS support for the Caching DNS server, do the following:

Before you begin

Before enabling TLS, you must place the public certificate and private key files in the CDNS data directory under the **tls** subdirectory (that is, <cnr.datadir>/cdns/tls), and set the *tls-service-key* and *tls-service-pem* attributes under the **TLS Settings** section on the Manage DNS Caching Server page. You can also use the managed certificates (see the "Certificate Management" section in *Cisco Prime Network Registrar 11.0 Administration Guide*).

-
- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page. Click **CDNS** on the Manage Servers pane.
- Step 2** On the Edit Local CDNS Server tab, under the **TLS Settings** section, enable the *TLS* attribute by selecting the **enabled** option.
- Step 3** Click **Save** to save the changes.
-



Note You must restart the Cisco Prime Network Registrar service whenever TLS settings are modified.

CLI Commands

Use the following commands to enable TLS support for the Caching DNS server:

```
nrcmd> cdns enable tls
```

Then, restart the Cisco Prime Network Registrar service using the following command:

```
# systemctl restart nwreglocal.service
```

Use `cdns set attribute=value` to set the TLS attributes in the Caching DNS server.



Note You must restart the Cisco Prime Network Registrar service whenever TLS settings are modified.

TLS Statistics

On the Manage DNS Caching Server page, click the **Statistics** tab to view the Server Statistics page. The *queries-over-tls* attribute appears under the **Query Details** section of both the Total Statistics and Sample Statistics categories. The *tls-errors-in* and *tls-errors-out* attributes appear under the **Server Statistics** section of both the Total Statistics and Sample Statistics categories.

Table 14: TLS Statistics Attributes

Attribute	Description
<i>queries-over-tls</i>	Total number of queries received over TLS by the CDNS Server.
<i>tls-errors-in</i>	Total number of TLS related errors on inbound DNS query attempts. An error may occur whether a query was successfully received or not.
<i>tls-errors-out</i>	Total number of TLS related errors on outbound DNS query attempts. An error may occur whether a query was successfully transmitted or not.

Setting Prefetch Timing

Use the *Prefetch* attribute to set whether message cache elements should be prefetched before they expire to keep the cache up to date. Turning it **on** gives about 10 percent more traffic and load on the machine, but can increase the query performance for popular DNS names.

When *Prefetch* is enabled, records are assigned a prefetch time that is within 10 percent of the expiration time. As the server processes client queries and looks up the records, it checks the prefetch time. Once the record is within 10 percent of its expiration, the server will issue a query for the record to keep it from expiring.

Setting Cache TTLs

Time to Live (TTL) is the amount of time that a DNS server is allowed to cache data learned from other nameservers. Each record added to the cache arrives with some TTL value. When the TTL period expires, the server must discard the cached data and get new data from the authoritative nameservers the next time it sends a query. TTL attributes, *cache-min-ttl* and *cache-max-ttl* defines the minimum and maximum time Cisco Prime Network Registrar retains the cached information. These parameters limit the lifetime of records in the cache whose TTL values are very large or very small.

Local Basic or Advanced Web UI

Step 1 On the Edit Local CDNS Server tab, you can find:

- The Maximum Cache TTL (*cache-max-ttl*) attribute, set it to the desired value (the default value is 24 hours)
- The Min Cache TTL (*cache-min-ttl*) attribute, set it to the desired value (the preset value is 0)

Step 2 Click **Save** to save the changes.

CLI Commands

Use **cdns set cache-max-ttl=***value* to set the maximum Cache TTL value.

Use **cdns set cache-min-ttl =***value* to set the minimum Cache TTL value.

Enabling Smart Caching

Whenever Authoritative DNS servers face an outage or are offline for other reasons, this could cause issues with being able to reach Internet services that are likely not impacted. Smart caching allows the Caching DNS server to continue to serve the expired data (last known answer) when it cannot reach the authoritative name servers. The Caching DNS server will still continue to contact the authoritative name servers and when the name servers are once again functional, the Caching DNS server will update its cached data.



Note Enabling Smart Cache (*smart-cache*) automatically enables prefetch.

Smart Cache Configuration Settings

In Cisco Prime Network Registrar, Caching DNS Smart Cache is not enabled by default. To use Smart Cache, the *smart-cache* attribute must be enabled at the Caching DNS server level.

When the Caching DNS server receives a query for data that has expired and if the *smart-cache* attribute is enabled, it will continue to respond with its expired cached data and increment the *smart-cache* counter under the **Query Details** section in the Statistics tab.



Note Smart Cache is available in Advanced mode and requires a Caching DNS server reload to take effect.

Table 15: Smart Cache Attributes

Attribute	Description
Smart Cache (<i>smart-cache</i>)	Specifies if the Caching DNS server should use Smart Caching. When <i>smart-cache</i> is enabled, the Caching DNS server continues to use its last best known answer when cached responses have expired and it cannot reach the authoritative name servers. The RRs in smart cache responses will have a 0 TTL. Smart Caching is useful to mitigate network outages and possible DDoS attacks that make the authoritative name servers unavailable. Enabling <i>smart-cache</i> automatically enables prefetch.
Smart Cache Expiration (<i>smart-cache-expiration</i>)	When <i>smart-cache</i> is enabled, specifies a time limit for responding with expired RRs. The default is 0, which allows the server to respond with expired answers as long as they remain in the cache.

Attribute	Description
Smart Cache Expiration Reset (<i>smart-cache-expiration-reset</i>)	When <i>smart-cache</i> is enabled and <i>smart-cache-expiration</i> is greater than 0, will reset the expiration time on active queries. This allows active queries to return expired answers, while allowing others to return SERVFAIL responses for a short period. Default is disabled.
Prefetch (<i>prefetch</i>)	<p>Sets whether message cache elements should be prefetched before they expire to keep the cache up to date. Turning it on gives about 10 percent more traffic and load on the machine, but popular items do not expire from the cache.</p> <p>When <i>Prefetch</i> is enabled, records are assigned a prefetch time that is within 10 percent of the expiration time. As the server processes client queries and looks up the records, it checks the prefetch time. Once the record is within 10 percent of its expiration, the server will issue a query for the record in order to keep it from expiring.</p>



Note From Cisco Prime Network Registrar 10.1, the *Prefetch* attribute is available under the **Smart Cache** section and it is an Advanced mode feature.

Local Advanced Web UI

To enable Smart Cache, do the following:

- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page. Click **CDNS** on the Manage Servers pane.
- Step 2** On the Edit Local CDNS Server tab, under the **Smart Cache** section, enable the *smart-cache* attribute by selecting the **enabled** option.
- Step 3** Click **Save** to save the changes.

CLI Commands

Use **cdns enable smart-cache** to enable Smart Caching.

Use **cdns set smart-cache-expiration=value** to specify a time limit for responding with expired RRs, when *smart-cache* is enabled. For example:

```
nrcmd> cdns set smart-cache-expiration=5m
```

Use **cdns enable smart-cache-expiration-reset** to reset the expiration time on active queries, when *smart-cache* is enabled and *smart-cache-expiration* is greater than 0.

Defining Root Nameservers

Root nameservers know the addresses of the authoritative nameservers for all the top-level domains. When you first start a newly installed Cisco Prime Network Registrar Caching DNS server, it uses a set of preconfigured root servers, called root hints, as authorities to ask for the current root nameservers.

When Cisco Prime Network Registrar gets a response to a root server query, it caches it and refers to the root hint list. When the cache expires, the server repeats the process. The TTL on the official root server records is preconfigured and you can specify a different cache TTL value (see [Setting Cache TTLs, on page 37](#)).

As the configured servers are only hints, they do not need to be a complete set. You should periodically (every month to six months) look up the root servers to see if the information needs to be altered or augmented.

Local Basic or Advanced Web UI

On the Edit Local CDNS Server tab, under the **Root Name Servers** section, enter the domain name and IP address of each additional root nameserver, clicking **Add Root Nameserver** after each one, then click **Save**.

CLI Commands

Use **cdns addRootHint** *name addr [addr ...]* to add the name of a root server and the root name server address(es).

Dynamic Allocation of UDP Ports

The Caching DNS server uses a large number of UDP port numbers, by default up to 48000 port numbers. These numbers are divided among the processing threads. The large number of port numbers reduce the risk of cache poisoning via Birthday Attacks. The Caching DNS server uses the default pool of UDP ports (2048) and the maximum allowable size of the default pool of UDP ports is 4096.

Currently, Cisco Prime Network Registrar uses the port range from 1024 to 65535. Based on the number of outstanding resolution queries, the Caching DNS server adjusts the pool size by adding or removing ports. The Caching DNS server allocates and releases the UDP ports dynamically when the server is running. If you reload the server, all the UDP ports are released and randomly picked again.

Setting Maximum Memory Cache Sizes

The maximum memory cache size property specifies how much memory space you want to reserve for the DNS in-memory cache. The larger the memory cache, the less frequently the Caching DNS server will need to re-resolve unexpired records.

Local Advanced Web UI

On the Edit Local CDNS Server tab, under the **Caching** section, set the desired value for the RRSets Cache Size (*rrset-cache-size*) attribute, then click **Save**. The default size is 1 GB.

To set the size of the message cache, use the Message Cache Size (*msg-cache-size*) attribute. The message cache stores query responses. The default size is 1 GB.

CLI Commands

- Use **cdns set rrset-cache-size** to set RRSets Cache Size.
- Use **cdns set msg-cache-size** to set Message Cache Size.

Specifying Resolver Settings

Glue record(s) is/are A record(s) for name server(s) that cannot be found through normal DNS processing because they are inside the zone they define. When the *harden-glue* attribute is enabled, the Caching DNS

server will ignore glue records that are not within the zone that is queried. The *harden-glue* attribute is on by default.

Domain randomization allows a DNS server to send upstream queries for resolution with a randomly generated query name. A valid name server responds with the query name unchanged and therefore this technique can be used to ensure that the response was valid.

In certain occasions, attacker issues a request and then flood the server with fake responses in an attempt to poison the DNS server's cache with rogue data. Randomizing the case gives the server another level of protection against types of attacks.

Cisco Prime Network Registrar supports randomizing upstream queries, but there are some name servers that do not maintain the randomized case. Therefore, if you enable case randomization, you may block out valid name servers. The *randomize-query-case-exclusion* attribute allows you to create an exclusion list, so that you can continue to use case randomization, but exclude name servers that do not maintain the case but still respond with a valid answer.

Table 16: Resolver Settings Attributes

Attribute	Description
<i>harden-glue</i>	Specified if glue should only be trusted if it is within the servers authority.
<i>randomize-query-case</i>	Enables the use of 0x20-encoded random bits in the query to foil spoof attempts. This perturbs the lowercase and uppercase of query names sent to authority servers and checks if the reply still has the correct casing.
<i>randomize-query-case-exclusion</i>	Allows to create an exclusion list for randomization of upstream queries. This attribute will be used when <i>randomize-query-case</i> is enabled.

Configuring Case Randomization Exclusions

The *randomize-query-case-exclusion* attribute is available under the **Resolver Settings** section on the Manage DNS Caching Server page. The *randomize-query-case* is not enabled by default. To use randomize query case exclusion, the *randomize-query-case* attribute must be enabled at the Caching DNS server level.

Both *randomize-query-case* and *randomize-query-case-exclusion* attributes are available in the web UI in Advanced mode.

Local Advanced Web UI

-
- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page. Click **CDNS** on the Manage Servers pane.
- Step 2** On the Edit Local CDNS Server tab, under the **Resolver Settings** section:
- Enable the *randomize-query-case* attribute by selecting the **enabled** option.
 - In the *randomize-query-case-exclusion* field, enter the list of domains (comma separated) that you want to exclude from case randomization.
- Step 3** Click **Save** to save the changes.
-



Note You must reload the Caching DNS server for the changes to take effect.

CLI Commands

Use **cdns enable randomize-query-case** to enable the case randomization.

Use the **cdns set** and **cdns unset** commands to set or unset *randomize-query-case-exclusion*. For example:

```
nrcmd> cdns set randomize-query-case-exclusion="cisco.com"
nrcmd> cdns set randomize-query-case-exclusion="cisco.com, example.com"
nrcmd> cdns unset randomize-query-case-exclusion
```

Specifying Network Settings

The *listen-ip-version* attribute lets you to choose the IP packets to accept and issue. You can check IPv4, IPv6, or both. The *listen-protocol* attribute lets you to choose the packet protocol to answer and issue. You can check UDP, TCP, or both.



Note The default *listen-ip-version* is both IPv4 and IPv6. You can change this to IPv4 if the server you are running does not support IPv6. Otherwise, you will likely experience query timeouts.

Specifying Advanced Settings

The *minimal-responses* attribute controls whether the DNS Caching server omits or includes records from the authority and data sections of query responses when these records are not required. Enabling this attribute may improve query performance such as when the DNS server is configured as a caching server.

The *remote-ns-host-ttl* attribute sets TTL for entries in the remote name server cache. The remote name server cache contains roundtrip timing (RTT), lameness and EDNS support information. Once an entry expires, it is removed from the remote name server cache and the next time the server is contacted a new entry will be added.

Note that RTT is used to decide which name server to query. If a timeout occurs, the RTT value of that server is doubled. If a server starts to become unresponsive, a probing scheme is applied in which a few queries are selected to probe the IP address. If that fails, the name server is blocked for 15 minutes (*remote-ns-host-ttl*) and re-probed with one query after that. Therefore, it may be necessary to decrease the *remote-ns-host-ttl* to allow probing more frequently. The remote name server cache is not flushed after a CDNS server reload, but can be flushed using the **cdns execute flush-ns-cache** command.

The *remote-ns-cache-numhosts* attribute lets you to set the number of hosts for which information is cached.

Enabling Round-Robin

A query might return multiple A or AAAA records for a name lookup. To compensate for most DNS clients starting with, and limiting their use to, the first record in the list, *round-robin* is enabled to share the load.

This ensures that successive clients resolving the same name will connect to different addresses on a revolving basis. The DNS server then rearranges the order of the records each time it is queried. It is a method of load sharing, rather than load balancing, which is based on the actual load on the server.

Local Advanced Web UI

On the Edit Local CDNS Server tab, under the **Advanced Settings** section, find the *round-robin* attribute.

CLI Commands

Use **cdns get round-robin** to see if round-robin is enabled (it is by default). If not, use **cdns enable round-robin**.

Flushing Caching DNS Cache

Cisco Prime Network Registrar cache flushing function lets you remove all or a portion of cached data in the memory cache of the server.

Local Basic or Advanced Web UI

Step 1 From the **Deploy** menu, choose **CDNS Server** under the **DNS** submenu to open the Manage DNS Caching Server page.

Step 2 On the Manage DNS Caching Server page, click the **Commands** button to open the CDNS Command dialog box. There will be two types of cache flushing commands.

- **Flush the CDNS cache**—Allows you to either flush all cache entries for a particular zone or the entire cache if no zone is provided. To remove all data for a specific zone, enter the zone name in the Zone field. To clear the whole cache, leave the Zone field empty.
- **Flush Resource Record**—Allows you to flush an RR name or an RRSet when the type field is specified.
 - Remove common RR types (A, AAAA, NS, SOA, CNAME, DNAME, MX, PTR, SRV, NAPTR, and TXT) from a specific domain—Enter the required RR name as the FQDN for the Flush Resource Record command and leave the RR type field empty.
 - Remove a specified RR type for a domain—Specify the domain in the FQDN field, and the RR type in the RR type field.

Note When no type is specified, the server flushes types A, AAAA, NS, SOA, CNAME, DNAME, MX, PTR, SRV, TXT, and NAPTR.

CLI Commands

- Use the following command to remove all cached entries at or below a given domain. If no domain is given, it flushes all RRs in the cache.

```
nrcmd> cdns flushCache domain
```

- Use the following command to flush RRs from the cache associated with the given RR name. When type is provided, it flushes all entries with the given name and type. If no type is provided, it flushes types A, AAAA, NS, SOA, CNAME, DNAME, MX, PTR, SRV, TXT, and NAPTR.

```
nrcmd> cdns flushName name type
```

Detecting and Preventing DNS Cache Poisoning

Cisco Prime Network Registrar enhances the Caching DNS server performance to address the CDNS related issues such as DNS cache poisoning attacks (CSCsq01298), as addressed in a Cisco Product Security Incident

Response Team (PSIRT) document number PSIRT-107064 with Advisory ID cisco-sa-20080708-dns, available at:

<http://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20080708-dns>

DNS Cache Poisoning Attacks

A cache poisoning attack can change an existing entry in the DNS cache as well as insert a new invalid record into the DNS cache. This attack causes a hostname to point to the wrong IP address. For example, let us say that `www.example.com` is mapped to the IP address `192.168.0.1`, and this mapping is present in the cache of a DNS server. An attacker can poison the DNS cache and map `www.example.com` to `10.0.0.1`. If this happens, if you try to visit `www.example.com`, you will end up contacting the wrong web server.

A DNS server that uses a single static port for receiving responses to forwarded queries are susceptible to malicious clients sending forged responses.

The DNS transaction ID and source port number used to validate DNS responses are not sufficiently randomized and can easily be predicted, which allows an attacker to create forged responses to DNS queries. The DNS server will consider such responses as valid.

Handling DNS Cache Poisoning Attacks

To reduce the susceptibility to the DNS cache poisoning attack, the DNS server randomizes the UDP source ports used for forwarded queries. Also, a resolver implementation must match responses to the following attributes of the query:

- Remote address
- Local address
- Query port
- Query ID
- Question name (not case-sensitive)
- Question class and type, before applying DNS trustworthiness rules (see [RFC2181], section 5.4.1)



Note

The response source IP address must match the query's destination IP address and the response destination IP address must match the query's source IP address. A mismatch must be considered as format error, and the response is invalid.

Resolver implementations must:

- Use an unpredictable source port for outgoing queries from a range (either 53, or > 1024) of available ports that is as large as possible and practicable.
- Use multiple different source ports simultaneously in case of multiple outstanding queries.
- Use an unpredictable query ID for outgoing queries, utilizing the full range available (0 to 65535). By default, CDNS uses up to 48000 port numbers.

The Caching DNS server attribute *randomize-query-case*, when enabled, specifies that when sending a recursive query, the query name is pseudo-randomly camel-cased and the response is checked to see if this camel-casing is unchanged. If *randomize-query-case* is enabled and the casing has changed, then the response is discarded. The *randomize-query-case* is disabled by default, disabling this feature.

Local Basic or Advanced Web UI

The Caching DNS server statistics appears on the Statistics tab of the Manage DNS Caching Server page. The Statistics displays the *answers-unwanted* values. You can refresh the DNS Caching Server Statistics by clicking the **Refresh Server Statistics** icon at the top of the statistics table.

Handling Unresponsive Nameservers

When trying to resolve query requests, Caching DNS servers may encounter unresponsive nameservers. A nameserver may be unresponsive to queries or respond late. This affects the performance of the local DNS server and remote nameservers.

Using Cisco Prime Network Registrar, you can resolve these problems by barring unresponsive nameservers. You can configure a global ACL of unresponsive nameservers that are to be barred, using the *acl-do-not-query* attribute.

When Cisco Prime Network Registrar receives a list of remote nameservers to transmit a DNS query request to, it checks for the nameservers listed in the *acl-do-not-query* list and removes them from this list. Conversely, all incoming DNS requests from clients or other nameservers are also filtered against *acl-blocklist*.

Use the *acl-query* attribute to specify which clients are allowed to query the server. By default, any client is allowed to query the server. A client that is not in this list will receive a reply with status REFUSED. Clients on the *acl-blocklist* do not get any response whatsoever.

Local Advanced Web UI

On the Edit Local CDNS Server tab, expand the **Query Access Control** section to view the various attributes and their values. For the Do Not Query (*acl-do-not-query*) attribute, enter the value (for example, 10.77.240.73). Then, click **Save**.

Running DNS Caching Server Commands

Access the DNS Caching server commands using the Commands button. Clicking the **Commands** button opens the CDNS Commands dialog box in the local web UI. Each command has its own Run icon (click it, then close the dialog box):

- **Flush the CDNS cache**— This command allows you to flush either all RRs or RRs for a particular zone from the in-memory cache. See [Flushing Caching DNS Cache, on page 43](#).
- **Flush Resource Record**— This command that lets you specify an RR name and optionally a type to remove from the in-memory cache.



Note To remove all the entries from the in-memory cache, you need to reload the Caching DNS server.



Note If you find a server error, investigate the server log file for a configuration error, correct the error, return to this page, and refresh the page.

Configuring Caching DNS Server Network Interfaces

You can configure the network interfaces for the Caching DNS server from the Manage Servers page in the local web UI. If no interfaces are explicitly configured, the server uses all the available interfaces.

Local Advanced Web UI

- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page.
 - Step 2** Select **CDNS** from the Manage Servers pane.
 - Step 3** Click the **Network Interfaces** tab to view the available network interfaces that you can configure for the server. By default, the server uses all of them.
 - Step 4** To configure an interface, click the **Configure** icon in the Configure column for the interface. This adds the interface to the Configured Interfaces table, where you can edit or delete it.
 - Step 5** Click the name of the configured interface to edit the configured interfaces, where you can change the address, direction and port of the interface.
 - Step 6** Click **Modify Interface** when you are done editing, then click **Go to Server Interfaces** to return to the Network Interfaces page.
-



CHAPTER 4

Advanced Caching DNS Server

This chapter explains how to set the Caching DNS parameters for the advanced features of the server. Before you proceed with the tasks in this chapter, see [Introduction to the Domain Name System, on page 1](#), which explains the basics of DNS.

- [Using Forwarders, on page 47](#)
- [Using Exceptions, on page 49](#)
- [Managing DNS64, on page 51](#)
- [Managing DNSSEC, on page 52](#)
- [Managing Caching Rate Limiting, on page 53](#)
- [Managing DNS Views, on page 56](#)
- [Setting up Caching DNS and Authoritative DNS Servers on the Same Operating System, on page 56](#)
- [Managing DNS Firewall, on page 56](#)
- [Configuring Caching DNS to Use Umbrella, on page 57](#)

Using Forwarders

You can specify a domain for which forwarding should occur. The forwarder definition is a list of IP addresses with an optional port number or a list of names of servers, or both. Typically forwarders are other DNS Caching servers that have access to Internet or external DNS resources.



Note We highly recommend using IP address rather than hostnames.

When forwarders are used, the Caching DNS server forwards user queries matching the forwarding domain to another Caching DNS server to perform the resolution. This can be useful in situations where the local Caching DNS server does not have Internet access (that is, inside a firewall). In these situations, it is typical for exceptions to be configured for local zones and then a root (.) forwarder to be created for all external queries. Forwarder name corresponds to the domains you would like to have forwarded. For example, to forward example.com queries, your forwarder will be named example.com.



Note You can specify IPv4 and/or IPv6 addresses and for the changes to take effect, you must reload the Caching DNS server.



Tip To force the Caching DNS server to forward all queries to one or more DNS forwarders, use the DNS root (.) as the forwarder name.



Note Caching DNS by default does not allow access to AS112 and RFC 1918 reverse zones. These are the reverse zones for IP address ranges that are reserved for local use only. To access these zones, define an exception or forwarder for the reverse zones that are defined locally.

In Cisco Prime Network Registrar 11.0, you can enable TLS at the individual forwarder object level. To do this, enable the *tls* attribute by selecting the **enabled** option. If you enable this, you should configure a *tls-cert-bundle* to load the CA certificates, otherwise, the connections cannot be authenticated. To add public key to the Certificate Authority bundle, copy the *public.pem* of forwarder server to the Caching DNS server, and update the same in *tls-upstream-cert-bundle* using the following commands:

```
scp -r public.pem @client-ip:/etc/pki/ca-trust/source/anchors/
# update-ca-trust
```

The *tls-auth-name* indicates the auth name for the forwarder server. If TLS is enabled, the Caching DNS server checks the TLS authentication certificates with that name sent by the forwarder server.

Local and Regional Web UI

To define a forwarder:

- Step 1** From the **Design** menu, choose **Forwarders** under the **Cache DNS** submenu. This opens the List/Add Forwarders page.
- Step 2** Click the **Add Forwarders** icon on the **Forwarders** pane to open the Add Forwarder dialog box.
- Step 3** Enter the name of the zone to be forwarded as the name and click **Add Forwarder**.
 - Note** To use a forwarder for all external queries, create a forwarder with the name ".".
- Step 4** In the Edit Forwarders page, enter the hostname, and click **Add Host** or enter the IP address for the forwarder, and then click **Add Address**.
- Step 5** Click **Save**.

CLI Commands

- To specify the address (or space-separated addresses) of nameservers to use as forwarders, use **cdns addForwarder** *domain* [**tls=on** | **off**] [**tls-auth-name=name**] *addr*.
If the **tls** flag is on, the server connects to the name server using TLS. If **tls-auth-name** is provided, the server verifies this name in the TLS certificate provided by the name server.
You can also use **cdns-forwarder** *name* **create** *attribute=value* to create the Caching DNS forwarder objects.
- To list the current forwarders, use **cdns listForwarders** or **cdns-forwarder list**.

- To modify the forwarder objects, use **cdns-forwarder name set attribute=value**.
- To remove a forwarder or list of forwarders, use **cdns removeForwarder domain [addr ...]** or **cdns-forwarder name delete**.



Note For any TLS related changes in the forwarders to take effect, you should restart the Caching DNS server.

Using Exceptions

If you do not want the Caching DNS server to use the standard resolution method to query the nameserver for certain domains, use exceptions. This bypasses the root nameservers and targets a specific server (or list of servers) to handle name resolution. Typically exceptions are used to access local DNS authoritative resources (that is, a company's internal zones).

Let us say that example.com has two subsidiaries: Red and Blue. Each has its own domain under the .com domain. When users at Red want to access resources at Blue, their Caching DNS server follows delegations starting at the root nameservers.

These queries cause unnecessary traffic, and in some cases fail because internal resources are often barred from external queries or sites that use unreachable private networks without unique addresses.

Exceptions solve this problem. The Red administrator can list all the other example.com domains that users might want to reach and at least one corresponding nameserver. When a Red user wants to reach a Blue server, the Red server queries the Blue server instead following delegations from the root servers down.

To enable resolution exceptions, simply create an exception for the domain listing the IP address(es) and/or hostname(s) of the authoritative nameserver(s).



Note Exceptions can contain both IPv4 and/or IPv6 addresses, and require a Caching DNS server reload to take effect.



Warning If the Authoritative DNS server is using a non-standard DNS port (a port other than 53) and if the exception zone has subzones, then the user must configure separate exceptions for each subzone that refers to the non-standard port. Otherwise, the Caching DNS server defaults to using port 53 for the subzones, leading to resolution failures.

In Cisco Prime Network Registrar 11.0, you can enable TLS at the individual exception object level. To do this, enable the *tls* attribute by selecting the **enabled** option. If you enable this, you should configure a *tls-cert-bundle* to load the CA certificates, otherwise, the connections cannot be authenticated. To add public key to the Certificate Authority bundle, copy the public.pem of exception server to the Caching DNS server, and update the same in *tls-upstream-cert-bundle* using the following commands:

```
scp -r public.pem @client-ip:/etc/pki/ca-trust/source/anchors/
# update-ca-trust
```

The *tls-auth-name* attribute indicates the auth name for the exception server. If TLS is enabled, the Caching DNS server checks the TLS authentication certificates with that name sent by the exception server.

Local and Regional Web UI

-
- Step 1** From the **Design** menu, choose **Exceptions** under the **Cache DNS** submenu. This opens the List/Add Exceptions page.
 - Step 2** Click the **Add Exceptions** icon in the **Exceptions** pane to open the Add Exception dialog box.
 - Step 3** In the Name field, enter the domain or zone for which an exception is wanted and click **Add Exception**.
 - Step 4** In the Edit Exceptions page, enter the hostname in the DNS Name field and click **Add Host**. To address, enter the IP address in the IP Address field and click **Add Address**.
 - Step 5** If the *prime* attribute is on, Caching DNS server queries the zone for the currently published name servers and use those. This is similar to how the server treats root hints.
 - Step 6** Click **Save**.
-

To delete an exception list, select the exception in the Exceptions pane and click the **Delete** icon. To add or remove name servers to an exception, click the name of the exception in the List/Add Exceptions page to open the Edit Exceptions page.

CLI Commands

Use the exception commands only if you do not want your Caching DNS server to use the standard name resolution for querying root name servers for names outside the domain. Cisco Prime Network Registrar sends non-recursive queries to these servers.

- To add the resolution exception domains and IP addresses of servers, separated by spaces, use **cdns addException domain [prime=on | off] [tls=on | off] [tls-auth-name=name] [views=on | off] [addr ...]**. The addresses can be IPv4 or IPv6 with an optional port number (that is, *addr[@port]*) or the name of a server (it must be possible to resolve the server name before it is used). Use this command only if you do not want your Caching DNS server to use the standard name resolution for a zone.

If the **tls** flag is on, the server connects to the name server using TLS. If **tls-auth-name** is provided, the server verifies this name in the TLS certificate provided by the name server.

You can also use **cdns-exception name create attribute=value** to create the Caching DNS exception objects.

- To list the domains that are configured to have exceptional resolution of their names, use **cdns listExceptions** or **cdns-exception list**.
- To remove an entry for exceptional resolution of addresses within a domain, use **cdns removeException domain [addr ...]** or **cdns-exception name delete**. You can remove an individual server by specifying it, or the exception itself by just specifying its name.
- To modify the exception objects, use **cdns-exception name set attribute=value**.



Note For any TLS related changes in the exceptions to take effect, you should restart the Caching DNS server.

Managing DNS64

DNS64 with NAT64 provides access to the IPv4 Internet and servers for hosts that have only IPv6 addresses. DNS64 synthesizes AAAA records from A records, when an IPv6 client queries for AAAA records, but none are found. It also handles reverse queries for the NAT64 prefix(es).

In Cisco Prime Network Registrar, you can define multiple prefixes for synthesizing AAAA record.



Note

- When you enable DNS64 on multiple Caching DNS servers, you must ensure that the same version of Cisco Prime Network Registrar is installed on all the Caching DNS servers.
- If DNS firewall redirect is also enabled, the Caching DNS redirect takes precedence over DNS64 functionality.
- If DNS64 is enabled, enabling DNSSEC is not recommended. DNS64 causes responses to be simulated which may cause DNSSEC validation to fail.
- For DNS64 to be useful, there must be a corresponding NAT64 service on the network.

Local and Regional Advanced Web UI

To add, edit, or view the DNS64 configuration items:

- Step 1** From the **Design** menu, choose **DNS64** under the **Cache DNS** submenu to open the List/Add DNS64 page.
- Step 2** Click the **Add DNS64** icon in the DNS64 pane to open the Add DNS64 dialog box.
- Step 3** Enter the name for the DNS64 configuration item in the Name field.
- Step 4** Click **Add Dns64** to add the configuration item. The Edit DNS64 page appears with the list of attributes that can be edited.
- Step 5** Edit the values of the attributes, as required. The value defined for *Priority* decides the search order for the client's DNS64 configuration.
- Step 6** Click **Save** to save your settings for the selected DNS64 configuration item.

To delete a DNS64 configuration item, select the DNS64 entry on the DNS64 pane, click the **Delete DNS64** icon, and then confirm the deletion.

CLI Commands

To create DNS64 in the Caching DNS server, use the **cdns64 name create [acl-match-clients=ACL prefix=IPv6 prefix]** command (see the **cdns64** command in the CLIGuide.html file in the /docs directory for syntax and attribute descriptions or use **help cdns64** in the CLI). For Example:

```
nrcmd> cdns64 dns64 create
nrcmd> cdns64 dns64 set acl-match-clients=baaa::56ff:febd:3d6
```

When connected to a regional cluster, you can use the following pull, push, and reclaim commands. For push and reclaim, a list of clusters or "all" may be specified.

- `cdns64 <name | all > pull < ensure | replace | exact > cluster-name [-report-only | -report]`
- `cdns64 <name | all > push < ensure | replace | exact > cluster-list [-report-only | -report]`
- `cdns64 name reclaim cluster-list [-report-only | -report]`

Managing DNSSEC

DNS Security Extensions (DNSSEC) enables the server to determine the security status of all Resource Records that are retrieved. You can manage DNSSEC in the Advanced and Expert modes. The `dnssec` attribute enables validation of DNS information. The `domain-insecure` attribute defines domain names to be insecure, DNSSEC chain of trust is ignored towards the domain names. So, a trust anchor above the domain name can not make the domain secure with a DS record, such a DS record is then ignored. DNSSEC requires a root trust anchor to establish trust for the DNS root servers. The initial DNSSEC root trust anchor, `root.anchor`, is stored in the `.../data/cdns` directory and is the default value of the `auto-trust-anchor-file` attribute. Additional trust anchors may be added by adding them to the `.../data/cdns` directory and to the `auto-trust-anchor-file` if the zone supports automated updates according to RFC 5011 or the `trust-anchor-file` attribute if not. The `cdnssec` command controls and configures DNSSEC processing in the Cisco Prime Network Registrar Caching DNS server.

To set the size of the aggressive negative cache in bytes, use the `neg-cache-size` attribute on the Manage DNS Caching Server page.

The `key-cache-size` attribute sets the size of the key cache in bytes. The `prefetch-key` attribute sets whether the Caching DNS server should fetch the DNSKEYs earlier in the validation process, when a DS record is encountered.



Note If DNS64 is enabled, enabling DNSSEC is not recommended. DNS64 causes responses to be simulated which may cause DNSSEC validation to fail.

Local Advanced Web UI

-
- Step 1** From the **Design** menu, choose **Caching DNSSEC** under the **Security** submenu to open the Manage Caching DNSSEC page.
 - Step 2** Enable DNSSEC validation by selecting the **enabled** option for the Enable DNSSEC validation (`dnssec`) attribute.
 - Step 3** The page displays all the Caching DNSSEC attributes. Modify the attributes as per your requirements.
 - Step 4** Click **Save** to save your settings.
-

CLI Commands

- To create DNSSEC in the Caching DNS server, use `cdnssec create attribute=value`. To enable DNSSEC, use `cdnssec enable dnssec` (see the `cdnssec` command in the `CLIGuide.html` file in the `/docs` directory for syntax and attribute descriptions or use `help cdnssec` in the CLI).
- Use `cdns set neg-cache-size` to set Negative Cache Size.

Managing Caching Rate Limiting

Rate limiting helps the DNS server from being overwhelmed by a small number of clients. It also protects against upstream query attacks against Authoritative DNS servers. The rate limiting feature helps to mitigate some of the DDoS attacks and prevents the server from being overwhelmed by a small number of clients. This feature allows you to limit the malevolent traffic.

You can manage rate limiting in Advanced mode in the local web UI. Rate limiting is divided into two separate categories, Client Rate Limiting and Domain Rate Limiting, which are managed separately.

Client Rate Limiting

Client Rate Limiting imposes limits on the QPS per client and when the limit is reached, new queries are dropped. When a client is rate limited, it is possible to still allow some queries through.

The *client-rate-limiting* attribute on the Rate Limiting Settings tab enables IP based client rate limiting. It is not enabled by default. The *client-rate-limit-qps* attribute specifies the maximum QPS for an incoming client IP before starting the rate limiting. Default value is 1000. The *client-rate-limiting-factor* specifies that one out of this many number of queries will be allowed through when a client IP is being rate limited. For information about all the Client Rate Limiting attributes, see [Table 17: Client Rate Limiting Attributes](#) below.

The Client Rate Limiting tab on the Manage Caching Rate Limiting page displays information about the current clients being rate limited and the limits they are hitting. The table on the page shows:

- **Client**—Rate limited client IP addresses.
- **Number of times rate limited**—Total number of times a client was rate limited.

Table 17: Client Rate Limiting Attributes

Attribute	Description
Client Rate Limiting (<i>client-rate-limiting</i>)	Enables IP based client rate limiting.
Client Rate Limiting QPS (<i>client-rate-limiting-qps</i>)	Specifies the rate limit for incoming DNS clients.
Client Rate Limiting Factor (<i>client-rate-limiting-factor</i>)	When <i>client-rate-limiting</i> is enabled and a client is being rate limited, specifies that one out of this number of queries from that client will be allowed to complete.
Client Report Max (<i>client-report-max-count</i>)	Specifies the maximum number of entries in the list of rate limited clients. This limit is applied to the lists of clients that are logged, returned as part of activity summary or included in statistics.

Domain Rate Limiting

Domain Rate Limiting imposes limits on the QPS the server may send to authoritative name server for a DNS zone. When a domain is rate limited, it is possible to still allow some queries through.

The *domain-rate-limiting* attribute on the Rate Limiting Settings tab enables domain based (name server zones) rate limiting. It is not enabled by default. The *domain-rate-limit-qps* specifies the maximum QPS for a domain/zone before starting the rate limiting. The default value is 1000. The *domain-rate-limiting-factor* specifies that one out of this many number of queries to the specified zone will be allowed through, when the zone is being rate limited. For information about all the Domain Rate Limiting attributes, see [Table 18: Domain Rate Limiting Attributes](#) below.

The Domain Rate Limiting tab on the Manage Caching Rate Limiting page displays information about the current domains being rate limited and the limits they are hitting. The table on the page shows:

- **Domain**—Rate limited domains.
- **Rate Limit Max QPS**—Maximum number of entries in the list of rate limited domains.
- **Number of times rate limited**—Total number of times a domain was rate limited.

Table 18: Domain Rate Limiting Attributes

Attribute	Description
Domain Rate Limiting (<i>domain-rate-limiting</i>)	Enables rate limiting for name server zones.
Domain Rate Limiting QPS (<i>domain-rate-limiting-qps</i>)	Specifies the rate limit for name server zones.
Domain Rate Limiting Factor (<i>domain-rate-limiting-factor</i>)	When <i>domain-rate-limiting</i> is enabled and a zone is being rate limited, specifies that one out of this number of queries to the specified zone will be allowed to complete.
Per Domain Limit	Specifies a list of domains that use a rate limit other than <i>domain-rate-limiting-qps</i> . The list entries have the following attributes: <ul style="list-style-type: none"> • domain—The name of the zone delegation point to which this entry applies. • applies-to—Specifies if this entry applies to only the zone designated by 'domain', only zones specified by subdomains of 'domain', or both. • rate-limit—The rate limit that applies to zones covered by this entry.
Domain Report Max (<i>domain-report-max-count</i>)	Specifies the maximum number of entries in the list of rate limited domains. This limit is applied to the lists of domains that are logged, returned as part of activity summary or included in statistics.

Managing Rate Limiting

You can manage both Client Rate Limiting and Domain Rate Limiting from the Manage Caching Rate Limiting page in the local web UI. This page contains the following three tabs:

- **Rate Limiting Settings**—Displays all the Rate Limiting attributes under their respective categories.
- **Domain Rate Limiting**—Displays a list of domains that are rate limited. This tab also contains information such as Rate Limit Max QPS and number of times a domain was rate limited.

- **Client Rate Limiting**—Displays a list of clients that are rate limited. This tab also contains information about the number of times a client was rate limited.



Note The length of the list is controlled by Client Report Max and Domain Report Max attributes.

Local Advanced Web UI

Step 1 From the **Design** menu, choose **Rate Limiting** under the **Cache DNS** submenu to open the Manage Caching Rate Limiting page.

Step 2 Modify the attributes in the **Client Rate Limiting** and **Domain Rate Limiting** categories as per your requirements:

- To enable Client Rate Limiting, find the *client-rate-limiting* attribute under the **Client Rate Limiting** section, and enable it by selecting the **on** option.
- To enable Domain Rate Limiting, find the *domain-rate-limiting* attribute under the **Domain Rate Limiting** section, and enable it by selecting the **on** option.

Step 3 Click **Save** to save the changes.



Note You must restart the Caching DNS server for these changes to take effect.

Per Domain Limit

You can specify a list of domains to be rate limited with their associated rate limit values. This applies to a domain, its subdomains, or both. These domains use a rate limit other than *domain-rate-limiting-qps*. You can specify a list by adding domains using the **Add** button under the **Per Domain Limit** section.



Note When specifying Per Domain Limit, it is important that the domain names match a DNS zone.

Local Advanced Web UI

On the Rate Limiting Settings tab, under the **Domain Rate Limiting** section, click the **Add** button next to **Per Domain Limit**. In the Add Domain dialog box, enter the domain name (the name of the zone), rate limit value, and specify whether it applies to a domain, its subdomains, or both. Then, click the **Add** button. Click **Save** on the Rate Limiting Settings tab to save the changes.

CLI Commands

- Use **cdns-rate-limit enable client-rate-limiting** to enable the client rate limiting feature.
- Use **cdns-rate-limit set client-rate-limiting-qps=value** to set the QPS value for the client rate limiting. For example:

```
nrcmd> cdns-rate-limit set client-rate-limiting-qps=1000
```

- Use **cdns-rate-limit set domain-rate-limiting-qps=value** to set the QPS value for the domain rate limiting. For example:

```
nrcmd> cdns-rate-limit set domain-rate-limiting-qps=500
```

- Use **cdns-rate-limit add [domain=]<domain> [[applies-to=]domain | subdomain | both] [[rate-limit=]rate-limit]** to specify Rate Limit for the *domain-rate-limiting-list* attribute. For example:

```
nrcmd> cdns-rate-limit add example.com both 1000
```

- Use **cdns-rate-limit list** to display the list of domains that use a rate limit other than *domain-rate-limiting-qps*.
- Use **cdns getStats rate-limit** to get rate limiting statistics.

Managing DNS Views

The Cisco Prime Network Registrar Caching DNS server can associate the client requests to the appropriate views on behalf of the Authoritative DNS server. This is done by configuring the DNS Views on the Caching DNS server and setting the *uses-views* attribute on the List/Add Exceptions page to **true**. The Caching DNS server maps the client to the appropriate view and tags the queries forwarded to the Authoritative DNS server with the appropriate view. Therefore, in these cases, the view mapping is done by the Caching DNS server.



Note The Caching DNS server only maps clients to *acl-match-clients*. The *acl-match-destinations* attribute is ignored.

DNS Views and Exception settings are automatically synced/set by zone distribution.

For more information on DNS Views, see [Managing DNS Views, on page 157](#).

Setting up Caching DNS and Authoritative DNS Servers on the Same Operating System

In Cisco Prime Network Registrar 10.0 and later, both the Caching DNS and Authoritative DNS servers can run on the same operating system, without the need for two separate virtual or physical machines. For more information on DNS firewall, see [Managing DNS Firewall, on page 115](#).

Managing DNS Firewall

Cisco Prime Network Registrar DNS Firewall provides a mechanism to control the domain names, IP addresses, and name servers that are allowed to function on the network. For more information on DNS firewall, see [Managing DNS Firewall, on page 115](#).

Configuring Caching DNS to Use Umbrella

Cisco Umbrella provides the first line of defense against threats on the Internet. To switch to Umbrella from Cisco Prime Network Registrar Caching DNS server, you need to create a forwarder for the “.” domain using the following CLI commands:

```
nrcmd> cdns addForwarder . 208.67.222.222 208.67.220.220
nrcmd> cdns reload
```

Once configured, the Cisco Prime Network Registrar Caching DNS server will forward all resolution queries to Cisco Umbrella (the server will still respond with locally cached answers). It can be used in conjunction with DNS firewall for queries not explicitly blocked by the firewall.



Note Exceptions will operate as usual. Local resolution through exceptions will bypass the Umbrella servers.



Note Umbrella also supports IPv6 addresses 2620:119:35::35 and 2620:119:53::53. For more information, go to umbrella.cisco.com.



CHAPTER 5

Caching DNS Metrics

Following Caching DNS metric elements are available in the dashboard. For the complete list of Caching DNS server statistics, see the *"CDNS Statistics" section under the "Server Statistics" appendix of Cisco Prime Network Registrar 11.0 Administration Guide*.

- [Caching DNS General Indicators, on page 59](#)
- [DNS Caching Activity, on page 60](#)
- [DNS Caching Server Queries Per Second, on page 60](#)
- [DNS Caching Server Recursion Rate Limit, on page 60](#)
- [DNS Incoming Queries, on page 60](#)
- [DNS Queries Responses, on page 61](#)
- [DNS Queries Type, on page 62](#)
- [DNS Recursive Query Time , on page 62](#)

Caching DNS General Indicators

The Caching DNS General Indicators dashboard element shows the server state, its last and startup reload time, and the total resource record (RR) count. The table is available if you choose **CDNS Metrics :Caching DNS General Indicators** in the Chart Selections page.

The resulting table displays:

- **Server State**—Up or Down (based on whether statistics are available), and how long the server has been in this state.
- **Last Reload**—How long since the last server reload.
- **Start Time**—Date and time of the last server process (Cisco Prime Network Registrar server agent) startup.

How to Interpret the Data

The data in this chart shows general server health and operational duration. The objective is to make decisions about the server, such as whether it might be time for another reload, perhaps warranted by the number of configured zones.

Troubleshooting Based on the Results

If the server state is Down, all the CDNS chart indicators show a red status box, so no data will be available. In the case of a server that is down, restart the server.

DNS Caching Activity

The DNS Caching dashboard element rendered as area chart traces the cache hits and cache misses. The chart is available if you choose **CDNS Metrics: DNS Caching Activity** in the Chart Selections page.

The resulting area chart plots the following trends:

- **Cache Hits**—The total number of queries that were answered from cache.
- **Cache Misses**—The total number of queries that were not found in the cache.
- **Prefetches**—Number of prefetches performed.

How to Interpret the Data

This chart indicates the number of queries that were successfully answered using a cache lookup against the number of queries that needed recursive processing.

Troubleshooting Based on the Results

If the cache misses are increasing exponentially, check the CDNS logs for errors. Increasing rates of cache misses can indicate that not enough space is available in memory to store the cached queries for more efficient responses.

DNS Caching Server Queries Per Second

The DNS Caching Server Queries Per Second dashboard element, rendered as chart, displays queries per second for the Caching DNS server. This chart is available if you choose **CDNS Metrics: DNS Caching Server Queries Per Second** in the Chart Selections page.

DNS Caching Server Recursion Rate Limit

The DNS Caching Server Recursion Rate Limit dashboard element, rendered as line chart, shows the number of queries limited for clients and domain. This chart is available if you choose **CDNS Metrics: DNS Caching Server Recursion Rate Limit** in the Chart Selections page.

The resulting line chart plots the following trends:

- **Client Rate Limit**—Number of times a client has been rate limited, when *client-rate-limiting* is enabled.
- **Domain Rate Limit**—Number of times a zone has been rate limited, when *domain-rate-limiting* is enabled.

DNS Incoming Queries

The CDNS Incoming queries by dashboard element rendered as area chart traces the TCP, IPv6, DNSSEC, EDNS and Total queries. The chart is available if you choose **CDNS Metrics: DNS Incoming Queries** in the Chart Selections page.

The resulting area chart plots the following trends:

- **TCP**—Total number of queries received over TCP by the CDNS Server.
- **IPv6**—Total number of queries received over IPv6 by the CDNS Server.
- **EDNS**—Number of queries with EDNS OPT RR present.
- **DNSSEC**—Number of queries with EDNS OPT RR with DO (DNSSEC OK) bit set.
- **Total**—Total number of queries received by the CDNS Server.

How to Interpret the Data

This chart shows the number of queries that were made using TCP, IPv6, and DNSSEC towards the CDNS server, number of queries that had an EDNS OPT record present, and the total number of queries received.

DNS Queries Responses

The CDNS Query Responses dashboard element rendered as area chart shows the number of responses with NOERROR, NODOMAIN, No Data, Other Errors, Secure, and Unsecure return codes. The display is available if you choose **CDNS Metrics: DNS Queries Responses** in the Chart Selections page.

The resulting area chart plots the following trends:

- **NOERROR**—Number of answers from cache or recursion that result in rcode of NOERROR being returned to client.
- **NXDOMAIN**—Number of answers from cache or recursion that result in rcode of NXDOMAIN being returned to client.
- **NODATA**—Number of answers that result in pseudo rcode of NODATA being returned to client.
- **Other Errors**—Other errors.
- **Secure**—Number of answers that were validated correctly by DNSSEC.
- **Unsecure**—Number of answers that failed validation by DNSSEC.

How to Interpret the Data

This chart shows the following:

- The number of answers to queries, from cache or from recursion, that had the return code NXDOMAIN.
- The number of answers to queries that had the pseudo return code NODATA. This means the actual return code was NOERROR, but additionally, no data was carried in the answer (making what is called a NOERROR/NODATA answer). These queries are also included in the NOERROR number. Common for AAAA lookups when an A record exists, and no AAAA.
- Number of answers that were secure. The answer validated correctly. The AD bit might have been set in some of these answers, where the client signalled (with DO or AD bit in the query) that they were ready to accept the AD bit in the answer.
- Number of answers that did not correctly validate.

In a normal scenario, NOERROR is the successful response code.

Troubleshooting Based on the Results

Check the CDNS server configuration if the errors are increasing.

DNS Queries Type

The DNS Queries Type dashboard element rendered as area chart traces the number queries by type. The chart is available if you choose **CDNS Metrics: DNS Queries Type** in the Chart Selections page.

The resulting area chart plots the following trends:

- **A**—Number of A queries received.
- **AAAA**—Number of AAAA queries received.
- **CNAME**—Number of CNAME queries received.

How to Interpret the Data

This chart shows the number of incoming queries of type A, AAAA, CNAME, PTR, and others.

DNS Recursive Query Time

The CDNS Queries by Type dashboard element rendered as area chart traces the average time to complete a recursive query and the median time to complete a query. The table is available if you choose **CDNS Metrics: DNS Recursive Query Time** in the Chart Selections page.

The resulting area chart plots the following trends:

- **Average**—The average time to complete a recursive query.
- **Median**—The median time to complete a recursive query.

How to Interpret the Data

Average indicates the time the server took to answer queries that needed recursive processing. Note that the queries that were answered from the cache are not in this average.

Median time indicates the median of the time the server took to answer the queries that needed recursive processing. The median means that 50% of the user queries were answered in less than this time. Because of big outliers (usually queries to non responsive servers), the average can be bigger than the median.

Troubleshooting Based on the Results

Check the connectivity and configuration for the name servers as forwarders or exception lists for the increasing values of the average and median time.



PART III

Authoritative DNS Server

- [Managing Authoritative DNS Server, on page 65](#)
- [DNS Host Health Check, on page 111](#)
- [Managing DNS Firewall, on page 115](#)
- [Managing High Availability DNS, on page 123](#)
- [Managing Zones, on page 129](#)
- [Managing DNS Views, on page 157](#)
- [Managing Resource Records, on page 163](#)
- [Managing Hosts, on page 175](#)
- [Authoritative DNS Metrics, on page 179](#)



CHAPTER 6

Managing Authoritative DNS Server

This chapter explains how to set the Authoritative DNS server parameters. Before you proceed with the tasks in this chapter, read [Managing Zones, on page 129](#) which explains how to set up the basic properties of a primary and secondary zone.

- [Setting DNS Server Properties, on page 65](#)
- [Running DNS Authoritative Server Commands, on page 98](#)
- [Configuring DNS Server Network Interfaces, on page 99](#)
- [Managing Authoritative DNSSEC, on page 99](#)
- [Managing Authoritative DNSSEC Keys, on page 102](#)
- [Setting Advanced Authoritative DNS Server Properties, on page 104](#)
- [Running Caching DNS and Authoritative DNS on the Same Server, on page 107](#)
- [Troubleshooting DNS Servers, on page 109](#)

Setting DNS Server Properties

You can set properties for the DNS server, along with those you already set for its zones. These include:

- **General server properties**—See [Setting General DNS Server Properties, on page 66](#)
- **Log Settings**—See [Specifying Log Settings, on page 66](#)
- **Packet Logging**—See [Enabling Packet Logging, on page 67](#)
- **Activity Summary Settings**—See [Specifying Activity Summary Settings, on page 69](#)
- **Top Names Settings**—See [Specifying Top Names Settings, on page 92](#)
- **TLS Settings**—See [Specifying TLS Settings, on page 92](#)
- **Round-Robin server processing**—See [Enabling Round-Robin, on page 94](#)
- **Enabling Weighted Round-Robin**—See [Enabling Weighted Round-Robin, on page 95](#)
- **Enabling incremental zone transfers**—See [Enabling Incremental Zone Transfers \(IXFR\), on page 96](#)
- **Restricting Zone Queries**—See [Restricting Zone Queries, on page 96](#)
- **Enabling NOTIFY packets**—See [Enabling NOTIFY, on page 96](#)



Note To enable GSS-TSIG support, you must set *tsig-processing* to none, and *gss-tsig-processing* to 'ddns, query' to support both ddns and query.

- **Blocking recursive queries**—See [Blocking Recursive Queries from Authoritative Server](#), on page 97

Setting General DNS Server Properties

You can display general DNS server properties, such as the name of the server cluster or host machine and the version number of the Cisco Prime Network Registrar DNS server software. You can change the internal name of the DNS server by deleting the current name and entering a new one. This name is used for notation and does not reflect the official name of the server. Cisco Prime Network Registrar uses the server IP address for official name lookups and for DNS updates (see the *"Managing DNS Update"* chapter in *Cisco Prime Network Registrar 11.0 DHCP User Guide*).

The following subsections describe some of the more common property settings. They are listed in [Setting DNS Server Properties](#), on page 65.

Local Basic or Advanced Web UI

-
- Step 1** To access the server properties, from the **Deploy** menu, choose **DNS Server** under the **DNS** submenu to open the Manage DNS Authoritative Server page. The page displays all the DNS server attributes.
 - Step 2** Modify the attributes as per your requirements.
 - Step 3** Click **Save** to save the DNS server attribute modifications.
-

CLI Commands

Use **dns show** to display the DNS server properties.

Specifying Log Settings

The *server-log-settings* attribute determines which events to log in the DNS log files. Default flags are activity-summary, config, update, xfr-in, xfr-out, scp, scavenge, server-operations, and ha.

Logging additional detail about events can help analyze a problem. However, leaving detailed logging enabled for a long period can fill up the log files.

The possible options are:

- **activity-summary**—This setting enables logging of DNS statistic messages at the interval specified by *activity-summary-interval*. The type of statistics logged can be controlled with *activity-counter-log-settings* and *activity-summary-type*.
- **config**—This setting enables logging of DNS server configuration and de-initialization messages.
- **config-detail**—This setting enables logging of detailed configuration messages (that is, detailed zone configuration logging).
- **db**—This setting enables logging of database processing messages. Enabling this flag provides insight into various events in the server's embedded databases.

- **dnssec**—This setting enables log messages associated with DNSSEC processing.
- **ha**—This setting enables logging of HA DNS messages.
- **host-health-check**—This setting enables logging associated with DNS Host Health Check.
- **notify**—This setting enables logging of messages associated with NOTIFY processing.
- **query**—This setting enabled logging of messages associated with QUERY processing.
- **scavenge**—This setting enables logging of DNS scavenging messages.
- **scp**—This setting enabled logging associated with SCP messages handling.
- **server-operations**—This setting enables logging of general server events, such as those pertaining to sockets and interfaces.
- **tsig**—This setting enables logging of events associated Transaction Signature (TSIG).
- **update**—This setting enables logging of DNS Update message processing.
- **xfr-in**—This setting enables logging of inbound full and incremental zone transfers.
- **xfr-out**—This setting enables logging of outbound full and incremental zone transfers.

Enabling Packet Logging

Cisco Prime Network Registrar supports packet logging for Authoritative DNS server to help analyze and debug the Authoritative DNS server activity. The packet logging settings determine the type of packet logging (summary or detail), the type of packets logged, and to which log file the messages are logged. By default, the Authoritative DNS server does not log any packet log messages.

Use the following server level attributes to enable packet logging for the Authoritative DNS server:

Table 19: Authoritative DNS Server Packet Logging Attributes

Attribute	Description
Packet Logging <i>(packet-logging)</i>	Determines the type of packet logging that is logged to the DNS logs. The type of DNS packets logged can be controlled with the <i>packet-log-settings</i> attribute. <ul style="list-style-type: none"> • disabled—This settings disables logging of DNS packets. • summary—This setting enables one line summary logging of DNS packets. • detail—This setting enables detailed packet tracing of DNS packets. <p>Note: While packet logging can be helpful for debugging and troubleshooting, it does have an impact on DNS server performance. Therefore, Cisco does not recommend leaving packet logging enabled in production environments.</p>
Packet Logging File <i>(packet-logging-file)</i>	Determines the destination log of packet log messages when packet logging is enabled. <ul style="list-style-type: none"> • dns—Packet logging messages are logged to the standard DNS log file (<i>name_dns_1_log*</i>). • packet—Packet logging messages are logged to a separate DNS packet log file (<i>dns_packet_log*</i>).

Attribute	Description
Packet Log Settings <i>(packet-log-settings)</i>	Determines the type of DNS messages to log if packet logging has been enabled. Packet logging can be enabled by configuring the <i>packet-logging</i> attribute. <ul style="list-style-type: none"> • all-in—This setting enables logging of all incoming packets. <i>Note:</i> This is equivalent to enabling all the -in settings. • all-out—This setting enabled logging of all outgoing packets. <i>Note:</i> This is equivalent to enabling all the -out settings. • ha-in, ha-out—These settings enable logging of HA DNS messages except for HA heartbeat and frame ACK messages which are controlled by the <i>ha-heartbeat-in</i>, <i>ha-heartbeat-out</i> and <i>ha-frameack-in</i>, <i>ha-frameack-out</i> settings, respectively. • ha-heartbeat-in, ha-heartbeat-out—These settings enable logging of HA DNS heartbeat messages. • ha-frameack-in, ha-frameack-out—These settings enable logging of HA DNS frame ACK messages. • notify-in, notify-out—These settings enable logging of DNS NOTIFY messages. • update-in, update-out—These settings enable logging of DNS UPDATE messages. • xfr-in, xfr-out—These settings enable logging of DNS IXFR and AXFR messages.

Local Advanced Web UI

- Step 1** On the Manage DNS Authoritative Server page, under the **Packet Logging** section, select the value for **packet-logging** from the drop-down list. The value can be **summary** or **detail**.
- Step 2** For the *packet-log-settings* attribute, check the desired check boxes.
- Step 3** Click **Save** to save the changes.

CLI Commands

Use **dns set packet-logging=summary** to enable one line summary packet logging.

Use **dns set packet-logging=detail** to enable detailed packet tracing.

Use **dns set packet-log-settings=value** to set the type of packets to log when packet logging is enabled.



Note Reloading of Authoritative DNS server is not required for the *packet-logging* and *packet-log-settings* attributes to take effect immediately (similar to log settings). However, the *packet-logging-file* attribute requires a Authoritative DNS server reload.

Specifying Activity Summary Settings



Note To specify the activity summary settings, you have to check *activity-summary* under the Log Settings.

You can specify the interval at which to log activity summary information using the Statistics Interval (*activity-summary-interval*) attribute. Enable the *activity-summary* attribute in the Log Settings (*server-log-settings*) attribute to set the seconds between DNS activity summary log messages. The *activity-summary-interval* attribute has a default value of 60 seconds.

The Authoritative DNS server logs sample and/or total statistics based on the option you check for the Statistics Type (*activity-summary-type*) attribute. The default value is "sample".

The option checked for the Statistics Settings (*activity-counter-log-settings*) attribute controls what activity counters a DNS server uses for logging.



Note *activity-summary-type* and *activity-counter-log-settings* take effect without a reload as soon as the DNS server object or the session is saved.

The possible settings are:

- **cache**—Log query cache related counters.

For the list of activity summary statistics that are displayed in the logs for the **cache** setting, see [Cache Statistics, on page 70](#).

- **db**—Log database counters.

For the list of activity summary statistics that are displayed in the logs for the **db** setting, see [DB Statistics, on page 71](#).

- **errors**—Log error related counters.

For the list of activity summary statistics that are displayed in the logs for the **errors** setting, see [Errors Statistics, on page 73](#).

- **ha**—Log HA related counters.

For the list of activity summary statistics that are displayed in the logs for the **ha** setting, see [HA Statistics, on page 74](#).

- **host-health-check**—Log DNS Host Health Check counters.

For the list of activity summary statistics that are displayed in the logs for the **host-health-check** setting, see [Host Health Check Statistics, on page 78](#).

- **ipv6**—Log IPv6 related counters.

For the list of activity summary statistics that are displayed in the logs for the **ipv6** setting, see [IPv6 Statistics, on page 80](#).

- **maxcounters**—Log maxcounters related counters.

For the list of activity summary statistics that are displayed in the logs for the **maxcounters** setting, see [Maxcounters Statistics, on page 80](#).

- **performance**—Log performance related counters.

For the list of activity summary statistics that are displayed in the logs for the **performance** setting, see [Performance Statistics, on page 81](#).

- **query**—Log query related counters.

For the list of activity summary statistics that are displayed in the logs for the **query** setting, see [Query Statistics, on page 83](#).

- **security**—Log security related counters.

For the list of activity summary statistics that are displayed in the logs for the **security** setting, see [Security Statistics, on page 86](#).

- **system**—Log system related counters.

For the list of activity summary statistics that are displayed in the logs for the **system** setting, see [System Statistics, on page 88](#).

- **top-names**—Log the top names queried and hit count.

For the list of activity summary statistics that are displayed in the logs for the **top-names** setting, see [Top Names Statistics, on page 89](#).

- **update**—Log DNS Update related counters.

For the list of activity summary statistics that are displayed in the logs for the **update** setting, see [Update Statistics, on page 89](#).

Activity Summary Statistics

Following sections describe the list of activity summary statistics that are displayed in the logs under each of the *activity-counter-log-settings* category.

Cache Statistics

The **cache** activity-counter-log-settings logs query cache related counters.

The cache activity summary statistics are logged under the **Query-Cache** sub category.

Sample log message:

```
10/22/2021 16:47:05 name/dns/1 Activity Stats 0 21333 [Query-Cache] Sample since Fri Oct
22 16:46:05 2021: size=number, #-records=number, #-rrs=number, nxdomain=number, hits=number,
misses=number, full=number, collisions=number
```

Table 20: Cache Statistics

Activity Summary Name	Statistic ¹⁰	Description
size	cache-size	Reports the size of in-memory query cache in bytes.
#-records	cache-records	Reports the total number of RR name sets stored in the query cache.
#-rrs	cache-rrs	Reports the total number of RRs stored in the query cache.

Activity Summary Name	Statistic ¹⁰	Description
nxdomain	cache-nxdomain	Reports the total number of NXDOMAIN entries in the query cache.
hits	cache-hits	Reports the number of times incoming client queries were found in the query cache.
misses	cache-misses	Reports the number of times incoming client queries were not found in the query cache.
full	cache-full	Reports the number of times the query cache was found to be at its configured limit (<i>mem-cache-size</i>).
collisions	N/A	Reports the number of times different FQDNs mapped to the same memory cache index. A high number of collisions indicates that the configured cache size may be too small.

¹⁰ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

DB Statistics

The **db** activity-counter-log-settings logs database counters.

Sample log message:

```
10/22/2021 16:44:05 name/dns/1 Activity Stats 0 21344 [Cset-DB] Sample since Fri Oct 22
16:43:05 2021: reads=number, writes=number, deletes=number, csets-trimmed=number,
conflicts=number, insufficient-history=number, txns=number, txn-commits=number,
txn-aborts=number, txn-locked=number, txn-unlocked=number, check-pts=number,
log-purges=number, #-logs-purged=number

10/22/2021 16:44:05 name/dns/1 Activity Stats 0 21345 [RR-DB] Sample since Fri Oct 22
16:43:05 2021: reads=number, writes=number, deletes=number, check-pts=number,
log-purges=number, #-logs-purged=number, txns=number, txn-commits=number, txn-aborts=number

10/22/2021 16:44:05 name/dns/1 Activity Stats 0 21352 [Cset-Queue] Sample since Fri Oct 22
16:43:05 2021: cset-count=number, cset-queue-max-size=number, commits=number,
commits-failed=number
```

Table 21: DB Statistics

Activity Summary Name	Logging Sub Category	Statistic ¹¹	Description
txn	RR-DB	rrdb-txn	Reports the total number of RR DB database transactions.

Activity Summary Name	Logging Sub Category	Statistic ¹¹	Description
txn-commits	RR-DB	rrdb-txn-commits	Reports the total number of RR DB database transactions committed.
txn-aborts	RR-DB	rrdb-txn-aborts	Reports the total number of RR DB database transactions aborted.
reads	RR-DB	rrdb-reads	Reports the total number of RR DB read operations.
writes	RR-DB	rrdb-writes	Reports the total number of RR DB write operations.
deletes	RR-DB	rrdb-deletes	Reports the total number of RR DB delete operations.
check-pts	RR-DB	rrdb-check-pts	Reports the total number of RR DB check point operations.
log-purges	RR-DB	rrdb-log-purges	Reports the total number of RR DB log purge operations.
#-logs-purged	RR-DB	rrdb-log-purges-count	Reports the total number of RR DB logs purged.
cset-count	Cset-Queue	csetq-count	Reports the total of number of change sets queued up to be written to the cset DB.
cset-queue-max-size	Cset-Queue	N/A	The maximum number of cset entries queued during this interval.
commits	Cset-Queue	N/A	Number of DB commits that happened in the last interval.
commits-failed	Cset-Queue	N/A	Number of DB commits that failed in the last interval.
txns	Cset-DB	csetdb-txn	Reports the total number of CSET DB database transactions.
txn-commits	Cset-DB	csetdb-txn-commits	Reports the total number of CSET DB database transactions committed.
txn-aborts	Cset-DB	csetdb-txn-aborts	Reports the total number of CSET DB database transactions aborted.
reads	Cset-DB	csetdb-reads	Reports the total number of CSET DB read operations.
writes	Cset-DB	csetdb-writes	Reports the total number of CSET DB write operations.

Activity Summary Name	Logging Sub Category	Statistic ¹¹	Description
deletes	Cset-DB	csetdb-deletes	Reports the total number of CSET DB delete operations.
csets-trimmed	Cset-DB	csetdb-csets-trimmed	Reports the total number of change sets trimmed from the CSET DB by the history trimming process or by inline trimming.
check-pts	Cset-DB	csetdb-check-pts	Reports the total number of CSET DB check point operations.
log-purges	Cset-DB	csetdb-log-purges	Reports the total number of CSET DB log purge operations.
#-logs-purged	Cset-DB	csetdb-log-purges-count	Reports the total number of CSET DB logs purged.

¹¹ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in Cisco Prime Network Registrar 11.0 Administration Guide.

Errors Statistics

The **errors** activity-counter-log-settings logs error related counters.

The errors activity summary statistics are logged under the **Errors** sub category.

Sample log message:

```
10/22/2021 16:44:05 name/dns/1 Activity Stats 0 21492 [Errors] Sample since Fri Oct 22
16:43:05 2021: update-errors=number, update-prereq-fail=number, ixfr-in-errors=number,
ixfr-out-errors=number, axfr-in-errors=number, axfr-out-errors=number,
xfer-in-auth-errors=number, xfer-failed-attempts=number, sent-total-errors=number,
sent-refusal-errors=number, sent-format-errors=number, exceeded-max-dns-packets=number
```

Table 22: Errors Statistics

Activity Summary Name	Statistic ¹²	Description
update-errors	update-errors	Reports the total number of updates resulting in errors. This excludes negative responses to update prerequisite checks, and TSIG responses. Both update packets and updates generated by the CNR UIs may be included in this count.
update-prereq-fail	update-prereq-fail	Reports the total number of updates resulting in prerequisite failures.
ixfr-in-errors	ixfr-in-errors	Reports the total in-bound IXFR errors, excluding packet format errors.

Activity Summary Name	Statistic ¹²	Description
ixfr-out-errors	ixfr-out-errors	Reports the total IXFR error responses sent, excluding packet format errors.
axfr-in-errors	axfr-in-errors	Reports the total in-bound AXFR errors, excluding packet format errors.
axfr-out-errors	axfr-out-errors	Reports the total AXFR error responses sent, excluding packet format errors.
sent-total-errors	sent-total-errors	Reports the total number of requests the server answered with errors (RCODE values other than 0,3,6,7, and 8). See RFC 1611.
sent-format-errors	sent-format-errors	Reports the number of requests received that were unparseable. See RFC 1611.
sent-refusal-errors	sent-refusal-errors	Reports the number of requests that resulted in REFUSED. See RFC1611.
xfer-in-auth-errors	xfer-in-auth-errors	Reports the number of secondary IXFR/AXFR requests that were refused because of authorization errors.
xfer-failed-attempts	xfer-failed-attempts	Reports the number of secondary IXFR/AXFR failures, excluding authorization refusals.
exceeded-max-dns-packets	exceeded-max-dns- packets	Reports the number of times inbound packets exceeded the maximum DNS packets defined by <i>max-dns-packets</i> .

¹² The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

HA Statistics

The **ha** activity-counter-log-settings logs HA related counters.

Sample log message:

```
name_dns_1_log:11/19/2021 11:43:23 name/dns/1 Activity Stats 0 20005 [HA-State] Sample since
Fri Nov 19 11:41:35 2021: current=state, last-state-change=time, normal=number,
comm-interrupted=number, negotiate=number, start-up=number, partner-down=number
```

```
name_dns_1_log:11/19/2021 12:09:23 name/dns/1 Activity Stats 0 21341 [HA-Requests-Sent]
Sample since Fri Nov 19 12:08:23 2021: requests-sent=number, last-req-sent=Heartbeat @ Fri
Nov 19 12:09:21 2021 (xid: 207), update=number, heart-beat=number, zone-sync=number,
rr-sync=number, rr-recon=number, connect=number, negotiate=number, shutdown=number,
truncated=number
```

```
name_dns_1_log:11/18/2021 13:07:26 name/dns/1 Activity Stats 0 21342 [HA-Requests-Rcvd]
```

Sample since Thu Nov 18 13:04:12 2021: requests-recv=number, last-req-recv=Heartbeat @ Thu Nov 18 13:07:07 2021 (xid: 207), update=number, heart-beat=number, zone-sync=number, rr-sync=number, rr-recon=number, connect=number, negotiate=number, shutdown=number, truncated=number

11/29/2021 9:02:44 name/dns/1 Activity Stats 0 21343 [HA-Errors] Sample since Mon Nov 29 09:01:44 2021: update-reject=number, resp-mismatch=number, resp-inconsistent=number, resp-servfail=number, resp-unknown=number

11/29/2021 14:49:32 name/dns/1 Activity Stats 0 20006 [HA-Zone-Sync] Sample since Mon Nov 29 14:47:32 2021: sync=number, sync-completed=number, sync-failed=number, zone-mismatch=number, full-resync=number, conflict=number, merge=number, discard=number

Table 23: HA Statistics

Activity Summary Name	Logging Sub Category	Statistic ¹³	Description
comm-interrupted	HA-State	ha-state-comm-interrupted	Number of occurrences where the server enters the communication-interrupted state (HA_STATE_COMMINTR).
partner-down	HA-State	ha-state-partner-down	Number of occurrences where the server enters the partner-down state (HA_STATE_PARTNERDOWN).
negotiate	HA-State	ha-state-negotiating	Number of occurrences where the server enters the Negotiating state (HA_STATE_NEGOTIATING).
current	HA-State	ha-state-current	Current HA server state.
last-state-change	HA-State	ha-state-last-change-time	Last time when HA state changed.
start-up	HA-State	ha-state-startup	Number of occurrences where the server enters Startup State (HA_STARTUP).
normal	HA-State	ha-state-normal	Number of occurrences where the server enters Normal State (HA_NORMAL).
connect	HA-Requests-Sent	ha-msg-connect-sent	Number of connection establishment request messages sent (HA_DNS_ESTABLISH_CONNECTION).
rr-recon	HA-Requests-Sent	ha-msg-reconcile-sent	Number of zone reconciliation request messages sent (HA_DNS_RECONCILIATION).
heart-beat	HA-Requests-Sent	ha-msg-heartbeat-sent	Number of heartbeat request messages sent (HA_DNS_HEARTBEAT).
zone-sync	HA-Requests-Sent	ha-msg-zonesync-sent	Number of zone synchronization request messages sent (HA_DNS_ZONE_SYNC).
rr-sync	HA-Requests-Sent	ha-msg-rrsync-sent	Number of rr-sync request messages sent (HA_DNS_RR_SYNC).

Activity Summary Name	Logging Sub Category	Statistic ¹³	Description
update	HA-Requests-Sent	ha-msg-rrupdate-sent	Number of rr-update request messages sent (HA_DNS_RR_UPDATE).
N/A	N/A	ha-msg-resp-sent	Number of response messages sent. Response messages are used to acknowledge all types of request messages.
shutdown	HA-Requests-Sent	ha-msg-shutdown-sent	Number of shutdown request messages sent.
requests-sent	HA-Requests-Sent	ha-msg-req-sent	Number of HA request messages sent to the HA partner.
last-req-sent	HA-Requests-Sent	ha-msg-req-sent-time	Specifies the date and time the HA server last sent a request message to the HA partner.
negotiate	HA-Requests-Sent	N/A	Number of negotiate HA message sent.
truncated	HA-Requests-Sent	N/A	Number of HA messages sent that were truncated.
connect	HA-Requests-Rcvd	ha-msg-connect-recv	Number of connection establishment request messages received (HA_DNS_ESTABLISH_CONNECTION).
rr-recon	HA-Requests-Rcvd	ha-msg-reconcile-recv	Number of zone reconciliation request messages received (HA_DNS_RECONCILIATION).
heart-beat	HA-Requests-Rcvd	ha-msg-heartbeat-recv	Number of heartbeat request messages received (HA_DNS_HEARTBEAT).
zone-sync	HA-Requests-Rcvd	ha-msg-zonesync-recv	Number of zone synchronization request messages received (HA_DNS_ZONE_SYNC).
rr-sync	HA-Requests-Rcvd	ha-msg-rrsync-recv	Number of rr-sync messages request received (HA_DNS_RR_SYNC).
update	HA-Requests-Rcvd	ha-msg-rrupdate-recv	Number of rr-update request messages received (HA_DNS_RR_UPDATE).
N/A	N/A	ha-msg-resp-recv	Number of response messages received. Response messages are used to acknowledge all types of request messages.
shutdown	HA-Requests-Rcvd	ha-msg-shutdown-recv	Number of shutdown request messages received.

Activity Summary Name	Logging Sub Category	Statistic ¹³	Description
requests-recv	HA-Requests-Rcvd	ha-msg-req-recv	Number of HA request messages received from the HA partner.
last-req-recv	HA-Requests-Rcvd	ha-msg-req-recv-time	Specifies the date and time the HA server last received a request message from the HA partner.
negotiate	HA-Requests-Rcvd	N/A	Number of negotiate HA message received.
truncated	HA-Requests-Rcvd	N/A	Number of HA messages received that were truncated.
update-reject	HA-Errors	ha-update-reject	Number of DNS updates rejected by the server.
resp-mismatch	HA-Errors	ha-zone-mismatch	Number of zones reporting a mismatch error (HA_DNS_RESP_ERR_MISMATCH).
resp-servfail	HA-Errors	ha-resp-servfail	Number of responses reporting a server failure error (HA_DNS_RESP_ERR_SERVFAIL).
resp-inconsistent	HA-Errors	ha-resp-inconsistent	Number of responses reporting an inconsistent server state (HA_DNS_RESP_ERR_INCONSISTENT_STATE).
resp-unknown	HA-Errors	ha-resp-unknown	Number of responses with an unknown message type (HA_DNS_RESP_ERR_UNKNOWN_MSG_TYPE).
full-resync	HA-Zone-Sync	ha-full-zone-resync	Number of zones requiring full-zone resynchronization for nameset reconciliation.
conflict	HA-Zone-Sync	ha-sync-conflict	Number of zones with name conflicts during nameset reconciliation.
discard	HA-Zone-Sync	ha-sync-discard-name	Number of name conflicts where one nameset must be discarded to synchronize the zone.
merge	HA-Zone-Sync	ha-sync-merge-name	Number of name conflicts which the namesets can be merged to synchronize the zone.
sync	HA-Zone-Sync	N/A	Number of zones that were requested to be synced.

Activity Summary Name	Logging Sub Category	Statistic ¹³	Description
sync-completed	HA-Zone-Sync	N/A	Number of zones where sync was completed.
sync-failed	HA-Zone-Sync	N/A	Number of zones where sync failed.
zone-mismatch	HA-Zone-Sync	N/A	Number of zones that do not match on HA Main and HA Backup.

¹³ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Host Health Check Statistics

The **host-health-check** activity-counter-log-settings logs DNS Host Health Check counters.

The host health check activity summary statistics are logged under the **HHC** sub category.

Sample log message:

```
10/22/2021 16:44:05 name/dns/1 Activity Stats 0 21509 [HHC] Sample since Fri Oct 22 16:43:05
2021: hhc-domains=number, hhc-domains-failed=number, hhc-domains-passed=number,
hhc-rrs=number, hhc-rrs-passed=number, hhc-rrs-failed=number, hhc-ping-domains=number,
hhc-ping-domains-failed=number, hhc-ping-domains-passed=number, hhc-ping-rrs=number,
hhc-ping-rrs-passed=number, hhc-ping-rrs-failed=number, hhc-gtp-echo-domains=number,
hhc-gtp-echo-domains-failed=number, hhc-gtp-echo-domains-passed=number,
hhc-gtp-echo-rrs=number, hhc-gtp-echo-rrs-passed=number, hhc-gtp-echo-rrs-failed=number
```

Table 24: Host Health Check Statistics

Activity Summary Name	Statistic ¹⁴	Description
hhc-domains	hhc-domains	Reports the total number of domains checked for Host Health Check.
hhc-domains-failed	hhc-domains-failed	Reports the total number of domains check failed for Host Health Check. When all the RRs in the RR set are down, this stat is incremented.
hhc-domains-passed	hhc-domains-passed	Reports the total number of domains check passed for Host Health Check. Any A/AAAA RR in the RR set is up, this stat is incremented.
hhc-rrs	hhc-rrs	Reports the total number of RRs checked for Host Health Check.
hhc-rrs-passed	hhc-rrs-passed	Reports the total number of RRs that have passed Host Health Check health check.

Activity Summary Name	Statistic ¹⁴	Description
hhc-rrs-failed	hhc-rrs-failed	Reports the total number of RRs that have failed Host Health Check health check.
hhc-ping-domains	hhc-ping-domains	Reports the total number of domains checked for ping Host Health Check.
hhc-ping-domains-failed	hhc-ping-domains-failed	Reports the total number of domains check failed for ping Host Health Check. When all the RRs in the RR set are down, this stat is incremented.
hhc-ping-domains-passed	hhc-ping-domains-passed	Reports the total number of domains check passed for ping Host Health Check. When any RR in the RR set is up, this stat is incremented.
hhc-ping-rrs	hhc-ping-rrs	Reports the total number of RRs checked for ping Host Health Check.
hhc-ping-rrs-failed	hhc-ping-rrs-failed	Reports the total number of RRs that have failed ping Host Health Check health check.
hhc-ping-rrs-passed	hhc-ping-rrs-passed	Reports the total number of RRs that have passed ping Host Health Check health check.
hhc-gtp-echo-domains	hhc-gtp-echo-domains	Reports the total number of domains checked for gtp-echo Host Health Check.
hhc-gtp-echo-domains-failed	hhc-gtp-echo-domains-failed	Reports the total number of domains check failed for gtp-echo Host Health Check. When all the RRs in the RR set are down, this stat is incremented.
hhc-gtp-echo-domains-passed	hhc-gtp-echo-domains-passed	Reports the total number of domains check passed for gtp-echo Host Health Check. When any RR in the RR set is up, this stat is incremented.
hhc-gtp-echo-rrs	hhc-gtp-echo-rrs	Reports the total number of RRs checked for gtp-echo Host Health Check.
hhc-gtp-echo-rrs-failed	hhc-gtp-echo-rrs-failed	Reports the total number of RRs that have failed gtp-echo Host Health Check health check.
hhc-gtp-echo-rrs-passed	hhc-gtp-echo-rrs-passed	Reports the total number of RRs that have passed gtp-echo Host Health Check health check.

¹⁴ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

IPv6 Statistics

The **ipv6** activity-counter-log-settings logs IPv6 related counters.

The IPv6 activity summary statistics are logged under the **Perform** sub category.

Sample log message:

```
11/26/2021 15:25:36 name/dns/1 Activity Stats 0 03523 [Perform] Sample since Fri Nov 26
15:24:36 2021: pkts-in=number, pkts-out=number, pkts-in-udp=number, pkts-out-udp=number,
pkts-in-tcp=number, pkts-out-tcp=number, ipv4-pkts-in=number, ipv4-pkts-out=number,
ipv6-pkts-in=number, ipv6-pkts-out=number, queries=number, updates=number, notifies-in=number,
notifies-out=number, notify-errors=number, ixfrs-in=number, ixfrs-out=number,
ixfrs-full-resp=number, axfrs-in=number, axfrs-out=number, xfrs-in-at-limit=number,
xfrs-out-at-limit=number, responses-with-NOTIMP=number, total-zones=number, total-rxs=number
```

Table 25: IPv6 Statistics

Activity Summary Name	Statistic ¹⁵	Description
ipv6-pkts-in	ipv6-packets-in	Total number of IPv6 packets received.
ipv6-pkts-out	ipv6-packets-out	Total number of IPv6 packets sent.

¹⁵ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Maxcounters Statistics

The **maxcounters** activity-counter-log-settings logs maxcounters related counters.

The maxcounters activity summary statistics are logged under the **Max-Counters** sub category.

Sample log message:

```
10/22/2021 16:40:05 name/dns/1 Activity Stats 0 21353 [Max-Counters] Sample since Tue Oct
19 19:32:39 2021: concurrent-xfrs-in=number, concurrent-xfrs-out=number,
ha-update-latency-max=number, ha-batch-count-limit=number, ha-rr-pending-list=number,
ha-rr-active-list=number, ha-persisted-edit-list=number, dns-concurrent-packets=number,
pn-conn-max-conns=number
```

Table 26: Maxcounters Statistics

Activity Summary Name	Statistic ¹⁶	Description
concurrent-xfrs-in	concurrent-xfrs-in	Reports the maximum number of concurrent threads processing inbound transfers during the last sampling period.
concurrent-xfrs-out	concurrent-xfrs-out	Reports the maximum number of concurrent threads processing outbound transfers during the last sampling period.

Activity Summary Name	Statistic ¹⁶	Description
ha-batch-count-limit	ha-batch-count-limit	Reports the number of times the <i>ha-dns-max-batch-count</i> limit was reached during the last sampling period.
ha-rr-pending-list	ha-rr-pending-list	Reports the maximum number of RRs in the pending List, waiting acknowledgement from the HA DNS backup server, during the last sampling period.
ha-rr-active-list	ha-rr-active-list	Reports the maximum number of RRs in the active list, waiting to be sent to the HA DNS backup server, during the last sampling period.
ha-persisted-edit-list	ha-persisted-edit-list	Reports the maximum number of names persisted in the edit list database during the last sampling period.
ha-update-latency- max	ha-update-latency-max	Reports the maximum DNS update latency in seconds, during the last sampling period. Latency is measured as the time an update remains in the pending List.
dns-concurrent- packets	dns-concurrent-packets	Reports the maximum number of concurrent packets processed by the DNS server during the sampling period.

¹⁶ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Performance Statistics

The **performance** activity-counter-log-settings logs performance related counters.

The performance activity summary statistics are logged under the **Perform** sub category.

Sample log message:

```
10/22/2021 16:40:05 name/dns/1 Activity Stats 0 03523 [Perform] Sample since Tue Oct 19
19:32:39 2021: pkts-in=number, pkts-out=number, pkts-in-udp=number,pkts-out-udp=number,
pkts-in-tcp=number, pkts-out-tcp=number, ipv4-pkts-in=number, ipv4-pkts-out=number,
ipv6-pkts-in=number, ipv6-pkts-out=number, queries=number, updates=number,notifies-in=number,
notifies-out=number, notify-errors=number, ixfrs-in=number, ixfrs-out=number,
ixfrs-full-resp=number, axfrs-in=number, axfrs-out=number, xfrs-in-at-limit=number,
xfrs-out-at-limit=number, responses-with-NOTIMP=number, total-zones=number, total-rrs=number
```

Table 27: Performance Statistics

Activity Summary Name	Statistic ¹⁷	Description
ipv4-pkts-in	ipv4-packets-in	Reports the total number of IPv4 packets received.

Activity Summary Name	Statistic ¹⁷	Description
ipv4-pkts-out	ipv4-packets-out	Reports the total number of IPv4 packets sent.
N/A	updated-rrs	Reports the total number of RRs added and deleted, including updates from the CPNR UIs, whether or not there were database errors.
updates	update-packets	Reports the number of successful DNS updates.
queries	queries-total	Total number of queries received by the DNS Server.
ixfrs-out	ixfrs-out	Reports the number of successful outbound incremental transfers.
ixfrs-in	ixfrs-in	Reports the number of successful inbound incremental transfers, including incremental requests that resulted in full zone transfers.
ixfrs-full-resp	ixfrs-full-resp	Reports the number of outbound full zone transfers in response to IXFR requests. These may have been due to IXFR errors, insufficient serial history, or too many changes in the zone.
axfrs-in	axfrs-in	Reports the number of successful inbound AXFRs.
axfrs-out	axfrs-out	Reports the number of successful outbound full zone transfers, including those counted in <i>ixfrs-full-resp</i> .
xfrs-in-at-limit	xfrs-in-at-limit	Reports the number of times that inbound transfers reached the concurrent limit.
xfrs-out-at-limit	xfrs-out-at-limit	Reports the number of times that outbound transfers reached the concurrent limit.
notifies-out	notifies-out	Reports the number of outbound notifies. Each notify packet sent is counted separately.
notifies-in	notifies-in	Reports the number of inbound notifies. Each notify packet received is counted separately.
notify-errors	N/A	Errors detected while processing notify requests.
total-zones	N/A	Total number of zones configured.
total-rrs	N/A	Total number of RRs across all configured zones.
responses-with-NOTIMP	responses-with-NOTIMP	Reports the numbers of requests with OP codes that are not implemented.
pkts-in	packets-in	Reports the total number of packets received.
pkts-out	packets-out	Reports the total number of packets sent.

Activity Summary Name	Statistic ¹⁷	Description
pkts-in-udp	packets-in-udp	Reports the total number of UDP packets received.
pkts-out-udp	packets-out-udp	Reports the total number of UDP packets sent.
pkts-in-tcp	packets-in-tcp	Reports the total number of TCP packets received.
pkts-out-tcp	packets-out-tcp	Reports the total number of TCP packets sent.
ipv6-pkts-in	ipv6-packets-in	Reports the total number of IPv6 packets received.
ipv6-pkts-out	ipv6-packets-out	Reports the total number of IPv6 packets sent.

¹⁷ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in Cisco Prime Network Registrar 11.0 Administration Guide.

Query Statistics

The **query** activity-counter-log-settings logs query related counters.

Sample log message:

```
10/22/2021 16:41:05 name/dns/1 Activity Stats 0 21168 [Query] Sample since Fri Oct 22
16:40:05 2021: total=number, dropped=number, acl-failures=number, udp=number, tcp=number,
ipv4=number, ipv6=number, tls=number, tls-failures=number, dropped-recursive=number,
dropped-unwanted-class=number, dropped-unwanted-type=number

10/22/2021 16:44:05 name/dns/1 Activity Stats 0 21333 [Query-Cache] Sample since Fri Oct
22 16:43:05 2021: size=number, #-records=number, #-rrs=number, nxdomain=number, hits=number,
misses=number, full=number, collisions=number

10/22/2021 16:41:05 name/dns/1 Activity Stats 0 21331 [Query-Type] Sample since Fri Oct 22
16:40:05 2021: A=number, AAAA=number, ANY=number, CNAME=number, MX=number, NAPTR=number,
NS=number, PTR=number, SOA=number, SRV=number, TXT=number, DNSKEY=number, DS=number,
RRSIG=number, NSEC=number, CAA=number, URI=number, other=number

10/22/2021 16:41:05 name/dns/1 Activity Stats 0 21332 [Query-Responses] Sample since Fri
Oct 22 16:40:05 2021: total=number, no-error=number, referrals=number, no-data=number,
nxdomain=number, refused=number, notauth=number, formerr=number, servfail=number, other=number

10/22/2021 16:41:05 name/dns/1 Activity Stats 0 21524 [DNSSEC] Sample since Fri Oct 22
16:40:05 2021: dnssec-zones=number, dnssec-sign-zone=number, dnssec-queries=number,
dnssec-responses=number, dnssec-requests-dropped=number
```

Table 28: Query Statistics

Activity Summary Name	Logging Sub Category	Statistic ¹⁸	Description
hits	Query-Cache	mem-cache-hits	Reports the number of mem-cache lookup hits.

Activity Summary Name	Logging Sub Category	Statistic ¹⁸	Description
misses	Query-Cache	mem-cache-misses	Reports the number of mem-cache lookup misses.
dropped	Query	queries-dropped	Reports the number of non-error dropped packets. Queries restricted by server, TSIG, or update policies are included, but DNS updates, xfer requests, and notifies are excluded.
N/A	N/A	queries-with-edns	Reports the number of OPT RR packets processed.
total	Query	queries-total	Total number of queries received by the DNS Server.
udp	Query	queries-over-udp	Total number of queries received over UDP by the DNS Server.
tcp	Query	queries-over-tcp	Total number of queries received over TCP by the DNS Server.
ipv4	Query	queries-over-ipv4	Total number of IPv4 queries received by the DNS Server.
ipv6	Query	queries-over-ipv6	Total number of IPv6 queries received by the DNS Server.
tls	Query	queries-over-tls	Total number of queries received over TLS by the DNS Server.
tls-failures	Query	queries-over-tls-failed	Total number of TLS queries failed during TLS handshake.
dropped-recursive	Query	queries-dropped-recursive	Number of recursive queries dropped.
dropped-unwanted-class	Query	queries-dropped-unwanted-class	Total number of queries dropped due to unwanted classes. Only queries of class IN are allowed.
dropped-unwanted-type	Query	queries-dropped-unwanted-type	Total number of queries dropped due to unwanted types. Unwanted RR types are specified in the <i>query-types-unwanted</i> DNS server attribute.
acl-failures	Query	queries-failed-acl	Reports the number of query ACL (<i>restrict-query-acl</i>) failures.
total	Query-Responses	query-answers-total	Reports the total number of query responses.

Activity Summary Name	Logging Sub Category	Statistic ¹⁸	Description
no-error	Query-Responses	query-answers-with-NOERROR	Reports the number of queries that were authoritatively answered.
nxdomain	Query-Responses	query-answers-with-NXDOMAIN	Reports the number of queries that failed with no such name responses.
no-data	Query-Responses	query-answers-with-NODATA	Reports the number of queries that failed with no data (empty answer) responses.
notauth	Query-Responses	query-answers-with-NOTAUTH	Reports the number of queries that failed with not authoritative responses.
referrals	Query-Responses	query-answers-with-referral	Reports the number of requests that were referred to other servers.
refused	Query-Responses	query-answers-with-REFUSED	Reports the number of queries refused.
formerror	Query-Responses	query-answers-with-FORMERR	Reports the number of query responses with rcode of FORMERR.
servfail	Query-Responses	query-answers-with-SERVFAIL	Reports the number of query responses with rcode of SERVFAIL.
other	Query-Responses	query-answers-with-other-errors	Reports the number of queries with other errors.
dnssec-queries	DNSSEC	queries-dnssec	Reports the total number of queries requesting that responses to include DNSSEC related RRs (EDNS option DO bit).
A	Query-Type	queries-type-A	Number of A queries received.
AAAA	Query-Type	queries-type-AAAA	Number of AAAA queries received.
CNAME	Query-Type	queries-type-CNAME	Number of CNAME queries received.
PTR	Query-Type	queries-type-PTR	Number of PTR queries received.
NS	Query-Type	queries-type-NS	Number of NS queries received.
SOA	Query-Type	queries-type-SOA	Number of SOA queries received.
MX	Query-Type	queries-type-MX	Number of MX queries received.
NAPTR	Query-Type	queries-type-NAPTR	Number of NAPTR queries received.
other	Query-Type	queries-type-other	All other queries received.
ANY	Query-Type	queries-type-ANY	Number of ANY queries received.

Activity Summary Name	Logging Sub Category	Statistic ¹⁸	Description
SRV	Query-Type	queries-type-SRV	Number of SRV queries received.
TXT	Query-Type	queries-type-TXT	Number of TXT queries received.
DNSKEY	Query-Type	queries-type-DNSKEY	Number of DNSKEY queries received.
DS	Query-Type	queries-type-DS	Number of DS queries received.
RRSIG	Query-Type	queries-type-RRSIG	Number of RRSIG queries received.
NSEC	Query-Type	queries-type-NSEC	Number of NSEC queries received.
CAA	Query-Type	queries-type-CAA	Number of CAA queries received.
URI	Query-Type	queries-type-URI	Number of URI queries received.

¹⁸ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Security Statistics

The **security** activity-counter-log-settings logs security related counters.

Sample log message:

```
10/22/2021 16:44:05 name/dns/1 Activity Stats 0 21492 [Errors] Sample since Fri Oct 22 16:43:05 2021: update-errors=number, update-prereq-fail=number, ixfr-in-errors=number, ixfr-out-errors=number, axfr-in-errors=number, axfr-out-errors=number, xfer-in-auth-errors=number, xfer-failed-attempts=number, sent-total-errors=number, sent-refusal-errors=number, sent-format-errors=number, exceeded-max-dns-packets=number
```

```
10/22/2021 16:41:05 name/dns/1 Activity Stats 0 21332 [Query-Responses] Sample since Fri Oct 22 16:40:05 2021: total=number, no-error=number, referrals=number, no-data=number, nxdomain=number, refused=number, notauth=number, formerr=number, servfail=number, other=number
```

```
11/19/2021 16:59:41 name/dns/1 Activity Stats 0 21524 [DNSSEC] Sample since Fri Nov 19 16:58:41 2021: dnssec-zones=number, dnssec-sign-zone=number, dnssec-queries=number, dnssec-responses=number, dnssec-requests-dropped=number
```

```
11/26/2021 16:16:45 name/dns/1 Activity Stats 0 21491 [TSIG] Sample since Fri Nov 26 16:15:45 2021: tsig-packets=number, badtime=number, badkey=number, badsig=number, badtime-resp=number, badkey-resp=number, badsig-resp=number
```

```
12/08/2021 12:58:42 name/dns/1 Activity Stats 0 21389 [RPZ] Sample since Wed Dec 8 12:57:03 2021: rpz-queries=number, rpz-hits=number, rpz-misses=number
```

Table 29: Security Statistics

Activity Summary Name	Logging Sub Category	Statistic ¹⁹	Description
xfer-in-auth-errors	Errors	unauth-xfer-reqs	Reports the number of ACL authorization failures in zone transfers.
N/A	N/A	unauth-update-reqs	Reports the number of ACL authorization failures in DNS updates. Administrative RR updates (from CPNR UIs) are excluded.
refused	Query-Responses	restrict-query-acl	Reports the number of ACL authorization failures in DNS queries.
N/A	N/A	blackhole-acl-dropped-requests	Reports the number of DNS requests dropped by the server subject to <i>blackhole-acl</i> .
tsig-packets	TSIG	rcvd-tsig-packets	Reports the number of TSIG RR packets processed, if TSIG processing is enabled for the type of packet.
badtime-resp	TSIG	detected-tsig-bad-time	Reports the number of bad timestamps in incoming TSIG packets.
badkey-resp	TSIG	detected-tsig-bad-key	Reports the number of bad keynames (those with an invalid or unknown key) in incoming TSIG packets.
badsig-resp	TSIG	detected-tsig-bad-sig	Reports the number of bad signatures in incoming TSIG packets.
badtime	TSIG	rcvd-tsig-bad-time	Reports the number of BADTIME errors received after sending a TSIG packet.
badkey	TSIG	rcvd-tsig-bad-key	Reports the number of BADKEY errors received after sending a TSIG packet.
badsig	TSIG	rcvd-tsig-bad-sig	Reports the number of BADSIG errors received after sending a TSIG packet.
dnssec-zones	DNSSEC	dnssec-zones	Reports the number of zones with DNSSEC enabled.
dnssec-sign-zone	DNSSEC	dnssec-sign-zone	Reports the number of times the server signed a DNSSEC zone.
dnssec-queries	DNSSEC	dnssec-queries	Reports the total number of queries requesting that responses to include DNSSEC related RRs (EDNS option DO bit).

Activity Summary Name	Logging Sub Category	Statistic ¹⁹	Description
dnssec-responses	DNSSEC	dnssec-responses	Reports the total number of responses to DNSSEC enabled queries (EDNS option DO bit).
dnssec-requests-dropped	DNSSEC	dnssec-requests-dropped	Reports the total number of DNS requests that were dropped due to the server being in the process of signing a DNSSEC zone.
rpz-queries	RPZ	queries-rpz	Reports the number of queries for Response Policy Zones (RPZ).
rpz-hits	RPZ	query-answers-rpz-hits	Reports the number of RPZ queries that matched RRs in Response Policy Zones.
rpz-misses	RPZ	query-answers-rpz-misses	Reports the number of RPZ queries that did not match RRs in Response Policy Zones.

¹⁹ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

System Statistics

The **system** activity-counter-log-settings logs system related counters.

The system activity summary statistics are logged under the **System** sub category.

Sample log message:

```
10/22/2021 16:41:05 name/dns/1 Activity Stats 0 21493 [System] Sample since Fri Oct 22
16:40:05 2021: pid=number, cpu=number, memory=number, virtual=number, conntrack-max=number,
conntrack-count=number, conntrack-usage=number
```

Table 30: System Statistics

Activity Summary Name	Description
pid	The PID of the ADNS process.
cpu	The amount of CPU used by the ADNS process.
memory	The amount of memory used by the ADNS process.
virtual	The amount of virtual memory used by the ADNS process.
conntrack-max	The maximum number of Linux firewall connections reached.
conntrack-count	The current number of Linux firewall connections.

Activity Summary Name	Description
contrack-usage	The percentage of Linux firewall connections in use.

Top Names Statistics

The **top-names** activity-counter-log-settings logs the top names queried and hit count.

The top names activity summary statistics are logged under the **Top-Names** sub category.

Sample log message:

```
10/22/2021 16:55:05 name/dns/1 Activity Stats 0 21508 [Top-Names] from 16:53:05 to 16:54:05;
interval=number, total-counted=number
```

Table 31: Top Names Statistics

Activity Summary Name	Logging Sub Category	Statistic ²⁰	Description
interval	Top-Names	N/A	Length of data collection period.
total-counted	Top-Names	total-counted	Reports the total number of queries counted in this collection period.

²⁰ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in Cisco Prime Network Registrar 11.0 Administration Guide.

Update Statistics

The **update** activity-counter-log-settings logs DNS Update related counters.

Sample log message:

```
10/29/2021 15:56:31 name/dns/1 Activity Stats 0 21550 [Update] Sample since Fri Oct 29
15:55:31 2021: total=number, failed-acl=number, prereq-only=number, dropped=number,
simulated=number, udp=number, tcp=number, ipv4=number, ipv6=number, deletes=number,
adds=number, refreshes=number, rrs=number, A=number, AAAA=number, DHCID=number, TXT=number,
other=number
```

```
10/29/2021 15:56:31 name/dns/1 Activity Stats 0 21551 [Update-Responses] Sample since Fri
Oct 29 15:55:31 2021: total=number, no-error=number, failures=number, refused=number,
notauth=number, notzone=number, formerr=number, servfail=number, prereq-failures=number,
yxdomain=number, yxrrset=number, nxdomain=number, nxrrset=number
```

Table 32: Update Statistics

Activity Summary Name	Logging Sub Category	Statistic ²¹	Description
total	Update	update-total	Total number of updates received by the DNS server.

Activity Summary Name	Logging Sub Category	Statistic ²¹	Description
failed-acl	Update	update-failed-acl	Total number of updates that refused due to failing ACL and/or Update Policy authorization.
prereq-only	Update	update-prereq-only	Total number of prereq-only updates received by the DNS server.
dropped	Update	update-dropped	Total number of updates that are dropped by the DNS server.
simulated	Update	update-simulated	Total number of updates that are simulated. Simulated RR updates return a NOERROR response, but don't cause any RR changes.
udp	Update	update-over-udp	Total number of updates received over UDP.
tcp	Update	update-over-tcp	Total number of updates received over TCP.
ipv4	Update	update-over-ipv4	Total number of updates received over IPv4.
ipv6	Update	update-over-ipv6	Total number of updates received over IPv6.
deletes	Update	update-delete	Total number of RRs deleted by DNS update.
adds	Update	update-add	Total number of RRs added by DNS update.
refreshes	Update	update-refresh	Total number of RRs refreshed by DNS update.
rrs	Update	update-total-rrs	The total number of RRs updated by DNS update requests.
A	Update	update-type-A	Total number of updates for A records.
AAAA	Update	update-type-AAAA	Total number of updates for AAAA records.
DHCID	Update	update-type-DHCID	Total number of updates for DHCID records.
TXT	Update	update-type-TXT	Total number of updates for TXT records.
other	Update	update-type-other	Total number of updates for all other record types that are not specifically counted.

Activity Summary Name	Logging Sub Category	Statistic ²¹	Description
total	Update-Responses	update-resp-total	Total number of update responses returned by the DNS server.
no-error	Update-Responses	update-resp-NOERROR	Total number of update responses with rcode of NOERROR.
failures	Update-Responses	update-resp-failures	Total number of updates that failed.
refused	Update-Responses	update-resp-REFUSED	Total number of update responses with rcode of REFUSED.
notauth	Update-Responses	update-resp-NOTAUTH	Total number of update responses with rcode of NOTAUTH.
notzone	Update-Responses	update-resp-NOTZONE	Total number of update responses with rcode of NOTZONE.
formerr	Update-Responses	update-resp-FORMERR	Total number of update responses with rcode of FORMERR.
servfail	Update-Responses	update-resp-SERVFAIL	Total number of update responses with rcode of SERVFAIL.
prereq-failures	Update-Responses	update-resp-prereq-failures	Total number of update responses with prereq failures (YXDOMAIN, YXRRSET, NXDOMAIN, NXRRSET).
yxdomain	Update-Responses	update-resp-YXDOMAIN	Total number of update responses with rcode of YXDOMAIN.
yxrrset	Update-Responses	update-resp-YXRRSET	Total number of update responses with rcode of YXRRSET.
nxdomain	Update-Responses	update-resp-NXDOMAIN	Total number of update responses with rcode of NXDOMAIN.
nxrrset	Update-Responses	update-resp-NXRRSET	Total number of update responses with rcode of NXRRSET.

²¹ The statistics listed in this column are the server statistics displayed in the web UI and CLI. The REST API calls will have the statistic name camel-cased without dashes (that is, queries-total is queriesTotal in the REST API). Note that the activity summary and statistics are keyed off the same server data, but the activity-summary names are shortened to conserve space in the log message. For the complete list of Authoritative DNS server statistics, see the "DNS Statistics" section of the "Server Statistics" appendix in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Specifying Top Names Settings

The *top-names* attribute specifies if top names data should be collected. When enabled, a snapshot of the cache hits for the top names that are queried is collected for each interval set by the *top-names-max-age* value. The list of top names that is reported with activity summary statistics is the most current snapshot.

You can specify the maximum age (based on last access time) of a queried name allowed in the list of top names by using the *top-names-max-age* attribute. It has a default value of 60 seconds.

You can specify the maximum number of entries in the list of top names queried by using the *top-names-max-count* attribute. This limit is applied to the lists of top names that are logged or returned as part of activity summary.

Local Basic or Advanced Web UI

To enable Top Names, on the Edit Local DNS Server tab, under the **Top Names Settings** section, find the *top-names* attribute, enable it by selecting the **enabled** option, and then click **Save** to save the changes.

Top Names Statistics

The Top Names tab displays the relevant information with respect to top N domains and other important statistics attributes.

Local Basic or Advanced Web UI

-
- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page.
 - Step 2** Select **DNS** in the Manage Servers pane.
 - Step 3** Click the **Top Names** tab available in the Local DNS Server page.
-

CLI Commands

Use `dns getStats top-names` to view the Top Names statistics.

Specifying TLS Settings

Cisco Prime Network Registrar 11.0 supports TLS in the Authoritative DNS server in addition to the Caching DNS server. The DNS server listens on configurable port 853 for TLS. On port 853, only TCP TLS connections are accepted and other connections are dropped. The DNS server has configurable parameters to enable or disable TLS, and to add TLS private and public key files.

For more information on DNS over TLS, see the [Specifying TLS Settings, on page 34](#) section in the "Managing Caching DNS Server" chapter.



Note

- Cisco Prime Network Registrar does not support a command for generating self-signed certificates. However, they can be generated using readily available command line tool like openssl. For example:

```
# openssl req -new -x509 -days 365 -nodes -out public.pem -keyout private.pem
```
- TLS is not supported in hybrid mode and in zone transfers.
- TLS keys are not supported with password phrase.

Table 33: TLS Attributes in the Authoritative DNS Server

Attribute	Description
TLS (<i>tls</i>)	Enables or disables TLS support for DNS. Before enabling TLS, the private key files must be placed in the DNS data directory under <i>dns/tls</i> and the <i>tls-service-key</i> attribute be set. If using managed DNS certificates, the certificate settings will be automatically set. Otherwise, the public certificate file must be placed in the DNS data directory under <i>dns/tls</i> and the <i>tls-service-pem</i> attribute be set. Enabling or disabling TLS service requires a Cisco Prime Network Registrar service restart for the change to take effect.
TLS Port (<i>tls-port</i>)	The port number on which to provide TCP TLS service. The DNS server will not serve non-TLS queries on this port.
TLS Private Key File (<i>tls-service-key</i>)	Defines the file name which contains the private key to be used by DNS for TLS sessions. The file must be in the DNS data directory under the tls subdirectory (that is, <i><cnr.datadir>/dns/tls</i>) and must not be encrypted with a passcode.
TLS Public Key File (<i>tls-service-pem</i>)	Defines the pem file name which contains the public key certificate to be used by DNS for TLS sessions. The file must be in the DNS data directory under the tls subdirectory (that is, <i><cnr.datadir>/dns/tls</i>). Note that if using managed DNS certificates, this attribute will be ignored and should be left unset.

Local Advanced Web UI

To enable TLS support for the Authoritative DNS server, do the following:

Before you begin

Before enabling TLS, you must place the public certificate and private key files in the DNS data directory under the **tls** subdirectory (that is, *<cnr.datadir>/dns/tls*) and set the *tls-service-key* and *tls-service-pem* attributes under the **TLS Settings** section on the Manage DNS Authoritative Server page. You can also use the managed certificates (see the "Certificate Management" section in *Cisco Prime Network Registrar 11.0 Administration Guide*).

- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page. Click **DNS** on the Manage Servers pane.
- Step 2** On the Edit Local DNS Server tab, under the **TLS Settings** section, enable the *TLS* attribute by selecting the **enabled** option.
- Step 3** Click **Save** to save the changes.



Note You must restart the Cisco Prime Network Registrar service whenever TLS settings are modified.

CLI Commands

Use the following commands to enable TLS support for the Authoritative DNS server:

```
nrcmd> dns enable tls
```

Then, restart the Cisco Prime Network Registrar service using the following command:

```
# systemctl restart nwreglocal.service
```

Use **dns set attribute=value** to set the TLS attributes in the Authoritative DNS server.



Note You must restart the Cisco Prime Network Registrar service whenever TLS settings are modified.

TLS Statistics

On the Manage DNS Authoritative Server page, click the **Statistics** tab to view the Server Statistics page. The TLS statistics appear under the **Security Statistics** section of both the Total Statistics and Sample Statistics categories.

Table 34: TLS Statistics Attributes

Attribute	Description
<i>tls-queries</i>	Total number of queries received over TLS by the DNS Server.
<i>tls-queries-failed</i>	Total number of TLS queries failed during TLS handshake.

Enabling Round-Robin

A query might return multiple A or AAAA records for a name lookup. To compensate for most DNS clients starting with, and limiting their use to, the first record in the list, *round-robin* is enabled to share the load. This ensures that successive clients resolving the same name will connect to different addresses on a revolving basis. The DNS server then rearranges the order of the records each time it is queried. It is a method of load sharing, rather than load balancing, which is based on the actual load on the server.

Local Basic or Advanced Web UI

On the Manage DNS Authoritative Server page, under the Miscellaneous Options and Settings section, find the Enable round-robin (*round-robin*) attribute. It is set to enabled by default in Basic mode.

CLI Commands

Use **dns get round-robin** to see if round-robin is enabled (it is by default). If not, use **dns enable round-robin**.

Enabling Weighted Round-Robin

When a nameset is configured with multiple RRs of the same type, a weighted round-robin algorithm can be used to determine the frequency with which an RR is the first RR in the query response. To control the response behavior, administrators must be able to set weighted values on these RRs. In addition, the order in which multiple records are returned may be used by client applications and need to be controlled by administrators.

The *order* and *weight* attributes are available in Advanced mode.

Order

The *order* attribute specifies the sort order for the RR, compared to other RRs of the same type in the nameset. RRs with same type will be listed in ascending order, this will also be the order that RRs are returned when queried.

Weight

RR weight can be used in situations where you want certain servers providing the same service to be returned more frequently and therefore get more of the load. The *weight* attribute specifies the relative importance of this RR, compared to other RRs of the same type in the nameset. RRs with higher weight will be used more often in query responses for the name and type. For example, if *weight* for the RR is set to 5 and *weight* for another RR is set to 1, then RR will be used 5 times before the other RR is used once. RRs with a *weight* of 0 (zero) are always listed last and not included in the round robin operation.



Note The default *weight* on RRs is 1. When round robin is enabled (either DNS server or zone level), the RRs are returned in the first position once for each query (that is, traditional round robin).

If all the weights on the RR set are 0, then the response is returned to the client based on *order*. Effectively disabling round-robin on the RR set level.

The *order* and *weight* attributes can only be set on primary zones. These are transferred to HA backup and to the secondary servers, these attributes are not transferred when one of the servers in HA or secondary server is prior to 9.0 cluster. If you wish not to transfer order and weight, then disable the Transfer RR Meta Data (*xfer-rr-meta-data*) attribute present in the Manage DNS Authoritative Server page (you must do this in secondary DNS server). In secondary zone, *order* and *weight* are available, and the "resource records" are non-editable.

Local Basic or Advanced Web UI

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

Step 2 In the Forward Zone or Reverse Zone pane, click the zone name to open the Edit Zone page.

- Step 3** Click the **Resource Records** tab.
- Step 4** Add the RR name, TTL (if not using the default TTL), type, and data as appropriate.
- Step 5** Once the RRs are created, *order* and *weight* can be set by editing the RRs (click the pencil icon next to the desired RR). You can find the *order* and *weight* attributes under the **RR Settings** section.

CLI Commands

Use **zone name addRR** *rr-name rr-type rr-ttl rr-data* [**weight=rr-weight**] [**order=rr-order**] to set weight and order.

Use **zone name modifyRR** *rr-name type [data] attribute=value [attribute=value ...]* to modify the resource records.

Enabling Incremental Zone Transfers (IXFR)

Incremental Zone Transfer (IXFR, described in RFC 1995) allows only changed data to transfer between servers, which is especially useful in dynamic environments. IXFR works together with NOTIFY (see [Enabling NOTIFY, on page 96](#)) to ensure more efficient zone updates. IXFR is enabled by default.

Primary zone servers always provide IXFR. You should explicitly enable IXFR on the server (you cannot set it for the primary zone) only if the server has secondary zones. The DNS server setting applies to the secondary zone if there is no specific secondary zone setting.

Local Basic or Advanced Web UI

On the Manage DNS Authoritative Server page, under the Zone Default Settings section, you can find the Request incremental transfers (IXFR) attribute. It is set to enabled by default. For a secondary zone, you can also fine-tune the incremental zone transfers by setting the *ixfr-expire-interval* attribute.

This value is the longest interval the server uses to maintain a secondary zone solely from IXFRs before forcing a full zone transfer (AXFR). The preset value is 0, as we always use IXFR and it is enabled, we don't periodically change to AXFR. Then, click **Save**.

CLI Commands

Use **dns enable ixfr-enable**. By default, the *ixfr-enable* attribute is enabled.

Restricting Zone Queries

You can restrict clients to query only certain zones based on an Access Control List (ACL). An ACL can contain source IP addresses, network addresses, TSIG keys (see the *"Transaction Security"* section in *Cisco Prime Network Registrar 11.0 DHCP User Guide*), or other ACLs. The *restrict-query-acl* attribute on the Manage DNS Authoritative Server page serves as a default value for zones that do not have the *restrict-query-acl* explicitly set.

Enabling NOTIFY

The NOTIFY protocol, described in RFC 1996, lets the Cisco Prime Network Registrar DNS primary server inform its secondaries that zone changes occurred. The NOTIFY packets also include the current SOA record

for the zone giving the secondaries a hint as to whether or not changes have occurred. In this case, the serial number would be different. Use NOTIFY in environments where the namespace is relatively dynamic.

Since a zone primary server cannot know specifically which secondary server transfers from it, Cisco Prime Network Registrar notifies all nameservers listed in the zone NS records. The only exception is the server named in the SOA field of the primary server. You can add additional servers to be notified by adding the IPv4 and IPv6 addresses to the *notify-list* on the zone configuration.



Note In order for notifies to be sent to hidden name servers (that is, those that are not listed as NS RRs in the zone), their IP addresses need to be listed in the *notify-list* and notify setting needs to be set to *notify-list* or *notify-all*.

You can use IXFR and NOTIFY together, but this is not necessary. You can disable NOTIFY for a quickly changing zone for which immediate updates on all secondaries does not warrant the constant NOTIFY traffic. Such a zone might benefit from having a short refresh time and a disabled NOTIFY.

Local Advanced Web UI

- Step 1** On the Manage DNS Authoritative Server page, under the **Zone Transfer Settings** section, find the *notify* attribute and select the value from the drop-down list.
- Step 2** Set any of the other NOTIFY attributes (*notify-min-interval*, *notify-rcv-interval*, *notify-send-stagger*, *notify-source-port*, and *notify-wait*).
- Step 3** Click **Save**.
- Step 4** To add nameservers in addition to those specified in NS records, from the **Design** menu, choose **Forward Zones** or **Reverse Zones** or **Secondary Zones** under the **Auth DNS** submenu.
- Step 5** Click the zone name in the Forward Zones or Reverse Zones or Secondary Zones pane to open the Edit Zones page.
- Step 6** Add a comma-separated list of IP addresses of the servers using the *notify-list* attribute on the Edit Zone page.
- Step 7** Select the value from the *notify* drop-down list.
- Step 8** Click **Save**.

CLI Commands

Use `dns set notify=value`. NOTIFY is enabled by default. You can also enable NOTIFY at the zone level, where you can use `zone name set notify-list` to specify an additional comma-separated list of servers to notify beyond those specified in NS records.

Blocking Recursive Queries from Authoritative Server

Blocking recursive queries allows the server to not spend resources trying to process these queries. The Drop Recursive Queries (*drop-recursive-queries*) attribute controls whether the DNS server accepts or drops the queries which have RD flag on. When this attribute is enabled, recursive queries will be dropped by the server. The default value of *drop-recursive-queries* is disabled, which means that no recursive queries will be dropped.

To enable *drop-recursive-queries*, do the following:

Local Advanced Web UI

- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page. Click **DNS** on the Manage Servers pane.
- Step 2** On the Edit Local DNS Server tab, under the **Query Settings** section, enable the *drop-recursive-queries* attribute by selecting the **enabled** option.
- Step 3** Click **Save** to save the changes.



Note The setting can be changed dynamically without a DNS server reload.

CLI Command

Use **dns enable drop-recursive-queries** to enable Drop Recursive Queries.

Drop Recursive Queries Statistics

On the Manage DNS Authoritative Server page, click the **Statistics** tab to view the *queries-dropped-recursive* statistic attribute under the **Query Statistics** section. This indicates the number of queries dropped due to recursion. The *queries-dropped* counter will be incremented when recursive queries are dropped.

Running DNS Authoritative Server Commands

Access the commands by using the Commands button. Clicking the **Commands** button opens the DNS Commands dialog box in the local web UI. Each command has its own Run icon (click it, then close the dialog box):

- **Force all zone transfers**—A secondary server periodically contacts its primary server for changes. See [Enabling Zone Transfers, on page 143](#).
- **Scavenge all zones**—Periodically purges stale records. See the *"Scavenging Dynamic Records" section in Cisco Prime Network Registrar 11.0 DHCP User Guide*.
- **Synchronize All HA Zones**—Synchronizes all the HA zones. You have the option to choose the type of synchronization. The **Push All Zones From Main to Backup** option is checked by default. You can override this by checking **Pull All Zones From Backup to Main** check box.



Note The **Synchronize All HA Zones** command is an Expert mode command which you can see only if the server is an HA main server. You cannot see this command if it is an HA backup server. You can also, synchronize zones separately, which you can do from the Zone Commands for Zone page (see [Synchronizing HA DNS Zones, on page 126](#)).



Note If you find a server error, investigate the server log file for a configuration error, correct the error, return to this page, and refresh the page.

Configuring DNS Server Network Interfaces

You can configure the network interfaces for the DNS server from the Manage Servers page in the local web UI.

Local Advanced Web UI

-
- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page.
 - Step 2** Click **DNS** on the Manage Servers pane to open the Local DNS Server page.
 - Step 3** Click the **Network Interfaces** tab to view the available network interfaces that you can configure for the server. By default, the server uses all of them.
 - Step 4** To configure an interface, click the Configure icon in the Configure column for the interface. This adds the interface to the Configured Interfaces table, where you can edit or delete it.
 - Step 5** Clicking the name of the configured interface opens a new page, where you can change the address of the interface.
 - Step 6** Click **Modify Interface** when you are done editing, then click **Go to Server Interfaces** to return to the Manage Servers page.
- Note** The IPv6 functionality in DNS requires IPv4 interfaces to be configured except if the DNS server is isolated and standalone (it is its own root and is authoritative for all queries).
-

Managing Authoritative DNSSEC

DNSSEC provides origin authority, data integrity, and authenticated denial of existence. With DNSSEC, the DNS protocol is much less susceptible to certain types of attacks, particularly DNS spoofing attacks. DNSSEC provides protection against malicious or forged answers by adding digital signatures into DNS data, so each DNS response can be verified for integrity and authenticity.

Cisco Prime Network Registrar 9.0 and earlier Authoritative DNS Servers do not support signing of zones. From Cisco Prime Network Registrar 10.0, Authoritative DNSSEC support adds authentication and integrity to DNS zones. With this support, Cisco Prime Network Registrar DNS server is able to support both secure and unsecure zones.

To add DNSSEC Security:

1. Choose regional or local management of DNSSEC keys and zones.
2. Review the algorithm, size, lifetime, and intervals set for Authoritative DNSSEC that will be used for default key generation.
3. Create Zone Signing and Key Signing keys if not using internally generated keys.
4. Enable DNSSEC for the required zones.

- Export the DS RR for the signed zone which must be added to the parent zone, if it is not configured on the same server.

Enabling Authoritative DNSSEC

DNSSEC is enabled by default on the Authoritative DNS Server. It can be disabled by using the DNSSEC (*dnssec*) attribute (available in Expert mode) in the Manage Authoritative DNSSEC page. Disabling this attribute will disable zone signing for all zones, regardless of the zone *dnssec* attribute. By default, zone signing is disabled for all zones. To enable zone signing, the DNSSEC (*dnssec*) attribute in the zone configuration must be enabled only after the zone has been published. Once DNSSEC is enabled on the zone, zone signing is performed using core keys by default, or tenant keys specific to the zone tenant, if defined. The CCM server will create new keys for the zones, if there are no keys available.



Note DNSSEC cannot be enabled on a zone if RPZ is enabled and vice versa.

Table 35: Authoritative DNSSEC Attributes

Attribute	Description
Name	Specifies the name of authoritative DNSSEC configuration.
Description	A description of the authoritative DNSSEC configuration.
Key Rollover (<i>key-rollover</i>)	Specifies whether the regional or local cluster should perform Zone Signing Key (ZSK) rollover. If using regional zone management, this setting should be set to regional in order to centrally manage key generation and rollover.

Table 36: Zone Signing Key Attributes

Attribute	Description
Algorithm (<i>zsk-algorithm</i>)	Specifies the cryptographic algorithm to be used for the ZSK. DSA : DSA/RSA-1, value: 3, range: 512-1024 RSASHA1 : RSA/SHA-1, value: 5, range: 512-2048 RSASHA256 : RSA/SHA-256, value: 8, range: 512-2048 RSASHA512 : RSA/SHA-512, value: 10, range: 512-2048
Signature Size (<i>zsk-bits</i>)	Specifies the number of bits in the key and must be a multiple of 64. The value depends on the ZSK algorithm (<i>zsk-algorithm</i>) chosen. DSA : DSA/RSA-1, value: 3, range: 512-1024 RSASHA1 : RSA/SHA-1, value: 5, range: 512-2048 RSASHA256 : RSA/SHA-256, value: 8, range: 512-2048 RSASHA512 : RSA/SHA-512, value: 10, range: 512-2048

Key Lifetime (<i>zsk-lifetime</i>)	Specifies the lifetime of a ZSK. This defines the time interval where the key is used to sign zones. It is used to determine the deactivation-date when a ZSK key is created. The configured value MUST be greater than the <i>zsk-rollover-interval</i> . A value that is 10 times greater is recommended.
Key Rollover Interval (<i>zsk-rollover-interval</i>)	Specifies the time interval for the ZSK rollover process. It determines the lead time for the new key prior to the current key deactivation-date. Configured interval should be more than maximum TTL of the zones plus the propagation delay, to avoid bogus zone information.

Table 37: Key Signing Key Attributes

Attribute	Description
Algorithm (<i>sk-algorithm</i>)	Specifies the cryptographic algorithm to be used for the Key Signing Key (KSK). DSA : DSA/RSA-1, value: 3, range: 512-1024 RSASHA1 : RSA/SHA-1, value: 5, range: 512-2048 RSASHA256 : RSA/SHA-256, value: 8, range: 512-2048 RSASHA512 : RSA/SHA-512, value: 10, range: 512-2048
Signature Size (<i>sk-bits</i>)	Specifies the number of bits in the key and must be a multiple of 64. The value depends on the KSK algorithm (<i>sk-algorithm</i>) chosen. DSA : DSA/RSA-1, value: 3, range: 512-1024 RSASHA1 : RSA/SHA-1, value: 5, range: 512-2048 RSASHA256 : RSA/SHA-256, value: 8, range: 512-2048 RSASHA512 : RSA/SHA-512, value: 10, range: 512-2048
Key Rollover Interval (<i>sk-rollover-interval</i>)	Specifies the time interval for the KSK rollover process. It determines the lead time for the new key prior to the current key deactivation-date.

Local Advanced Web UI

- Step 1** From the **Design** menu, choose **Authoritative DNSSEC** under the **Security** submenu to open the Manage Authoritative DNSSEC page.
- Step 2** Modify the attributes in the **Zone Signing Key** and **Key Signing Key** sections as per your requirements.
- Step 3** Click **Save** to save your settings.

CLI Commands

Use `dnssec set attribute=value [attribute=value...]` to configure DNSSEC processing in the Authoritative DNS server. For example:

```
nrcmd> dnssec set zsk-algorithm=RSASHA1
```

Use **zone *zonename* signZone** to enable DNSSEC for the zone and add signatures for all RRs of the zone, when executed in Expert mode.

When connected to a regional cluster, you can use the following pull and push commands. Push allows a list of clusters or "all".

dnssec pull *cluster-name* [-report-only | -report]

dnssec push *cluster-list* [-report-only | -report]

Managing Authoritative DNSSEC Keys

To configure DNSSEC protected zones, a key must first be created. The zone is then signed using the key. You can create a key manually to customize the key attributes. Otherwise, the CCM server will create new keys automatically, as needed.

The *key-rollover* attribute in the Authoritative DNSSEC page can be set to local or regional management. The default is local. The *key-rollover* attribute specifies whether the regional or local cluster should perform ZSK rollover. With local rollover management, keys are managed on the local primary or HA main. The keys are copied to the HA backup via CCM HA sync. If zones are distributed across several primary servers, there will be many more keys to manage. With regional rollover management, keys are managed on the regional server and pushed to the local clusters. This lets you manage a common set of keys for your distributed primary servers. With central zone management, you can also stage zone edits and pre-sign zones before synchronizing the changes with the local DNS servers. Keys are auto-synched from regional to local when DNS edit mode is set to synchronous in the regional CCM server.

Rollover of ZSK is an automated process. Rollover of KSK has to be performed manually, the **rollover-ksk** command is used to start the KSK rollover process. You can provide your own key or allow CCM to generate keys.

dns rollover-ksk [*tenant-id=value*] [*next-key=keyname* | *key-group=value*]



Note In a lab setting, you can use the Expert mode command **zone *name* removeSignature** to remove all signature RRs and disable DNSSEC for the zone. This command should not be used for operational DNSSEC zones. Operational DNSSEC zones that will no longer be signed need to let signature records expire before they are deleted, following the guidelines in RFC 6781 - DNSSEC Operational Practices, Version 2.

Table 38: Key Timelines Attributes

Attribute	Description
Activation Date (<i>activation-date</i>)	Specifies the activation date and time for this key. Beginning at this date and time, the key will be used to sign RR sets.
Deactivation Date (<i>deactivation-date</i>)	Specifies the deactivation date and time for this key. Until this date and time, the key will be used to sign RR sets. This attribute must be 0 for KSKs. KSKs remain active until the key rollover process is initiated.

Removal Date (<i>expiration-date</i>)	Specifies the date and time this ZSK is scheduled to be removed. If 0, automatic removal is disabled and the key must be deleted by user action. This attribute must be 0 for KSKs. KSKs remain active until the key rollover process is initiated. When the rollover process is complete, the key can be deleted by user action.
Rollover Due Date (<i>rollover-due-date</i>)	Specifies the date and time, when this key should be (or was) rolled over. This transient attribute is used only for reporting.
Key Status (<i>status</i>)	Specifies the current status of the key. This transient attribute is used only for reporting.

Local and Regional Advanced Web UI

- Step 1** From the **Design** menu, choose **Auth DNSSEC Keys** under the **Security** submenu to open the List/Add Authoritative DNSSEC Keys page.
- Step 2** Set the *enable-signing* attribute value to **true** to enable the key and to sign the zones.
- Step 3** In the **Key Timelines** section, you can enter the deactivation date and removal date as required.
- Step 4** Click **Save** to save your settings.

CLI Commands

Use the following **dnssec-key** commands to create and manage Authoritative DNSSEC keys for zone signing.

dnssec-key name create [*attribute=value...*]

dnssec-key name delete [-force]

dnssec-key name show

dnssec-key name set attribute=value [*attribute=value...*]

Use **dnssec-key getStatus** to check the current status of DNSSEC keys related to rollover process.

When connected to a regional cluster, you can use the following pull, push, and reclaim commands. For push and reclaim, a list of clusters or "all" may be specified.

dnssec-key < name | all > pull < replace | exact > cluster-name [-report-only | -report]

dnssec-key < name | all > push < replace | exact > cluster-list [-report-only | -report]

dnssec-key name reclaim cluster-list [-report-only | -report]

Exporting DS Record

Export Delegation Signer (DS) record is available for the DNSSEC enabled zones. If the parent zone is found on the Authoritative DNS server, the DS record will be added to the zone automatically. If multiple authoritative servers are deployed, and the parent zone is on another local cluster, you can manage the zones on the regional server to automatically update the parent zone. If the parent zone is externally-owned, you must provide the DS record to be added by the external organization.

Local and Regional Advanced Web UI

To export DS record, do the following:

-
- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the Edit Zone page.
 - Step 2** On the Edit Zone page, under the **DNSSEC Settings** section, set the **DNSSEC** value to **true** to enable the DNSSEC.
 - Step 3** Click **Save** to save your settings.
 - Step 4** Click the **save** icon available next to the **DS Record** to export DS record.
-

CLI Commands

After you export the DS record, you need to publish the same to parent zone using the **export dnssec-ds zonename filename** command.

Setting Advanced Authoritative DNS Server Properties

You can set these advanced server properties:

- **SOA time-to-live**—See [Setting SOA Time to Live](#), on page 104
- **Secondary server attributes**—See [Setting Secondary Refresh Times](#), on page 105
- **Port numbers**—See [Setting Local and External Port Numbers](#), on page 106
- **Handle Malicious DNS Clients**—See [Handling Malicious DNS Clients](#), on page 106

Setting SOA Time to Live

The SOA record TTL is usually determined by the zone default TTL. However, you can explicitly set the SOA TTL, which sets the maximum number of seconds a server can cache the SOA record data. For example, if the SOA TTL is set for 3600 seconds (one hour), an external server must remove the SOA record from its cache after an hour and then query your nameserver again.

Cisco Prime Network Registrar responds to authoritative queries with an explicit TTL value. If there is no explicit TTL value, it uses the default TTL for the zone, as set by the value of the *defttl* zone attribute.

Normally, Cisco Prime Network Registrar assumes the default TTL when responding with a zone transfer with RRs that do not have explicit TTL values. If the default TTL value for the zone is administratively altered, Cisco Prime Network Registrar automatically forces a full zone transfer to any secondary DNS server requesting a zone transfer.

Local and Regional Web UI

-
- Step 1** On the List/Add Zones page, set the Zone Default TTL attribute value, which defaults to 24 hours.
 - Step 2** If you want, set the SOA TTL attribute value, which is the TTL for the SOA records only. It defaults to the Zone Default TTL value.
 - Step 3** You can also set a TTL value specifically for the NS records of the zone. Set the NS TTL attribute value under Nameservers. This value also defaults to the Zone Default TTL attribute value.

Step 4 Click **Save**.

CLI Commands

Use `zone name set defttl`.

Setting Secondary Refresh Times

The secondary refresh time is how often a secondary server communicates with its primary about the potential need for a zone transfer. A good range is from an hour to a day, depending on how often you expect to change zone data.

If you use NOTIFY, you can set the refresh time to a larger value without causing long delays between transfers, because NOTIFY forces the secondary servers to notice when the primary data changes. For details about NOTIFY, see [Enabling NOTIFY, on page 96](#).

Local and Regional Web UI

On the List/Add Zones page, set the Secondary Refresh field to the refresh time, which defaults to 3 hours. Make any other changes, then click **Save**.

CLI Commands

Use `zone name set refresh`. The default value is 10800 seconds (3 hours).

Setting Secondary Retry Times

The DNS server uses the secondary retry time between successive failures of a zone transfer. If the refresh interval expires and an attempt to poll for a zone transfer fails, the server continues to retry until it succeeds. A good value is between one-third and one-tenth of the refresh time. The default value is 60 minutes.

Local and Regional Web UI

On the List/Add Zones page, set the Secondary Retry field to the retry time, which defaults to one hour. Make any other changes, then click **Save**.

CLI Commands

Use `zone name set retry`. The default value is 60 minutes.

Setting Secondary Expiration Times

The secondary expiration time is the longest time a secondary server can claim authority for zone data when responding to queries after it cannot receive zone updates during a zone transfer. Set this to a large number that provides enough time to survive extended primary server failure. The default value is seven days (1 week).

Local and Regional Web UI

On the List/Add Zones page, set the Secondary Expire field to the expiration time, which defaults to seven days (1 week). Make any other changes, then click **Save**.

CLI Commands

Use **zone name set expire**. The default value is seven days (1 week).

Setting Local and External Port Numbers

If you are experimenting with a new group of nameservers, you might want to use non-standard ports for answering requests and asking for remote data. The local port and external port settings control the TCP and UDP ports on which the server listens for name resolution requests, and to which port it connects when making requests to other nameservers. The standard value for both is port 53. If you change these values during normal operation, the server will appear to be unavailable.

To see the full list of default ports, see the "*Default Ports for Cisco Prime Network Registrar Services*" section in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Local Advanced Web UI

On the Manage DNS Authoritative Server page, under the Network Settings section, set the Listening Port (*local-port-num*) and Remote DNS Servers Port (*remote-port-num*) attributes to the desired values (they both have default values of 53), then click **Save**.

Handling Malicious DNS Clients

When trying to resolve query requests, DNS servers may encounter malicious DNS clients. A client may flood the network with suspicious DNS requests. This affects the performance of the local DNS server and remote nameservers.

Using Cisco Prime Network Registrar, you can resolve this problem by barring malicious clients. You can configure a global ACL of malicious clients that are to be barred, using the *acl-blocklist* attribute.

Local Advanced Web UI

On the Manage DNS Authoritative Server page, expand the **Advanced Settings** section to view various attributes and their values. For the *acl-blocklist* attribute, enter the value (for example, 10.77.240.73). Then click **Save**.

Tuning DNS Properties

Here are some tips to tune some of the DNS server properties:

- **NOTIFY send min. interval DNS server attribute (*notify-min-interval*)**—Minimum interval required before sending notification of consecutive changes on the same zone to a server. The preset value is two seconds. For very large zones, you might want to increase this value to exceed the maximum time to send an outbound full zone transfer. This is recommended for secondary servers that receive inbound incremental zone transfers and send out full transfers to other secondaries. These include older BIND servers that do not support incremental zone transfers. Inbound incremental transfers may abort outbound full transfers.
- **NOTIFY delay between servers DNS server attribute (*notify-send-stagger*)**—Interval to stagger notification of multiple servers of a change. The preset value is one second, but you may want to raise it to up to five seconds if you need to support a large number of zone transfers distributed to multiple servers.

- **NOTIFY wait for more changes DNS server attribute (*notify-wait*)**—Time to delay, after an initial zone change, before sending change notification to other nameservers. The preset value is five seconds, but you may want to raise it to 15, for the same reason as given for the *notify-min-interval* attribute.
- **Maximum Memory Cache Size DNS server attribute (*mem-cache-size*)**—Size of the in-memory record cache, in kilobytes. The preset value is 500000 KB (500 MB) and this is used to make queries for Authoritative DNS server faster. The rule of thumb is to make it as large as the number of authoritative RRs.
- **EDNS Maximum Packet Size DNS server attribute (*edns-max-payload*)**— Specifies the sender's maximum UDP payload size, which is defined as the number of octets of the largest UDP packet that can be handled by a requestor. You can modify this attribute from a minimum of 512 bytes to a maximum of 4 KB. The default value for this attribute is 1232 bytes on the DNS server.

Running Caching DNS and Authoritative DNS on the Same Server

Cisco Prime Network Registrar includes a Hybrid DNS feature that allows you to run both the Caching DNS and Authoritative DNS servers on the same operating system without two separate virtual or physical machines. This feature allows the Caching DNS to auto-detect the Authoritative DNS server and its zones without creating DNS exceptions.



Note Cisco recommends that hybrid mode is only for smaller sized deployments. For larger deployments, Cisco recommends separating Caching DNS and Authoritative DNS on separate physical machines or VMs. For more information, see the "*Authoritative DNS Capacity and Performance Guidelines*" and "*Caching DNS Capacity and Performance Guidelines*" appendices in *Cisco Prime Network Registrar 11.0 Installation Guide*.



Note When you are in Hybrid mode configuration, SNMP queries to Cisco Prime Network Registrar will retrieve only the Caching DNS server static values and not the Authoritative DNS server static values.

Following prerequisites must be met for hybrid mode to work correctly:

- The local cluster must be licensed for both Caching DNS and Authoritative DNS servers.
- Caching DNS and Authoritative DNS servers must have their own configured unique and separate network interfaces. If there are no separate interfaces available and if only one interface is available, the loopback interface (127.0.0.1/8, ::1/128) should be configured on the Authoritative DNS server and the other interface (for example, eth0, eth1, ens192, and so on) should be configured for the Caching DNS server.

Once the prerequisites have been met, hybrid mode can be enabled on the Authoritative DNS server.

When you enable hybrid mode, the following results occur:

1. Whenever the Authoritative DNS server is reloaded, it causes the Caching DNS server to be reloaded.
2. The Caching DNS server reads the Authoritative DNS servers interface list to detect which IP to send requests to.
3. The Caching DNS server auto detects all zones (forward, reverse, and secondary) and auto creates in-memory exceptions for those zones.

4. The Caching DNS server will not cache hybrid mode responses regardless of the RRs TTL value. This ensures that the responses it returns to clients reflect the most up-to-date information.

Local Advanced Web UI

Step 1 To configure the network interfaces on the Authoritative DNS and Caching DNS servers, do the following:

Note In Hybrid mode, the Caching DNS and Authoritative DNS servers must be configured with their own separate network interfaces. Using the loopback interface for Authoritative DNS server is supported only when the Authoritative DNS server does not require direct access for queries, notifies, or zone transfers.

- a. From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page.
- b. Click **DNS** in the Manage Servers pane.
- c. Click the **Network Interfaces** tab and configure the available network interfaces for DNS.

Note The loopback interface (127.0.0.1/8, ::1/128) should be configured on the Authoritative DNS server for the DNS hybrid mode.

- d. Click **CDNS** in the Manage Servers pane.
- e. Click the **Network Interfaces** tab and configure the available network interfaces for the Caching DNS server.

Step 2 To enable the hybrid mode configuration on the Authoritative DNS server, do the following:

- a. From the **Deploy** menu, choose **DNS Server** under the **DNS** submenu to open the Manage DNS Authoritative Server page.
- b. Enable the *hybrid-mode* and *hybrid-use-adns-addr*s attributes available under the **Hybrid Mode** section:
 - Select the **enabled** option for the Hybrid Mode (*hybrid-mode*) attribute.
 - Select the **true** option for the Hybrid Use ADNS Addresses (*hybrid-use-adns-addr*s) attribute.

Note When the *hybrid-use-adns-addr*s attribute is enabled, the Caching DNS server will setup hybrid exceptions to forward to the Authoritative DNS server via *hybrid-adns-addr*s. The *hybrid-adns-addr*s attribute defaults to the loopback address (127.0.0.1) which is the recommended interface for hybrid DNS communication. If the *hybrid-use-adns-addr*s attribute is disabled, the Caching DNS server will use all of the Authoritative DNS server's configured network interfaces.

The *hybrid-adns-addr*s attribute specifies a list of one or more IP addresses to use for hybrid mode communication. Note that these addresses should match one or more of the Authoritative DNS server's configured interfaces. If using addresses other than the default loopback address (127.0.0.1), it may be necessary to also configure these interfaces in the Caching DNS Server for outbound traffic.

Step 3 Reload the Authoritative DNS server to enable the hybrid mode configuration.

CLI Commands

Use `dns set hybrid-mode=enabled` to enable the hybrid mode configuration on the Authoritative DNS server. Use `dns set hybrid-use-adns-addr=true` to enable the `hybrid-use-adns-addr` attribute. Use `dns-interface name set attribute=value` or `cdns-interface name set attribute=value` to set the interfaces.

Troubleshooting DNS Servers

Useful troubleshooting hints and tools to diagnose the DNS server and ways to increase performance include:

- **Restoring a loopback zone**—A loopback zone is a reverse zone that enables a host to resolve the loopback address (127.0.0.1) to the name `localhost`. The loopback address is used by the host to enable it to direct network traffic to itself. You can configure a loopback zone manually or you can import it from an existing BIND zone file.
- **Listing the values of the DNS server attributes**—From the **Deploy** menu, choose **DNS Server** under the **DNS** submenu to open the Manage DNS Authoritative Server page in the web UI. In the CLI, use `dns show`.
- **Adjusting certain attribute values that could have inherited preset values from previous releases during an upgrade**— These preset values are probably not optimal for current systems and can cause performance issues. We strongly recommend that you to update the settings to use the new preset values. Example: The present value of Maximum Memory Cache Size DNS server attribute (`mem-cache-size`) is updated to 500 MB.

Be sure to reload the DNS server after saving the settings.

- **Choosing from the DNS log settings to give you greater control over existing log messages**—Use the Log Settings (`server-log-settings`) attribute on the Edit DNS Server page in the web UI, or `dns set server-log-settings=value` in the CLI, with one or more of these keyword or numeric values, separated by commas (see the table below). Restart the server if you make any changes to the log settings.

Table 39: DNS Log Settings

Log Setting	Description
activity-summary	This setting enables logging of DNS statistic messages at the interval specified by <code>activity-summary-interval</code> . The type of statistics logged can be controlled with <code>activity-counter-log-settings</code> and <code>activity-summary-type</code> .
config	This setting enables logging of DNS server configuration and de-initialization messages.
config-detail	This setting enables logging of detailed configuration messages (that is, detailed zone configuration logging).
dnssec	This setting enables log messages associated with DNSSEC processing.
host-health-check	This setting enables logging associated with DNS Host Health Check.
db	This setting enables logging of database processing messages. Enabling this flag provides insight into various events in the server's embedded databases.

Log Setting	Description
ha	This setting enables logging of HA DNS messages.
notify	This setting enables logging of messages associated with NOTIFY processing.
query	This setting enabled logging of messages associated with QUERY processing.
scavenge	This setting enables logging of DNS scavenging messages.
scp	This setting enabled logging associated with SCP messages handling.
server-operations	This setting enables logging of general server events, such as those pertaining to sockets and interfaces.
tsig	This setting enables logging of events associated Transaction Signature (TSIG).
update	This setting enables logging of DNS Update message processing.
xfr-in	This setting enables logging of inbound full and incremental zone transfers.
xfr-out	This setting enables logging of outbound full and incremental zone transfers.

- **Using the dig utility to troubleshoot DNS Server**—dig (domain information groper) is a flexible tool for interrogating DNS name servers. It performs DNS lookups and displays the answers that are returned from the name server(s) that were queried. Most DNS administrators use dig to troubleshoot DNS problems because of its flexibility, ease of use, and clarity of output. To obtain help for the **dig** utility, use **dig -h** or use **man dig**.
- **Using the nslookup utility to test and confirm the DNS configuration**—This utility is a simple resolver that sends queries to Internet nameservers. To obtain help for the **nslookup** utility, enter **help** at the prompt after you invoke the command. Use only fully qualified names with a trailing dot to ensure that the lookup is the intended one. An **nslookup** begins with a reverse query for the nameserver itself, which may fail if the server cannot resolve this due to its configuration. Use the **server** command, or specify the server on the command line, to ensure that you query the proper server. Use the **-debug**, or better yet, the **-d2**, flag to dump the responses and (with **-d2**) the queries being sent.

Although dig is normally used with command-line arguments, it also has a batch mode of operation for reading lookup requests from a file. Unlike earlier versions, the BIND9 implementation of dig allows multiple lookups to be issued from the command line. Unless you specifically query a specific name server, dig tries each of the servers listed in /etc/resolv.conf. When no command line arguments or options are given, dig performs an NS query for the root ".". A typical invocation of dig looks like: dig @server name type where server is the name or IP address of the name server to query.



CHAPTER 7

DNS Host Health Check

In Cisco Prime Network Registrar 9.0 and earlier, DNS replies to A/AAAA queries with the RRs in its authoritative configuration regardless of whether or not the destination addresses are reachable. The returned IP address may or may not be reachable at the time when the DNS query is made. This outage may not be known to the DNS servers, or to the DNS client. In Cisco Prime Network Registrar 9.1 and later, an authoritative DNS server can periodically check the availability of a host or set of hosts for which it operates as the DNS authority, by pinging the addresses using ICMP echo messages (ping). In Cisco Prime Network Registrar 10.0 and later, DNS host health check supports the GTP-C protocol echo message using UDP v4 and UDP v6 to find out host availability. Hosts which are identified as unavailable are not sent in the query reply. The server responds with all RRs in the RR Set for the first query, with TTL set as *hhc-max-init-ttl*. The DNS server sends the pings (ICMP ping or GTP-C echo ping) for RRs in an RR Set only after receiving a query for that RR, and then the subsequent A/AAAA queries will respond with the reachable RRs.



Note All RRs which have *host-health-check* attribute set to **ping** or **gtp-echo** are monitored periodically. Monitoring will start only after receiving the first query for RR with *host-health-check* set to **ping** or **gtp-echo**. When *host-health-check* is set to **ping**, ICMP protocol will be used for monitoring.

To make the feature work effectively, the pinged systems should have default security settings that allow ping response. When *host-health-check* is set to **gtp-echo**, GTP-C v2 protocol (GTP-C Echo request and response) will be used for monitoring.

- [DNS Host Health Check Configuration Settings, on page 111](#)
- [Enabling Host Health Check, on page 112](#)
- [Host Health Check RR Set Settings, on page 113](#)
- [Viewing DNS Host Health Check Statistics, on page 113](#)

DNS Host Health Check Configuration Settings

DNS Host Health Check comes with preconfigured settings, and is disabled by default on the DNS server.

Use the following DNS server level attributes to enable DNS Host Health Check:

Table 40: DNS Server Level Attributes

Attribute	Description
-----------	-------------

Host Health Check (<i>host-health-check</i>)	Enables or disables DNS Host Health Check in the DNS server. When Host Health Check is enabled, DNS server sends <i>hhc-max-ttl</i> as TTL in query reply for active RRs. When DNSSEC is enabled, DNS server will add RRs which are not active at the end of RR list in the query reply. When DNSSEC is not enabled, DNS server will not add RRs which are not active in RR list in the query reply. <i>host-health-check</i> is disabled on the DNS server by default. Reload the DNS server after enabling <i>host-health-check</i> .
Host Health Check Interval (<i>hhc-interval</i>)	Specifies the time interval (in seconds) to check RR Sets for reachability.
Max TTL (<i>hhc-max-ttl</i>)	Specifies the maximum TTL (in seconds) to send in query reply when RR health status is up. By default the <i>hhc-interval</i> value will be used. Note If the RR Set has a TTL less than <i>hhc-interval</i> or <i>hhc-max-ttl</i> , the RR Set's TTL will be used in the response.
Max Initial TTL (<i>hhc-max-init-ttl</i>)	Specifies the maximum initial TTL (in seconds) to send in query reply when Host Health Check RR is queried for the first time. Note If the RR Set has a TTL less than <i>hhc-max-init-ttl</i> , the RR Set's TTL will be used in the response.

Enabling Host Health Check

To enable DNS Host Health Check, do the following:

Local Advanced Web UI

-
- Step 1** On the Manage DNS Authoritative Server page, under the **Host Health Check** section, select the **enabled** option for the *host-health-check* attribute.
- Step 2** Click **Save** to save the changes and reload the Authoritative DNS server.
-

CLI Commands

Use the **dns enable host-health-check** to enable host health check and use **dns reload** to restart the DNS server.



Note Restart the DNS server to apply the configuration changes successfully.

Host Health Check RR Set Settings

Local Advanced Web UI

From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page and click the **Resource Records** tab. Click the RR name. Under the **RR Set Settings** section, select the value as **ping** from the **host-health-check** drop down list. This attribute change on the RR Set does not require a reload.



Note If DNSSEC is enabled on the zone, DNS server will add the RRs which are not active at the end of the RR list in the query reply.

CLI Commands

The **rrSet** command sets/unsets the *host-health-check* flag on resource records for the *rr-name*. When this flag is set, the A and AAAA record's health will be monitored.

```
zone name rrSet rr-name [set <host-health-check=off/ping/gtp-echo>] [get <host-health-check>] [unset <host-health-check>] [show]
```



Note DNS server supports Global Unicast Address for IPv6 host health monitoring.

Viewing DNS Host Health Check Statistics

You can view the DNS Host Health Check statistics in the following ways:

Local Advanced Web UI

On the Manage DNS Authoritative Server page, click the **Statistics** tab to view the Server Statistics page. The DNS Host Health Check statistics appear under the **Host Health Check Statistics** section of both the Total Statistics and Sample Statistics categories.

Table 41: DNS Host Health Check Statistics Attributes

Attribute	Description
<i>hhc-domains</i>	Reports the number of domains checked for ping and gtp-echo Host Health Check.
<i>hhc-domains-failed</i>	Reports the number of domains check failed for ping and gtp-echo Host Health Check. When all the RRs in the RR set are down, this stat is incremented.

<i>hhc-domains-passed</i>	Reports the number of domains check passed for ping and gtp-echo Host Health Check. When any A/AAAA RR in the RR set is up, this stat is incremented.
<i>hhc-rrs</i>	Reports the number of RRs checked for ping and gtp-echo Host Health Check.
<i>hhc-rrs-passed</i>	Reports the number of RRs that have passed ping and gtp-echo health check.
<i>hhc-rrs-failed</i>	Reports the number of RRs that have failed ping and gtp-echo health check.
<i>hhc-ping-domains</i>	Reports the number of domains checked for ping Host Health Check.
<i>hhc-ping-domains-failed</i>	Reports the number of domains check failed for ping Host Health Check. When all the RRs in the RR set are down, this stat is incremented.
<i>hhc-ping-domains-passed</i>	Reports the number of domains check passed for ping Host Health Check. When any RR in the RR set is up, this stat is incremented.
<i>hhc-ping-rrs</i>	Reports the number of RRs checked for ping Host Health Check.
<i>hhc-ping-rrs-failed</i>	Reports the number of RRs that have failed ping Host Health Check.
<i>hhc-ping-rrs-passed</i>	Reports the number of RRs that have passed ping Host Health Check.
<i>hhc-gtp-echo-domains</i>	Reports the number of domains checked for gtp-echo Host Health Check.
<i>hhc-gtp-echo-domains-failed</i>	Reports the number of domains check failed for gtp-echo Host Health Check. When all the RRs in the RR set are down, this stat is incremented.
<i>hhc-gtp-echo-domains-passed</i>	Reports the number of domains check passed for gtp-echo Host Health Check. When any RR in the RR set is up, this stat is incremented.
<i>hhc-gtp-echo-rrs</i>	Reports the number of RRs checked for gtp-echo Host Health Check.
<i>hhc-gtp-echo-rrs-passed</i>	Reports the number of RRs that have passed gtp-echo Host Health Check.
<i>hhc-gtp-echo-rrs-failed</i>	Reports the number of RRs that have failed gtp-echo Host Health Check.

DNS Host Health Check statistics can also be logged in the server by enabling the **host-health-check** option present in the **Activity Summary Settings** section of the Edit Local DNS Server page.

CLI Commands

Use **dns getStats dns-hhc total** to view the host health check Total statistics and **dns getStats dns-hhc sample** to view the sampled counters statistics.



Note Restart the DNS Server to apply the configuration changes successfully.



CHAPTER 8

Managing DNS Firewall

- [Managing DNS Firewall, on page 115](#)

Managing DNS Firewall

DNS firewall controls the domain names, IP addresses, and name servers that are allowed to function on the network. This enables Internet Service Providers (ISP), enterprises, or organizations to define lists of FQDNs, IP addresses, subnets and prefixes of end nodes, and configure rules to secure the network by redirecting the resolution of DNS name away from known bad domains or non-existing domains (NXDOMAIN).

Every query to a Caching DNS server is first verified against the list of DNS firewall rules in the order of priority. To ensure that the Caching DNS server redirects queries for non-existing or known bad domains, you can create DNS firewall rules. The DNS firewall rule comprises of a priority, an ACL, an action, and a list of domains and takes precedence over exceptions and forwarders. You can configure the following actions for these queries:

- **Drop**—Drops the resource record query.
- **Refuse**—Responds with no data and the REFUSED status.
- **Redirect**—Redirects A or AAAA queries to the specified IP address.
- **Redirect-nxdomain**—Redirects to a specific A or AAAA address if the queried domain does not exist.
- **RPZ**—Uses RPZ rules.

When the incoming query matches the DNS firewall rule, the specified action will be taken unless the rule is for redirect-nxdomain. A redirect-nxdomain rule takes effect only for incoming queries that would result in an NXDOMAIN response.



Note The firewall rules such as Drop, Refuse, Redirect, and the RPZ query-name trigger take place before regular query processing and therefore take precedence over forwarders and exceptions. The other actions and triggers are applied during or after regular query processing.

DNS RPZ Firewall Rules

Cisco Prime Network Registrar supports RPZ. The DNS firewall rules can be set up for specially designated zones on the Authoritative DNS server. The RPZ and RR data combined with DNS resolver effectively creates

a DNS firewall to prevent misuse of the DNS server. The RPZ firewall rule comprises of a trigger (query-name, ip-answers, ns-name, and ns-ip) and a corresponding action.

The RPZ firewall rules utilize both the Authoritative DNS and Caching DNS servers to provide the RPZ functionality. The Authoritative DNS server stores the data for RPZ and the rules, whereas the Caching DNS server takes the client queries and applies these rules.

DNS RPZ Zones

We recommend that you create a separate forward zone on the Authoritative DNS server for RPZ. The zone can be either primary or secondary, and the data can either be manually entered or transferred from a third party RPZ provider. The zones can be named as **rpz.<customer-domain>** to avoid conflict with domain names in the Global DNS space. In the zone's **Query Settings** section, enable the *rpz* attribute to make it an RPZ zone.



Note If the RPZ comes via zone transfer, it must be named the same as at the source. If using a commercial RPZ provider, the name is specified by the provider.

The RPZ RR names can take the following forms:

Table 42: RPZ Triggers

RPZ Trigger	RR Name	Example	Example RR Name
Domain being queried	<domain>.rpz. <customer-domain>	Domain www.baddomain.com	www.baddomain.com.rpz.cisco.com
Name Server to query	<ns-domain-name>.rpz- nsdname.rpz.<customer-domain>	Name Server ns.baddomain.com	ns.baddomain.com.rpz-nsdname.rpz. cisco.com
Name Server IP to query	32.<reversed-ip>.rpz-nsip.rpz. <customer-domain>	Name Server Address 192.168.2.10	32.10.2.168.192.rpz-nsip.rpz.cisco.com
Name Server IP to query	32.<reversed-ip>.rpz-nsip.rpz. customer-domain>	Name Server Address 2001:db8:0:1::57	12857zz.10.db82001.rpz-nsip.rpz.cisco.com
A Records in Answer Section of Response	32.<reversed-ip>.rpz-ip.rpz. <customer-domain>	A answer record 192.168.2.10	32.10.2.168.192.rpz-ip.rpz.cisco.com
A Records in Answer Section of Response	<subnet-mask>.<reversed-ip>. rpz-ip.rpz.<customer-domain>	A answer record in subnet 192.168.2.0/24	24.0.2.168.192.rpz-ip.rpz.cisco.com

AAAA Records in Answer Section of Response	128.<reversed-ip>.rpz-ip.rpz.<customer-domain>	AAAA answer record 2001:db8:0:1::57	128.57.zz.1.0.db8.2001.rpz-ip.rpz.cisco.com
AAAA Records in Answer Section of Response	<prefix-length>.<reversed-ip>.rpz-ip.rpz.customer-domain>	AAAA answer record in prefix 2001:db8.0.1::/48	27.zz.1.0.db8.2001.rpz-ip.rpz.cisco.com

This zone contains all the RRs related to query names which are in block list. Blocking IP addresses and ranges must be done within the *rpz-ip* label (that is, *rpz-ip.rpz.cisco.com*). The same logic can be applied to blocking name servers using the *rpz-nsdname* and *rpz-nsip* labels.



Note *rpz-ip*, *rpz-nsdname*, and *rpz-nsip* are just another labels and are not real subdomains or separate zones. No delegation points will exist at this level and Caching DNS server relies on finding all the data within the referenced zone.



Note When using *rpz-nsdname* and *rpz-nsip*, the corresponding rule is applied to the original query and will therefore change the answer section. In cases when the final answer is determined from the RPZ rule(s), the RPZ zone SOA will be included in the authority section.

When the Caching DNS server is configured to use RPZ, it queries the Authoritative DNS server to lookup the RPZ rules. The Caching DNS server formulates the correct query name, interprets the query response as an RPZ rule, and applies the rule to the client query. If the RPZ rule causes Caching DNS server to rewrite the client response, this data is cached to make future lookups faster. The Caching DNS server RPZ configuration determines which RPZ trigger should be used. If no RPZ rule is found, the query proceeds normally.

In addition, RPZ overrides can be configured on the Caching DNS server. This enables the Caching DNS server to override the RPZ action returned by the Authoritative DNS server. This is useful when you do not have control over the Authoritative DNS data as is the case when the data is pulled from a third party. When the Caching DNS server gets a match from the Authoritative DNS server for the RPZ query, it performs the override action rather than the rule action specified in the RR data.

DNS RPZ Actions

RPZ rules are created using standard DNS RRs, mostly CNAME RRs. However, for redirecting, you can use any type of RR. The RR name follows the format based on the RPZ trigger as described in the [Table 42: RPZ Triggers, on page 116](#) section. The rdata defines the rule action to be taken. The following table describes the RPZ actions.

Table 43: RPZ Actions

RPZ Rule Action	RPZ RR RData	RPZ RR Example
-----------------	--------------	----------------

NXDOMAIN	CNAME .	www.baddomain.com.rpz.cisco.com. 300 CNAME .
NODATA	CNAME *.	www.baddomain.com.rpz.cisco.com. 300 CNAME *.
NO-OP (allowed list)	CNAME rpz-passthru. CNAME FQDN	www.gooddomain.com.rpz.cisco.com. 300 CNAME rpz-passthru. www.gooddomain.com.rpz.cisco.com. 300 CNAME www.gooddomain.com.
DROP	CNAME rpz-drop.	www.baddomain.com.rpz.cisco.com. 300 CNAME rpz-drop.
Redirect	<any RR type> <redirect-data>	www.wrongdomain.com.rpz.cisco.com. 300 CNAME walledgarden.cisco.com. www.baddomain.com.rpz.cisco.com. 300 A 192.168.2.10 www.baddomain.com.rpz.cisco.com. 300 AAAA 2001:db8:0:1::57

DNS RPZ Requirements and Best Practices

- All RPZ zones must have the *rpz* attribute enabled. A DNS reload is necessary for this change to take effect.
- Both Cisco Prime Network Registrar Authoritative DNS and Caching DNS must be used for end to end RPZ solutions.
- The *restrict-query-acl* on the RPZ zone must include only the Caching DNS address and localhost.
- Zone transfers (*restrict-xfer-acl*) must be either completely denied or restricted only to a specific set of servers.
- RPZ zone must not be delegated from the parent zone. It must be hidden and only available to a specially configured Caching DNS.
- There must be no RPZ nameserver address record to avoid caching and keeping the name server.
- The name server record must point to "localhost".
- The number of RPZ Firewall rules on a Caching DNS server should be limited to 2-3. The time to process a query increases linearly for each RPZ Firewall rule specified.
- The default TTL, for manually created RPZ zones, must reflect the rate of change in the zone data. The recommended rate ranges from 5m to 2h.
- The Caching DNS server must revise its *max-cache-ttl* setting to assure that the cached information is from a reliable source and can be trusted. This setting should be in line with the default TTL of 5m to 2h.
- The Authoritative DNS servers must enable NOTIFY, IXFR, AXFR, and TSIG for zone transfers of the distributed RPZ data.

- An RPZ zone may contain data for domains (allowed list or block list), but can also be separated into two distinct zones. This can be helpful when there is overlapping data or the block list zone is maintained by a third party (that is, RPZ subscription).

Setting Up RPZ Primary Zones on the Authoritative DNS Server

Local Basic or Advanced Web UI

- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
- Step 2** Click the **Add Forward Zone** icon in the **Forward Zones** pane to open the Add Zone dialog box
- Step 3** Enter the name of the zone (that is, **rpz.zonename**), specify **localhost** as the name server, add a contact E-mail, and a starting serial number.
- Step 4** Make the following changes in the Edit Zone page:
- a) Set the Zone Default TTL (recommended setting is between 5m and 2h).
 - b) Under the **Query Settings** section, set the *rpz* attribute to **true** and restrict queries using the *restrict-query-acl* attribute.
Note Queries should be restricted to localhost and the Caching DNS server address(es), **restrict-query-acl=localhost, cdns-address**.
 - c) Under the **Zone Transfer Settings** section, restrict zone transfers and notifies.
Note Zone transfers and notifies should only be allowed to other RPZ secondaries and localhost.
- Step 5** From the **Deploy** menu, choose **DNS Server** under the **DNS** submenu to open the Local DNS Server page.
- Step 6** Click the **Restart Server** icon to reload the DNS server and publish the RPZ zone.
-

CLI Commands

Use the following CLI commands:

- To create an RPZ zone, the zone name should indicate that it is an RPZ zone. For example, rpz.example.com.

```
nrcmd> zone rpz.example.com. create primary localhost admin
```
- Enable the RPZ zone attribute (*rpz*).

```
nrcmd> zone rpz.example.com. enable rpz
```
- Restrict queries to only be allowed from Caching DNS and localhost.

```
nrcmd> zone rpz.example.com. set restrict-query-acl="localhost, cdns-server"
```
- Restrict or completely deny zone transfers depending on deployment.

```
nrcmd> zone rpz.example.com. set restrict-xfer-acl=none
```
- Set the default TTL between 5m and 2h.

```
nrcmd> zone rpz.example.com. set defttl=5m
```
- Reload the DNS server to publish the RPZ zone and for the configuration changes to take effect.

```
nrcmd> dns reload
```

Setting Up DNS Firewall Rules

To add or edit DNS firewall rules:

Local and Regional Advanced Web UI

Step 1 From the **Design** menu, choose **DNS Firewall** under the **Cache DNS** submenu to open the List/Add DNS Firewall Rules page.

Step 2 Click the **Add DNS Firewall Rule** icon in the DNS Firewall pane to open the Add DNS Firewall dialog box.

Step 3 Enter a rule name in the Rule Name field and specify the action type.

Note The **drop** and **refuse** actions are applicable to all the queries for the specified domains, while the redirect and redirect-NXDOMAIN rules are applicable only to the queries of A and AAAA records.

Step 4 Click **Add DNS Firewall** to save the firewall rule. The List/Add DNS Firewall Rules page appears with the newly added firewall rule.

Note The rules with the action **refuse** do not use a domain or destination IP address.

Step 5 If you selected the **drop** or **redirect** action:

- Enter the ACL List, and click the **Add** icon to add the domains that need to be monitored for the drop or redirection.
- For the **redirect** action, you also need to enter the IPv4 Destination or IPv6 Destination.

Step 6 If you selected the **rpz** action:

a. Enter the RPZ Zone Name and the name of RPZ server.

Note The recommended RPZ zone name should be **rpz.customer-domain** to avoid conflicting with domain names in the Global DNS space.

b. Select the RPZ Trigger from the options and the corresponding override action.

Step 7 Click **Save** to save your settings, or click **Revert** to cancel the changes.

To delete a DNS firewall rule, select the rule on the DNS Firewall pane, click the **Delete** icon, and then confirm the deletion.

CLI Commands

Use **cdns-firewall rule-name create** to add the DNS firewall rules, separated by spaces.

Use **cdns-firewall list** to list the domains the domain redirect rule.

Use **cdns-firewall rule-name delete** to remove domain redirect rule.

Changing Priority of DNS Firewall Rules

When you create a set of DNS firewall rules, you can specify the priority in which order the rules will apply.



Note When using more than one DNS firewall rule, it is recommended to set the rules priority to control the order in which rules are processed. The lowest non-zero priority will be processed first. DNS firewall rules with a priority of 0 (default), will be processed last.

Local and Regional Advanced Web UI

To set the priority or reorder the rules:

-
- Step 1** From the **Design** menu, choose **DNS Firewall** under the **Cache DNS** submenu to open the List/Add DNS Firewall Rules page.
- Step 2** Click the **Reorder DNS Firewall Rules** icon in the DNS Firewall pane to open the Reorder dialog box.
- Step 3** Set the priority for the DNS firewall rules by either of the following methods:
- Select the rule and click the Move up or Move down icon to reorder the rules.
 - Select the rule and click the Move to button, and enter the row number to move the rule.
- Step 4** Click **Save** to save the reordered list.
-

CLI Commands

Use `cdns-firewall name set priority=value` to specify the rule priority relative to the other rules.

Enabling TLS for RPZ

In Cisco Prime Network Registrar 11.0, the Caching DNS Firewall RPZ action supports TLS for communication with the RPZ server.

Local and Regional Advanced Web UI

To enable TLS for the RPZ server, do the following:

-
- Step 1** From the **Design** menu, choose **DNS Firewall** under the **Cache DNS** submenu to open the List/Add DNS Firewall Rules page.
- Step 2** Enable the `rpz-tls` attribute by selecting the **enabled** option. If you enable this, you should configure a `tls-cert-bundle` to load the CA certificates, otherwise the connections cannot be authenticated.
- The `rpz-tls-auth-name` attribute defines the auth name for the RPZ server. If TLS is enabled, the Caching DNS server checks the TLS authentication certificates with that name sent by the RPZ server.
-

CLI Command

Use `cdns-firewall name set rpz-tls=true` to enable TLS for the RPZ server.



CHAPTER 9

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.



Note When running HA, we recommend having only primary zones on the server.

- [Introduction to HA DNS Processing, on page 123](#)
- [Creating High Availability DNS Pairs, on page 125](#)
- [Synchronizing HA DNS Zones, on page 126](#)
- [Enable Logging of HA DNS Information, on page 127](#)
- [Viewing HA DNS Statistics, on page 127](#)

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. 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 HA pair synchronize and exchange RRs that were changed or deleted during communications interrupted state.

Whenever you add a new zone, both the primary and backup 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. The backup ignores DNS updates, refuses RR edits, but processes RR Update messages from the main server. Scavenging is suspended on zones while they are still synchronizing.
- **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. The server continues listening for communication from the partner (they both send heartbeat messages at the rate specified by *ha-dns-poll-interval*) and tries to connect, meanwhile accepting DNS updates and RR edits and disabling scavenging.
- **Partner-Down**—It is similar to Communications-Interrupted, but does not continue to track RR changes. Once the partner returns, the entire zone will be sent to the partner. This allows for better performance and limits the disk space needed to track changes since the partner will get a copy of the zone when it becomes operational again.

When a DNS server starts up, it:

1. Opens its configured HA DNS listening ports and listens for connections from its partner.
2. Transitions to Negotiating state. In the Negotiating state, RR edits are not allowed.
3. Transitions to Normal state, the servers start synchronizing changes to each primary zone. The main starts allowing updates to zones and sending the update information to the backup.

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 Cisco Prime Network Registrar 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 11.0 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 starts sending updates to the backup. The DHCP server tries to contact the main DNS server, twice. It tries the backup partner if both 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 establish automatically and the servers will get into Normal state where they carry out zone synchronization and make sure that both have the same RRs, and so on.

If both the 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. In Expert mode, the following two options are provided:

- 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.

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 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 Pairs** under the **DNS** submenu to open the List/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** Select the cluster name of the main DNS server from the **main** drop-down list.
- Note** If you change the IP address (IPv4 or IPv6) of your local host machine, you must modify the localhost cluster (on the Edit Cluster page) to change the IP address (IPv4 or IPv6) in the IPv4 Address or IPv6 Address field. Do not set the value to 127.0.0.1 and ::1.
- Step 6** Select the cluster name of the backup DNS server from the **backup** drop-down list. This cannot be the same as the main server cluster. Set the *ha-dns-main-address* and *ha-dns-backup-address* attributes (for IPv4) and *ha-dns-main-ip6address* and *ha-dns-backup-ip6address* (for IPv6) 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:
- Select the HA in the HA Pairs pane and click the **Sync HA DNS Server Pair** tab.
 - Choose the direction of synchronization (Main to Backup or Backup to Main).
 - Choose the operation type (Update, Complete, or Exact). See the table on the page for details on the operations for each operation type.
 - Click the **Report** button to display the prospective synchronization changes on the View HA DNS Sync Report page.
 - Click **Run Complete** to complete the synchronization.
 - Click **Return** to return to the List/Add HA DNS Server Pair page.
- Step 9** Reload both DNS servers to begin HA communication.

CLI Commands

Create the HA DNS server pair (**ha-dns-pair name create main-cluster/address backup-cluster/address**). The *address* can be IPv4 or IPv6. 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 example-ha-pair create localhost test-cluster
nrcmd> ha-dns-pair example-ha-pair 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. The partner that was up will send all RRs to the server that was down.

Synchronizing HA DNS Zones

Local Advanced Web UI

To manually synchronize an HA DNS zone:

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

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 setting, *ha*, enables logging of HA DNS related information.

Local Basic or Advanced Web UI

On the Manage DNS Authoritative Server page, under the **Log Settings** section, check the **ha** check box. Click **Save** to save the changes.

CLI Command

Use **dns set server-log-settings=ha** to enable logging of HA DNS related information.

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.



CHAPTER 10

Managing Zones

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 [Introduction to the Domain Name System, on page 1](#).

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

DNS also supports creation of the Internationalized Domain Names (IDN). The full set of Unicode characters are supported to name DNS domains in the web UI, web-services (REST), and Java SDK with limited sort and search capabilities. For more information, see *Cisco Prime Network Registrar 11.0 Release Notes*.



Note You must set the locale parameters on UNIX to en_US.UTF-8 when running Java tools that use Java SDK, such as `cnr_rules`. For more information, see the *"Running Data Consistency Rules"* section in the *Cisco Prime Network Registrar 11.0 Administration Guide*.

This chapter describes the basics of configuring the Cisco Prime Network Registrar DNS servers, and their primary and secondary zones. [Managing Resource Records, on page 163](#) describes how to manage DNS resource records (RRs) and hosts, and [Managing Authoritative DNS Server, on page 65](#) describes how to set some of the more advanced zone and DNS server properties.

- [Managing Primary DNS Servers, on page 130](#)
- [Creating and Applying Zone Templates, on page 130](#)
- [Staged and Synchronous Modes, on page 132](#)
- [Configuring Primary Forward Zones, on page 133](#)
- [Configuring Primary Reverse Zones, on page 139](#)
- [Getting Zone Counts on the Server, on page 141](#)
- [Enabling DNS Updates, on page 142](#)
- [Managing Secondary Servers, on page 142](#)
- [Configuring Subzones, on page 144](#)
- [Managing Zone Distributions, on page 147](#)
- [Managing DNS ENUM Domain, on page 151](#)

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.



Note You do not need to create a loopback zone for the local host, because Cisco Prime 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.

Related Topics

[Configuring Primary Forward Zones, on page 133](#)

[Configuring Primary Reverse Zones, on page 139](#)

[Getting Zone Counts on the Server, on page 141](#)

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 attributes with those of the template. You can create zone templates in the local and regional cluster web UIs and in the CLI.



Caution Be careful while applying a template to an existing zone. The template overwrites all explicitly set attributes for the zone (other than its name), possibly causing severe consequences if the zone is already configured in a network. 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.

Local and Regional Web UI

Step 1 From the **Design** menu, choose **Zone Templates** under the **Auth DNS** submenu to open the List/Add Zone Templates page.

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 the **Add Zone Templates** icon in the Zone Templates pane. This opens the Add Zone Template dialog box, enter the name and click **Add Zone Template**.

To make the zone template meaningful, enter 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 [Creating Primary Zones, on page 133](#).)

After you enter these values, click **Save** at the bottom of the page.

- At the regional cluster, to pull a zone template from one or more local clusters, click the **Pull Replica** icon in the Zone Templates pane. This opens the Select Replica Zone Template Data to Pull dialog box which shows a tree view of the regional server 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**.
 - To update all the replica data for a cluster, click the **Pull Replica** icon.

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/Add Zone Templates page—Click the **Push All** icon in the Zone Templates pane.
 - To push individual zone templates on the page List/Add Zone Templates page—Click **Push**.

Both of these actions open a version of the Push Zone Template Data to Local Clusters page.

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:

- a. New zone**—Select the template from the Template drop-down list when you create the zone, as described in [Configuring Primary Forward Zones, on page 133](#).
- b. Existing zone**—After you create a zone (see [Configuring Primary Forward Zones, on page 133](#)), you can apply the template when you edit the zone on the Edit Zone page. Select the template name from the **Template** drop-down list, then click **Apply Template**.

CLI Commands

Use **zone-template name create** to create the zone template. (See [Configuring Primary Forward Zones](#), on page 133 for how to apply the template to a zone.) For example:

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

To apply a template to a zone, 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
```

When connected to a regional cluster, you can use the following pull, push, and reclaim commands. For push and reclaim, a list of clusters or "all" may be specified.

```
zone-template <name | all> pull <ensure | replace | exact> cluster-name [-report-only | -report]
```

```
zone-template <name | all> push <ensure | replace | exact> cluster-list [-report-only | -report]
```

```
zone-template name reclaim cluster-list [-report-only | -report]
```

Staged and Synchronous Modes

You can perform additions or edits to DNS zones, RRs, and hosts in one of two modes in regional cluster—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.
- **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.

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.

To choose the mode, select **Session Settings** from the **Settings** drop-down list at the top of the web UI.



Note Synchronous mode is the only DNS edit mode at the local cluster level. RR edits performed at the local cluster are immediately available via DNS.

Local and Regional Web UI

Staged or synchronous modes are preset based on the Session Edit Modes setting in **Session Settings** on the web UI main page under the **Settings** drop-down menu:

- The regional web UI is preset to **staged**.
- The local web UI is preset to **synchronous**.

CLI Commands

Set the session `dns-edit-mode` attribute to staged or synchronous. For example:

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

Configuring 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 [Configuring Primary Reverse Zones, on page 139](#) to configure a reverse zone for each network that you use.



Tip For an example of adding a forward zone, see the *"Create the Zone Infrastructure"* section in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Creating Primary Zones

Creating a primary zone requires, at a minimum, adding certain key SOA attributes and nameservers for the zone. The advantage of Basic mode in the web UI is that many of these settings are already done for you.

Local Basic Web UI

-
- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
- Step 2** Click the **Add Forward Zone** icon in the Forward Zones pane, enter the zone name (in domain name format).
- Step 3** Enter the name of the nameserver host, such as **ns1**.
- Step 4** Enter the contact e-mail name, such as **hostadmin**.
- Step 5** Click **Add DNS Zone**. Basic mode creates the zone with preset values:
- Zone default TTL—**24h**
 - Start of Authority (SOA) serial number—**1**
 - SOA secondary refresh time—**3h**
 - SOA secondary retry time—**60m**
 - SOA secondary expiration time—**1w**
 - SOA minimum TTL—**10m**
-

Local Advanced and Regional Web UI

-
- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
- Step 2** Click the Add Forward Zone icon in the Forward Zones pane, enter the zone name (in domain name format).

Step 3 Enter the name of the nameserver host, such as **ns1**.

Step 4 Enter the contact e-mail name, such as **hostadmin**.

Step 5 Enter the serial number.

Step 6 Click **Add Zone**.

Step 7 Choose an owner or region, if necessary, from the drop-down list.

Step 8 Apply an existing zone template, if necessary (see [Creating and Applying Zone Templates, on page 130](#)). Click the name of the configured template in the drop-down list.

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 9 Modify the top attributes, if necessary:

- a) Owner and region
- b) Preconfigured zone distribution (see [Managing Zone Distributions, on page 147](#))
- c) Zone default TTL

Step 10 In the SOA attributes, enter a:

- a) Serial number, such as **1**.

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. 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. You must explicitly enter this suggested serial number when creating a zone.

- b) Nameserver host, such as **ns1**.

Enter either just the hostname or its fully qualified name (such as **ns1.example.com.**, but you must end it 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 SOA record. 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) Contact e-mail name, such as **hostadmin**.

The fully qualified contact e-mail name becomes a slightly altered version of the e-mail address in that dots (.) are substituted for the at symbol (@). If using the fully qualified value, end the address with a trailing dot (for example, enter **hostadmin@example.com** as **hostadmin.example.com**).

Step 11 Enter an authoritative nameserver name under Nameservers further down the page, then click **Add Nameserver**.

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 configuration database, and from DNS updates, typically from a DHCP server. A secondary server obtains the zone data from its designated primary servers by way of a zone transfer.

You must add at least one nameserver for a zone—Cisco Prime 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. 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.

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

- a) Click **Host** to open the List Zones page.
- b) Click the zone name to open the List/Add Hosts for Zone page.
- c) Enter the hostname of the authoritative server.
- d) Enter its IP address.
- e) Click **Add Host**. The server hostname and address appear in the list.
- f) To edit the host, click its name to open the Edit Host page. Click **Modify** to implement the changes.

Step 12 Configure additional attributes as needed.

Step 13 Click **Save**.

CLI Commands

To create a primary zone, use **zone name create primary nameserver contact**. You must specify a primary DNS server; this server becomes the first authoritative DNS nameserver. For example:

```
nrcmd> zone example.com create primary ns1 hostadmin
```

The serial number defaults to 1. 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).

To add additional authoritative nameservers for the zone, 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.

Use **zone name addRR hostname A address** to add the authoritative server hostname and address. To list the host, use **zone name listHosts**. To remove the host, use **zone name removeRR hostname A**.

If you want to apply an existing template while creating a zone, use the *template* attribute. For example:

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



Note In this example, even though you need to specify the nameserver and contact as part of the syntax, the template definition (if any) overwrites them.

To apply a template after creating the zone, use **zone name applyTemplate template**. For example:

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

Editing Primary Zones

You can edit a primary zone to modify its properties, apply a template to it, or use the zone definition to create a template from it.

Local Advanced and Regional Web UI

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

- Step 2** Select the zone in the Forward Zones pane to open the Edit Zone page.
- Step 3** Make attribute changes as necessary.
- Step 4** To apply a template to the zone, choose a template name from the drop-down list at the bottom of the page, then click **Apply Template**.
- 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** To use the zone definitions to create a template from them while modifying the zone, click **Modify Zone and Save Template**. On the Save New Zone Template page, give the template a name in the Value field, then click **Save Zone Template**. You return to the List/Add Zones page.

Confirming Zone Nameserver Configuration

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

Local Advanced and Regional Web UI

Select the zone from the Forward Zones pane, and click the **Resource Records** tab. There should be an A record for each nameserver host in the zone. Edit these records or add more on this page.

See [Adding Resource Record to Zone](#), on page 164.

CLI Commands

Use `zone name listRR` to check the RRs you added.

Synchronizing Zones

Use manual zone synchronization only when there is an inconsistency between the HA main and HA backup that is not being resolved automatically by the servers. If a zone needs to be synchronized, do the following:

Regional Advanced Web UI

- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
- Step 2** Select the **Zone Sync** tab for the Primary Forward or Reverse zone.
- Step 3** Click the **Sync Zone - Report** button to open the Synchronize Zone page.
- Step 4** Expert mode includes an additional **Sync CCM Hosts from RR Data - Report** button.

CLI Commands

The `zone name sync <update | complete> [-report-only | -report]` command is available when connected to a regional cluster.

Zone Commands

The List/Add Zones (Forward or Reverse zone) page includes a **Commands** button. When clicked, this opens the Commands dialog box. These commands serve specific purposes:

- **Scavenge zone**—See the *"Scavenging Dynamic Records" section in Cisco Prime Network Registrar 11.0 DHCP User Guide.*
- **Get scavenge start time**—See the *"Scavenging Dynamic Records" section in Cisco Prime Network Registrar 11.0 DHCP User Guide.*
- **Synchronize HA Zone (Forward Zones)**—See [Synchronizing HA DNS Zones, on page 126.](#)



Note You can see the **Synchronize HA Zone** command only if the server is an HA main server. You cannot see this command if it is an HA backup 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 in the CLI. BIND 8 and BIND 9 use a configuration file, called named.conf, 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 Cisco Prime 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 primary ;
```

```
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).

The following table describes the named.boot and named.conf file directives that BIND 4 and BIND 9 support, and the corresponding Cisco Prime Network Registrar user interface location or syntax, if any.

Table 44: BIND-to-CLI Command Mappings

BIND 4 Command	BIND 9 Command	Mapping to User Interface
—	acl name { addr-match-list };	Web UI: List/Add Access Control Lists page fields (see the "Assigning ACLs on DNS Caching Servers or Zones" section in Cisco Prime Network Registrar 11.0 DHCP User Guide). CLI: acl name create value match-list=addr-match-list
—	key id { algorithm string ; secret string ; };	Web UI: List/Add Encryption Keys page fields. CLI: key name create secret algorithm=alg
limit transfers-in num	options { transfers-in num };;	Web UI: Edit DNS Server page, set <i>xfer-client-concurrent-limit</i> . CLI: session set visibility=3 dns set xfer-client-concurrent-limit=number
—	options { allow-query addr-match-list };;	Web UI: Edit DNS Server page, enable <i>restrict-query-acl</i> . CLI: dns set restrict-query-acl
options listen-on port	options { listen-on port {addr-match-list } };;	Web UI: Edit DNS Server page, set <i>Listening port</i> . CLI: dns set local-port-number=port
options max-cache-ttl num	options { max-cache-ttl num };;	Web UI: Edit DNS Server, set <i>Max. RR caching TTL</i> . CLI: dns set max-cache-ttl=num

BIND 4 Command	BIND 9 Command	Mapping to User Interface
options no-fetch-glue	options { fetch-glue no ;};	Web UI: Edit DNS Server page, enable <i>Don't fetch missing glue records</i> . CLI: dns enable no-fetch-glue
options notify yes	options { notify yes ;};	Web UI: Edit DNS Server page, enable <i>Send zone change notification (NOTIFY)</i> . CLI: dns enable notify
<i>options rrsset-order order order ...</i>	options { rrsset-order order ; order ; ... ;};	Web UI: Edit DNS Server page, enable <i>Enable round-robin</i> . CLI: dns enable round-robin
options support-ixfr yes	options { request-ixfr yes ;};	Web UI: Edit DNS Server page, enable <i>Request incremental transfers (IXFR)</i> . CLI: dns enable ixfr-enable
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: dns enable axfr-multirec-default
primary zonename file	zone "name " { type primary ;};	Web UI: Add Zone page fields. CLI: zone name create primary file=file
secondary zonename addr list [backupfile]	zone "name " { type secondary ;};	Web UI: Add Secondary Zone page fields. CLI: zone name create secondary ip-addr [,ip-addr ...]
—	zone "name " { allow-query { addr ; ... } ;};	Web UI: Edit Zone page, set <i>restrict-query-acl</i> . CLI: zone name set restrict-query-acl=addr [,addr ...]
tcplist addrlistxfernets addrlist	zone "name " { allow-transfer { addr ; ... } ;};	Web UI: Edit Zone page, enable <i>restrict-xfer</i> and set <i>restrict-xfer-acl</i> . CLI: zone name enable restrict-xfer zone name set restrict-xfer-acl=addr [,addr ...]

Configuring 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 [Adding Reverse Zones from Subnets, on page 141](#)).

Related Topics

[Adding Reverse Zones as Zones, on page 140](#)

[Adding Reverse Zones from Subnets, on page 141](#)

Adding Reverse Zones as Zones

You can manually add a reverse zone as a zone.

Local Basic or Advanced and Regional Web UI

From the **Design** menu, choose **Reverse Zones** under the **Auth DNS** submenu to open the List/Add Reverse Zones page. This page is almost identical to the List/Add Forward Zones page. Then, add a reverse zone the same way you would add a forward zone, as described in [Configuring Primary Forward Zones, on page 133](#), except use the reverse of the forward zone 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.

You can enter a DHCPv4 subnet or DHCPv6 prefix value in the Name field, which converts the subnet or prefix into an appropriate reverse zone name.

To create a reverse zone by using an IPv4 subnet or an IPv6 prefix, do the following:

-
- Step 1** From the **Design** menu, choose **Reverse Zones** under the **Auth DNS** submenu to open the List/Add Reverse Zones page.
- Step 2** Click the **Add Reverse Zone** icon in the Reverse Zones pane, enter values in the Name field. For example:
- **209.165.201.1/24**—Creates a reverse zone by using an IPv4 subnet.
 - **2001:db8:ff80:ff80::/64**—Creates a reverse zone by using an IPv6 prefix.
- Step 3** Enter the required fields to create the reverse zone:
- **Nameserver**—Enter **ns1.example.com.** (include the trailing dot).
 - **Contact E-Mail**—Enter **hostadmin.example.com.** (include the trailing dot).
 - **Serial Number**—Enter **1.**
- Step 4** Click **Add Reverse Zone**. The List/Add Reverse Zones page appears.
-

Local Basic or Advanced and Regional Web UI

To create a reverse zone by using the name of an IPv6 prefix, do the following:

-
- Step 1** From the **Design** menu, choose **Prefixes** under the **DHCPv6** submenu to open the List/Add DHCP v6 Prefixes page.
- Step 2** Click the **Add Prefixes** icon in the Prefixes pane to open the Add IPv6 Prefix dialog box.
- Step 3** Enter a prefix name (for example, **prefix-1**) and address (for example, **2001:db8:ff80:ff80::**).
- Step 4** Choose a prefix length from the drop-down list (for example, **64**).
- Step 5** Click **Add IPv6 Prefix**. The prefix is added to the List/Add DHCP v6 Prefixes page.

To create a reverse zone from the prefix,

- a) Click the **Reverse Zone** tab.

- b) Select a zone template.
- c) Click **Report**, and then click **Run**.

CLI Commands

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.

To create a reverse zone by using:

- An IPv4 subnet

For example, you can enter:

```
nrcmd> zone 209.165.201.1/24 create primary ns1.example.com. hostadmin.example.com.
```

- An IPv6 prefix

For example, you can enter:

```
nrcmd> zone 2001:db8::/64 create primary ns1.example.com. hostadmin.example.com.
```

- The name of an IPv6 prefix

For example, you can enter:

```
nrcmd> prefix prefix-1 create 2001:db8:ff80:ff80::/64
nrcmd> zone prefix-1 create primary ns1.example.com. hostadmin.example.com.
```

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.

Local Advanced and Regional Web UI

- Step 1** From the **Design** menu, choose **Subnets** under the **DHCPv4** submenu to open the List/Add Subnets page.
- Step 2** Create a subnet for the reverse zone, or use one of the existing subnets.
- Step 3** Click the **Reverse Zone** tab, and choose an existing zone template.
- Step 4** Click **Report** to show the changesets for the creation.
- Step 5** Click **Revert** to return to the List/Add Subnets page.
- Step 6** Confirm the creation by clicking **Run**, then **Reverse Zones** to see the newly created zone on the List/Add Reverse Zones page.

Getting Zone Counts on the Server

You can view the created zones associated with the DNS server, hence obtain a count, in the web UI.

Using the CLI, you can get an exact count of the total zones for the DNS server by using **dns getZoneCount** [**forward** | **reverse** | **primary** | **secondary** | **all**]. With no options specified, the command returns the total number of published zones only.

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 in detail in the *"Managing DNS Update" chapter in Cisco Prime Network Registrar 11.0 DHCP User Guide*. 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. (See the *"Configuring DNS Update Policies" section in Cisco Prime Network Registrar 11.0 DHCP User Guide*.)
- **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 *"Creating DNS Update Maps" section in Cisco Prime Network Registrar 11.0 DHCP User Guide*.)

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.



Note Zone transfer in secure mode supports both HMAC-MD5 based TSIG and GSS-TSIG.



Tip If the authoritative server for your secondary zones is also running Cisco Prime Network Registrar 6.0 or later, see [Managing Zone Distributions, on page 147](#) for how to avoid entering these zones manually. If you have only one secondary server, remove it geographically from the primary. They should not 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 primary server from which to perform zone transfers. Cisco Prime Network Registrar must know about this primary server.

Adding Secondary Forward Zones

You can add a secondary forward zone at the local cluster.

Local Basic or Advanced Web UI

-
- Step 1** From **Design** menu, choose **Secondary Zones** under the **Auth DNS** submenu to open the List/Add Secondary Zones page.
- Step 2** Click the **Add Secondary Zone** icon in the Secondary Zones pane to open the Add Secondary Zone dialog box.
- A secondary zone requires a name and a list of one or more primary servers. You can also enable restricting zone transfers to a set of hosts, then enter the access control list (ACL) of the restricted hosts in the *restrict-xfer-acl* field. Enter other attribute values as necessary.
- Step 3** Click **Add Secondary Zone**.
- Clicking the name of the secondary zone in the Secondary Zones pane opens the Edit Secondary Zone page where you can edit the secondary zone. Click **Save** on this page.
- You can add the secondary reverse zone the same way you do a secondary forward zone, except that the address must be a reverse zone address.
-

CLI Commands

To create a secondary zone, use **zone name create secondary address**. You must specify the primary DNS server IP address to perform the zone transfer.

For example:

```
nrcmd> zone shark.zone. create secondary 172.18.123.177
```

If you are using HA DNS server pair, the IP addresses must be provided by separating the addresses with comma. The HA DNS backup server will be used when the primary server is unavailable.

For example:

```
nrcmd> zone shark.zone. create secondary 172.18.123.177,172.18.123.45
```

Enabling Zone Transfers

A secondary server periodically contacts its primary server for changes, called a zone transfer. The interval is defined in the server SOA record as the secondary refresh time. You can restrict zone transfers by setting the *restrict-xfer* attribute to true (the preset value is false) on the primary server. You have to set the *restrict-xfer-acl* setting accordingly.



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 *restrict-xfer-acl* list.

Local Advanced and Regional Web UI

-
- Step 1** In the Forward Zones pane, click the name of the primary zone to open the Edit Zone page.

Step 2 In the zone attributes area, you can set the *restrict-xfer* attribute to false (the preset value). 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.

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.

Step 3 Click **Save**.

Step 4 You can force zone transfers for the DNS server in two ways:

- On the Secondary Zones pane, click the **Full Zone Transfer** button.
- To force all zone transfers from the primary server, on the Manage DNS Authoritative Server page, click the **Commands** button to Force all zone transfers.

CLI Commands

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**.

Configuring 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.

Related Topics

[Choosing Subzone Names and Servers, on page 144](#)

[Creating and Delegating Subzones, on page 145](#)

[Undelegating Subzones, on page 146](#)

[Editing Subzone Delegation, on page 146](#)

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 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.

Whenever a subzone nameserver changes its name or address, the subzone administrator must inform its parent zone so that the parent zone administrator can change the subzone nameserver and *glue records*. A glue record is an A record with the address of a subzone authoritative nameserver. If the subzone 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 Cisco Prime 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. Such a zone which creates the NS RRs and corresponding A records (glue records) for point of delegation in the parent zone is called a parented zone. A zone that does not create the NS RRs and corresponding A records (glue records) for point of delegation in the parent zone is called an unparented zone.

Consider a zone *example.com* with a parent zone *.com* and a subzone *subdomain.example.com*. If *example.com* is a parented zone, NS RRs for the *example.com* appears in two places; within the *example.com* and within its parent zone *.com*. Within *example.com* are authoritative records for the nameservers for the zone, at the point of delegation for either a subdomain of the zone or in the parent zone. The parent zone *.com* will contain non-authoritative NS RRs for *example.com* at its point of delegation and *subdomain.example.com* will have non-authoritative NS RRs in *example.com* at its point of delegation.

See [Choosing Subzone Names and Servers](#), on page 144.

Local Basic or Advanced Web UI

- Step 1** Create a zone as a subdomain of the parent domain on the List/Add Forward Zones page:
- If applying a zone template, go to **Step 2**.
 - If not applying a zone template, on the List/Add Forward Zones page, click the **Add Forward Zone** icon and add the SOA records and the nameserver with its address.
- Step 2** If Cisco Prime 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.
- Creating as subzone will create the NS RRs and corresponding A records (glue records) for point of delegation in the parent zone.

- 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.)
- Step 4** Click **Report** to show the intended changesets for the added records.
- Step 5** Click **Return** after viewing the actual changesets implemented.
- Step 6** 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**.
- Step 7** 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**.

CLI Commands

On the subzone primary nameserver machine, create the subdomain:

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

On the parent zone nameserver machine, add an NS record for the subzone nameserver, then Create a glue A record for the subzone nameserver:

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

Editing Subzone Delegation

You can edit the subzone RRs.

Local and Regional Web UI

- Step 1** On the corresponding Edit Zone page, click the **Resource Records** tab, 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.
- Step 3** Click **Modify Resource Record**.
- Step 4** Edit the glue A RR for the subzone server in the same way as in the previous steps.

CLI Commands

Use **zone name removeRR** to delete the NS and glue A records, then use **zone name addRR** to replace them.

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, Cisco Prime Network Registrar cleans up the delegation records automatically.

Local Basic or Advanced and Regional Web UI

On the corresponding Edit Zone page, click the **Resource Records** tab, delete the NS record for the subzone, then delete the glue A record for the subzone server host.

CLI Commands

Use `zone name removeRR NS` and `zone name removeRR A` to remove the subzone NS and glue A records.

Managing Zone Distributions

Creating a zone distribution simplifies creating multiple zones that share the same secondary zone attributes. It simplifies to a great extent the setup and management of multiple clusters sharing zone relationships such as primary to secondary or main to backup in the case of DNS HA.

The zone distribution requires adding one or more predefined secondary servers. Running a zone distribution synchronization adds secondary zones managed by secondary 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. Synchronizing the zone data will always sync the associated views and named ACLs for both primary and secondary zones.

The distribution must be in a star topology, that is, one primary server and multiple secondary servers. The authoritative server can only be the local primary server where the zone distribution default is defined. You can manage one zone distribution at the local cluster and multiple distributions at the regional clusters.

Related Topics

[Preparing the Zone Distribution Map, on page 147](#)

[Creating a Zone Distribution, on page 148](#)

[Pulling Zone Distributions from Replica Data, on page 150](#)

Preparing the Zone Distribution Map

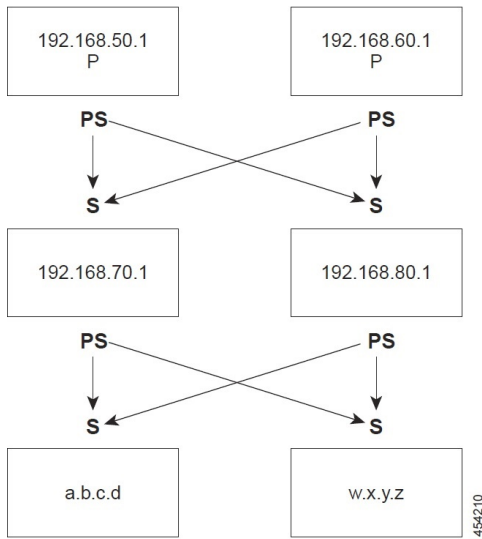
To prepare for creating a zone distribution, draw a zone distribution map diagram on paper.

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 the image below).

Figure 17: Diagramming a Zone Distribution Map



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

In the CLI, set the *primary-servers* attribute to the list of IP addresses; for example:

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

- Step 8** From the Secondary Servers drop-down list on the Add or Edit Zone Distribution Secondary Server page, choose the cluster associated with the secondary 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



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.

Local Basic or Advanced and Regional Web UI

- Step 1** From **Deploy** menu, choose **Zone Distributions** (for the regional cluster) or **Zone Distribution** (for the local cluster) under the **DNS** submenu. The option is available if the server is configured with authoritative service. This opens the regional List/Add Zone Distributions page or the local View Zone Distribution page. Note that the default zone distribution is predefined at both clusters; however, the default cluster is the only one available at the local cluster.
- Step 2** To add a new zone distribution, click the **Add Zone Distribution** icon to open the Add Zone Distribution dialog box. To edit an existing zone distribution, select its name to open the Edit Zone Distribution page.
- Step 3** In the Primary Server field, enter the cluster (or configured HA DNS pair) that has the primary server. 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 Primary Servers list, add the IP address (and optional key) for each primary server of secondary. This 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 primary servers for each of the secondary relationships. You might also want to determine the HA DNS server pairs from the Primary Servers list. You can also add an optional TSIG key or GSS-TSIG keys (see the *"Transaction Security" or "GSS-TSIG" section in the Cisco Prime Network Registrar 11.0 DHCP User Guide*) to the primary server address by hyphenating the entry in the format *address-key*. For each entry, click **Add IP Key**.
- Step 5** For a zone distribution, you need to add at least one secondary server. Click **Add Secondary Server** on the Edit Zone Distribution page. Choose the cluster of the secondary server. Optionally, if the primary servers of secondary are other than the primary servers indicated for the zone distribution, add the addresses of the secondary's primary servers, separating multiple addresses with commas. After clicking **Add Server** returns you to the Edit page, you can connect to the secondary server cluster, delete it, or edit it to change the primary servers of secondary.
- To manage the secondary servers in the zone distribution, click the **View** icon in the Manage Servers column to open the List Secondary Servers page. You can also edit the secondary server on an Edit Zone Distribution Secondary Server page.
- Step 6** 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.
- Step 7** Click **Save**.
- 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** Click the **Synchronize Zone Distribution** tab, and 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 primary servers 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 Click **Report** in the Synchronize Zone Distribution tab (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.

CLI Commands

To create the zone distribution, 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
```

To set the primary servers of secondary, use **zone-dist name set primary-servers=addresses**, separating the addresses with commas. For example:

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

To add the secondary server, use **zone-dist name addSecondary secondary-cluster**. For example:

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

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
```

To synchronize the zone distributions, 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 Primary Servers list (primary servers of secondary) 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
```

Pulling Zone Distributions from Replica Data

You can pull zone distributions from the local replica data instead of explicitly creating them.



Tip

For an example of pulling local zone data to create a zone distribution, see the *"Pull Zone Data and Create a Zone Distribution"* section in *Cisco Prime Network Registrar 11.0 Administration Guide*.

Regional Web UI

-
- Step 1** From **Deploy** menu, choose **Zone Distribution** under the **DNS** submenu to open the List/Add Zone Distribution page.
 - Step 2** On the List/Add Zone Distribution page, click the **Synchronize Zone Distribution** tab.
 - Step 3** Choose the zone synchronization mode (**Update**, **Complete**, or **Exact**). These modes are described in the table on that page.
 - Step 4** Click **Report** at the top of the dialog box.
 - Step 5** Click **Run**.
-

Managing DNS ENUM Domain

Creating separate ENUM domains simplifies the management of Naming Authority Pointer (NAPTR) Electronic Numbering (ENUM). It simplifies to a great extent the setup and management of E.164 numbers and how available services are connected to the E.164 numbers. When you create an ENUM zone and add the corresponding E.164 numbers, Cisco Prime Network Registrar automatically creates a forward zone and the respective NAPTR resource records.

Managing DNS ENUM Defaults

To configure the default ENUM settings, do the following:

Local Basic or Advanced Web UI

-
- Step 1** From the **Design** menu, choose **Defaults** under the **DNS ENUM** submenu to open the Manage DNS ENUM Defaults page.
 - Step 2** Enter the Top-level Domain.
 - Step 3** Enter the Local Prefix, such as +46.
 - Step 4** Enter the Default Services values: Click the **Add** button in the **Services** section, select a service type, enter a URI, and click **Add**.
 - Step 5** Select a Zone Template from the drop-down list.
 - Step 6** Click **Save**.
-

CLI Commands

Use `dns-enum-config set [number-prefix prefix | zone-template name]` to set the default ENUM domain, default top-level domain and local prefix, service, and zone template.

Use `dns-enum-config addService type subtype URI [order [preference]]` to add the default service.

Use `dns-enum-config removeService type subtype URI` to remove the default service user.

Adding DNS ENUM Domains

Adding an ENUM domain involves creating a domain name. You can also define an owner and use a zone template.

When you create an ENUM zone, Cisco Prime Network Registrar automatically creates a forward zone. For example, if you create an ENUM domain for E.164 number prefix 100 and the default top-level domain is set to e164enum.net., a forward zone 0.0.1.e164enum.net. is automatically created and appears in the list of forward zones.

To configure an ENUM domain, do the following:

Local and Regional Web UI

- Step 1** From the **Design** menu, choose **Domains** under the **DNS ENUM** submenu to open the List/Add DNS ENUM Domains page.
- Step 2** Click the **Add Domains** icon in the Domains pane to open the Add ENUM Domain dialog box.
- Step 3** Enter the E.164 number prefix for the domain, such as 897.
- Step 4** Enter the name of the nameserver host, such as ns1.
- Step 5** Enter the contact e-mail name, such as hostadmin.
- Step 6** Click **Add ENUM Domain**. The domain will be created with the default local prefix such as +4689. The Basic mode creates the zone with the following preset values:
- Zone default TTL-24h
 - Start of Authority (SOA) serial number-1
 - SOA secondary refresh time-3h
 - SOA secondary retry time-60m
 - SOA secondary expiration time-1w
 - SOA minimum TTL-10m
-

CLI Commands

Use **dns-enum-domain name create [zone-template=name] [nameservers [person]]** to create ENUM domain.

Use **dns-enum-domain name delete** to delete ENUM domain.

When connected to a regional cluster, you can use the following pull, push, and reclaim commands.

dns-enum-domain <name | all > pull < ensure | replace | exact > cluster-name [-report-only | -report]

dns-enum-domain <name | all > push < ensure | replace | exact > cluster-list [-report-only | -report]

dns-enum-domain name reclaim cluster-list [-report-only | -report]

Adding DNS ENUM Numbers

Cisco Prime Network Registrar supports NAPTR RRs. These records help with name resolution in a particular namespace and are processed to get to a resolution service.

In addition to the option of adding NAPTR resource records, you can now directly add the E.164 numbers and associate the corresponding services with the numbers. When you add a DNS ENUM number, you need to specify either the E.164 number prefix of the parent domain or the Zone templates, and a NAPTR resource record is created for the E.164 number. This approach uses a reversed E.164 number and treats every digit as a node on the DNS name hierarchy. For example, the E.164 address +4689761234 creates a NAPTR RR 4.3.2.1.6.7.9.8 for the +46 E.164 prefix domain.

For more information on NAPTR resource records, see the [Name Resolution in a Namespace Using NAPTR Resource Records, on page 170](#).

Local and Regional Web UI

-
- Step 1** From the **Design** menu, choose **Numbers** under the **DNS ENUM** submenu to open the List/Add DNS ENUM Numbers page.
 - Step 2** Click the **Add Numbers** icon in the Numbers pane to open the Add ENUM Number dialog box.
 - Step 3** Enter the E.164 number along with the E.164 number prefix, such as 1234.
 - Step 4** Click the **Add** button in the **Services** section, select a service type, enter URI, and click **Add**.
 - Step 5** Enter the E.164 number prefix for the parent domain.
 - Step 6** Select the Zone Template if you have not specified the E.164 prefix.
 - Step 7** Select a Ported option and enter the Ported Nameserver FQDN.
 - Step 8** Click **Add ENUM Number**. The number will be created and added under the domain +4689.
-

CLI Commands

Use `dns-enum-number number create type subtype URI [zone-template=name] [domain-prefix]` to create ENUM number.

When connected to a regional cluster, you can use the following pull, push, and reclaim commands. For push and reclaim, a list of clusters or "all" may be specified.

`dns-enum-number <name | all > pull < ensure | replace | exact > cluster-name [-report-only | -report]`

`dns-enum-number <name | all > push < ensure | replace | exact > cluster-list [-report-only | -report]`

`dns-enum-number name reclaim cluster-list [-report-only | -report]`

Pulling and Pushing ENUM Domains

You can push ENUM Domains to and pull ENUM Domains from local clusters on the List/Add DNS ENUM Domains page in the regional cluster web UI.

Pushing ENUM Domains to Local Clusters

To push ENUM domains to the local cluster, do the following:

Regional Web UI

-
- Step 1** From the **Design** menu, choose **Domains** under the **DNS ENUM** submenu to view the List/Add DNS ENUM Domains page in the regional web UI.
 - Step 2** Click the **Push All** icon in the Domains pane to push all the ENUM domains listed on the page, or select the ENUM domain on the Domains pane and click the **Push** icon to open the Push ENUM Domain page.
 - Step 3** Choose a push mode using one of the Data Synchronization Mode radio buttons.
 - If you are pushing all the ENUM Domains, you can choose Ensure, Replace, or Exact.
 - If you are pushing an ENUM Domain, you can choose Ensure or Replace.

In both cases, Ensure is the default mode.

Choose **Replace** only if you want to replace the ENUM domain data at the local cluster. Choose **Exact** only if you want to create an exact copy of the ENUM domain data at the local cluster, thereby deleting all ENUM domain data that is not defined at the regional cluster.

Step 4 Click **Push Data to Clusters**.

CLI Command

When connected to a regional cluster, you can use **dns-enum-domain** *<name | all >* **push** *< ensure | replace | exact >* *cluster-list* [**-report-only**| **-report**].

Pulling ENUM Domains from the Replica Database

To pull ENUM domains from the replica database, do the following:

Regional Web UI

- Step 1** From the **Design** menu, choose **Domains** under the **DNS ENUM** submenu to view the List/Add DNS ENUM Domains page in the regional web UI.
- Step 2** Click the **Pull Replica** icon in the Domains pane.
- Step 3** Click the **Replica** icon in the Update Replica Data column for the cluster. (For the automatic replication interval, see the "Replicating Local Cluster Data" section in *Cisco Prime Network Registrar 11.0 Administration Guide*.)
- Step 4** Choose a replication mode using one of the Mode radio buttons.
- Step 5** Leave the default **Replace** mode enabled, unless you want to preserve any existing ENUM domains data at the local cluster by choosing **Ensure**.
- Step 6** Click the **Pull all ENUM Domains** button to view the pull details, and then click **Run**.

CLI Command

When connected to a regional cluster, you can use **dns-enum-domain** *<name | all >* **pull** *< ensure | replace | exact >* *cluster-name* [**-report-only**| **-report**].

Pulling and Pushing ENUM Numbers

You can push ENUM numbers to and pull ENUM numbers from local clusters on the List/Add DNS ENUM Numbers page in the regional cluster web UI.

Pushing ENUM Numbers to Local Clusters

To push ENUM numbers to the local cluster, do the following:

Regional Basic and Advanced Web UI

- Step 1** From the **Design** menu, choose **Numbers** under the **DNS ENUM** submenu to view the List/Add DNS ENUM Numbers page in the regional web UI.
- Step 2** Click the **Push All** icon in the Numbers pane to push all the ENUM numbers listed on the page, or select the ENUM number on the Numbers pane and click the **Push** icon to open the Push ENUM Number page.

- Step 3** Choose a push mode using one of the Data Synchronization Mode radio buttons.
- If you are pushing all the ENUM numbers, you can choose Ensure, Replace, or Exact.
 - If you are pushing an ENUM number, you can choose Ensure or Replace.

In both cases, Ensure is the default mode.

Choose Replace only if you want to replace the ENUM number data at the local cluster. Choose Exact only if you want to create an exact copy of the ENUM number data at the local cluster, thereby deleting all ENUM number data that is not defined at the regional cluster.

- Step 4** Click **Push Data to Clusters**.

CLI Commands

When connected to a regional cluster, you can use **dns-enum-number** <name | all > **push** < ensure | replace | exact > cluster-list [-report-only] -report].

Pulling ENUM Numbers from the Replica Database

To pull ENUM numbers from the replica database, do the following:

Regional Basic and Advanced Web UI

- Step 1** From the **Design** menu, choose **Numbers** under the **DNS ENUM** submenu to view the List/Add DNS ENUM Number page in the regional web UI.
- Step 2** Click the **Pull Replica** icon in the Numbers pane.
- Step 3** Click the **Replica** icon in the Update Replica Data column for the cluster. (For the automatic replication interval, see the "Replicating Local Cluster Data " section in *Cisco Prime Network Registrar 11.0 Administration Guide*.)
- Step 4** Choose a replication mode using one of the Mode radio buttons.
- Step 5** Leave the default Replace mode enabled, unless you want to preserve any existing ENUM number data at the local cluster by choosing Ensure.
- Step 6** Click the **Pull all ENUM Numbers** button to view the pull details, and then click **Run**.

CLI Commands

When connected to a regional cluster, you can use **dns-enum-number** <name | all > **pull** < ensure | replace | exact > cluster-name [-report-only] -report].



CHAPTER 11

Managing DNS Views

DNS Views let you present alternate versions of zone data to different communities of clients using a single name server. For example, a DNS server for `example.com` can maintain two views of the zone, where the view of `example.com` that can be queried internally includes many hosts that do not exist in the external view. Each zone view is treated as an independent copy of the zone. The DNS server, when answering queries on the zone, uses the match criteria defined in each view to determine the matching zone for the client. The query is answered based on that zone contents. In some cases, the zone contents may only vary slightly between views.

- [DNS Views Processing, on page 157](#)
- [Key Points to Remember While Working on DNS Views, on page 158](#)
- [Managing DNS Views, on page 159](#)
- [Reorder DNS Views, on page 160](#)
- [Synchronizing DNS Views, on page 160](#)
- [Pushing and Pulling DNS Views, on page 160](#)

DNS Views Processing

DNS Views allow a name server to segregate the data and provide a different view of the data based on the clients accessing it. When DNS receives a DNS request, the request is processed to associate it with a DNS View. The association is performed by matching the client source and/or destination address to the source and destination ACLs configured on the view. Views are matched in priority order with the lowest non-zero priority being matched first. Once a request is matched to a DNS View, only the data in that view is available to the request. There is a one-to-one mapping between zones and views—a zone can only exist in one view. If the zone must exist in more than one view, make copies of the zone and associate with different views.

If you have an internal view and an external view, a typical setup is to set the priority of the internal view to one and set the ACLs (typically `acl-match-clients`) to match the criteria for internal clients. For the external view, leaving the default priority and ACLs will allow all requests not matching the internal view to match the external view.



Note Getting a NOTAUTH rcode response when DNS Views are configured, typically indicates that the request matched a view where the zone does not exist.



Note The auto-view detection is only applicable for Cisco Prime Network Registrar servers.

Views for the DNS client servers such as Caching DNS, Secondary DNS, Primary for Notices, DHCP, and so on, are easily defined with minimal configuration.

From Cisco Prime Network Registrar 10.1, DNS Views that are not associated with any zones are automatically ignored by default. However, in earlier versions, they were still processed and could possibly associate clients with empty views.

Key Points to Remember While Working on DNS Views

Following are some of the key points or attributes to know while working on DNS Views:

- **View ID**—Defines a unique integer identifier for the view that is assigned by the CCM server or the user while creating DNS Views.
- **View Priority (*priority attribute*)**—Each DNS View is assigned a unique priority to determine the view processing order. The lowest non-zero priority is processed first, followed by the second lowest, and so on. A zero priority is reserved for the Default view, which is always processed last. The web UI provides a mechanism to reorder views without explicitly setting the priority.
- **Default View**—The default view is created with *view-id=0*, *priority=0*, and client and destination ACLs set to any. Requests that do not match a named view always falls into the default view. By default, zones are created with a *view-id=0*, which automatically places them in the default view. The default view cannot be modified or deleted.
- ***acl-match-clients attribute***—Specifies the ACLs that map clients to a view based on the client source address. The default is any and must be modified in order to have the clients associated with the appropriate views.
- ***acl-match-destinations (Expert mode attribute)***—Specifies the ACLs that map clients to a view based on the client destination address. The default is any and should only be changed if the DNS server is using different network interfaces for different views.
- ***ignore-unused-views attribute***—Controls whether or not the DNS server uses configured DNS Views that are not associated with any of its configured zones.
- **Alternate Views**—In Cisco Prime Network Registrar 11.0, zones can be referenced by multiple views without the need to make copies of the zone.

This can be useful in a viewed configuration where a subset of zones are common across multiple views. To make the zones visible to other views, set the *alternate-view-ids* attribute for the zone and reload the DNS server. It is recommended that common zones have their *view-id* set to the Default view. Note that when changing the UI session's DNS View, only the zones with a matching *view-id* are displayed.

- The Cisco Prime Network Registrar Caching DNS server can associate the client requests to the appropriate views on behalf of the Authoritative DNS server. This is done by configuring the DNS Views on the Caching DNS server and setting the *uses-views* attribute on the List/Add Exceptions page to **true**. The Caching DNS server maps the client to the appropriate view and tags the queries forwarded to the Authoritative DNS server with the appropriate view. Therefore, in these cases, the view mapping is done by the Caching DNS server.



Note The Caching DNS server only maps clients to *acl-match-clients*. The *acl-match-destinations* attribute is ignored.

DNS Views and Exception settings are automatically synced/set by zone distribution.

Managing DNS Views

You can create, edit, and delete DNS Views from local or regional cluster. You can also push or pull views and ACLs in Ensure, Replace, and Exact modes from or to the regional CCM server.



Note You can create a maximum of 100 views.

Local and Regional Web UI

To create DNS Views:

- Step 1** From the **Design** menu, choose **Views** under the **Auth DNS** (or **Cache DNS** (Local web UI)) submenu.
- Step 2** On the Views pane, click the **Add View** icon.
- Step 3** Specify the name for the DNS View.
- Step 4** Specify the view ID (in Advanced mode). If you do not specify, the application automatically assigns a view id to the view.
- Step 5** You can specify the ACL that maps the client to this view in the *acl-match-clients* field.
- Step 6** Click the **Add DnsView** button.
- Step 7** To edit a DNS View, click its name in the Views pane on the left and then edit the attributes as required..

CLI Commands

The **view** command is used to control and manage DNS Views for the DNS servers. For example:

```
nrcmd> view MyView create
```

When connected to a regional cluster, you can use the following pull, push, and reclaim commands. For push and reclaim, a list of clusters or "all" may be specified.

```
view <name | all> pull <ensure | replace | exact> cluster-name [-report-only | -report]
```

```
view <name | all> push <ensure | replace | exact> cluster-list [-report-only | -report]
```

```
view name reclaim cluster-list [-report-only | -report]
```

Reorder DNS Views

When you create a set of DNS Views, you can specify the priority order. To specify the priority order:

-
- Step 1** From the **Design** menu, choose **Views** under the **Auth DNS** submenu to open the List/Add Zone Views page.
- Step 2** Click the **Reorder Views** icon in the Views pane to open the Reorder dialog box.
- Step 3** Set the priority for the DNS Views rules by either of the following methods:
- Select the view and click the **Move up** or **Move down** icon to reorder the rules.
 - Select the view and click the **Move to** button, and enter the row number to move the view.
- Step 4** Click **Save** to save the reordered list.
- If you delete a view, you get a choice to delete all zones.
-

CLI Commands

Use `dns-view name create` to add DNS Views (see the `dns-view` command in the CLIGuide.html file in the install-path/docs directory for syntax and attribute descriptions).

Synchronizing DNS Views

Zone distribution sync, single zone sync, and HA DNS zone sync will always sync associated views and named ACLs for both primary and secondary zones. The synchronization modes applied while running zone distribution or HA DNS sync vary. When you run:

- **Zone Distribution Sync**—views will be synchronized in Replace mode for all zone distribution sync types (Update, Complete, and Exact), while ACLs will use Ensure mode. If caching DNS servers are included in the zone distribution, the associated views and named ACLs will be synchronized to these servers and the primary servers list will be configured as exceptions for the unique set of domain names in the distribution. The user must exclude secondaries and/or caching servers.
- **HA DNS Sync**—views will be updated in Replace mode for both Update and Complete sync, while Exact sync will sync views in Exact mode.

Pushing and Pulling DNS Views

You can also push views and ACLs to and pull views and ACLs from the regional cluster in Ensure, Replace, and Exact modes.

Pushing DNS Views to Local Clusters

You can push the views you create from the regional cluster to any of the local clusters.

Regional Web UI

- Step 1** From the **Design** menu, choose **Views** under the **Auth DNS** submenu to open the List/Add Zone Views page.
- Step 2** On the Views pane, click the **Push All** icon in the left pane, or select a DNS View and click **Push** at the top of the Edit Zone View page. This opens the Push Data to Local Clusters or Push Zone View page.
- Step 3** Choose a push mode using one of the Data Synchronization Mode radio buttons.
 - If you are pushing all the DNS Views, you can choose Ensure, Replace, or Exact.
 - If you are pushing a DNS View, you can choose Ensure or Replace.

In both the above cases, Ensure is the default mode.

Choose Replace only if you want to replace the existing DNS View data at the local cluster. Choose Exact only if you want to create an exact copy of the DNS View at the local cluster, thereby deleting all DNS Views that are not defined at the regional cluster.

- Step 4** Choose one or more local clusters in the Available field of the Destination Clusters and move it or them to the Selected field.
 - Step 5** Click **Push Data to Clusters**.
-

CLI Commands

When connected to a regional cluster, you can use `view <name | all> push <ensure | replace | exact> cluster-list [-report-only | -report]`.

Pulling DNS Views from Local Clusters

Instead of explicitly creating views, you can pull them from the local clusters. In the regional web UI, you may first want to update the view replica data by clicking the Replica icon next to the cluster name.

Regional Web UI

- Step 1** From the **Design** menu, choose **Views** under the **Auth DNS** submenu to open the List/Add Zone Views page.
 - Step 2** Click the **Pull Replica** icon in the Views pane.
 - Step 3** Choose the data synchronization mode (**Update**, **Complete**, or **Exact**). These modes are described in the table on that page.
 - Step 4** Click **Report** at the bottom of the dialog box.
 - Step 5** Click **Run**.
-

CLI Commands

When connected to a regional cluster, you can use `view <name | all> pull <ensure | replace | exact> cluster-name [-report-only | -report]`.



CHAPTER 12

Managing Resource Records

This chapter explains how to configure some of the more advanced DNS zone and server parameters by using the Cisco Prime Network Registrar web UI and CLI. Before you proceed with the concepts in this chapter, read [Managing Zones, on page 129](#) which explains how to set up the basic properties of a primary and secondary DNS server and its zones.

- [Managing Resource Records for Zone, on page 163](#)
- [Adding Resource Record to Zone, on page 164](#)
- [Editing Resource Records, on page 165](#)
- [Removing Resource Records from Zone, on page 165](#)
- [Managing Resource Records for Host, on page 166](#)
- [Protecting Resource Record Sets, on page 166](#)
- [Searching Server-Wide for Records and Addresses, on page 167](#)
- [Filtering Resource Records, on page 169](#)
- [Advertising Services to Network Using Service Location \(SRV\) Records, on page 169](#)
- [Name Resolution in a Namespace Using NAPTR Resource Records, on page 170](#)
- [DNS Certification Authority Authorization \(CAA\) Resource Record, on page 171](#)
- [Uniform Resource Identifier \(URI\) Resource Records, on page 172](#)

Managing Resource Records for Zone

Resource records (RRs) comprise the data within a DNS zone. Although there is no fixed limit to the number of RRs a zone may own, in general, a zone may own one or more RRs of a given type (the zone always has a Start of Authority, or SOA, record). There are some exceptions depending on the types involved. All RRs have the entries described in the following table.

Table 45: Resource Record Common Entries

RR Entry	Description
Name	Owner of the record, such as a zone or hostname.
Class (not required for all formats)	Cisco Prime Network Registrar supports only the IN (Internet) class.
TTL (time to live)	Amount of time to store the record in a cache, in seconds. If you do not include a TTL, Cisco Prime Network Registrar uses the zone default TTL, defined as a zone attribute.

RR Entry	Description
Type	Type of the record, such as A (AAAA for IPv6), NS, SOA, MX, and so on. There are many types that various RFCs define, although fewer than ten are in common use.
Record data	Data types whose format and meaning varies with record type.

Related Topics

[Adding Resource Record to Zone, on page 164](#)

[Protecting Resource Record Sets, on page 166](#)

[Editing Resource Records, on page 165](#)

[Removing Resource Records from Zone, on page 165](#)

[Searching Server-Wide for Records and Addresses, on page 167](#)

[Filtering Resource Records, on page 169](#)

[Advertising Services to Network Using Service Location \(SRV\) Records, on page 169](#)

[Name Resolution in a Namespace Using NAPTR Resource Records, on page 170](#)

Adding Resource Record to Zone

Before adding or modifying RRs, keep in mind the two distinct dns edit modes that you can set and work in: staged and synchronous (see the *"Staged and Synchronous Modes"* section in *Cisco Prime Network Registrar 11.0 DHCP User Guide*).

Administrator roles required for RR management are the dns-admin role at the local cluster and the central-dns-admin role at the regional cluster. The host-admin role at the local cluster and the central-host-admin role at the regional cluster can view host records only.

Local and Regional Web UI

-
- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
- Step 2** In the Forward Zones pane, click the zone name to open the Edit Zone page. Note that resource record edits is managed jointly by CCM and DNS, and a system lock is used to prevent DNS and CCM from accessing the resource record database at the same time.
- Tip** Records are listed in the formats that their respective RFCs specify, with only the first record in a set labeled with its name, and in DNSSEC order. To reduce or increase the items in the table, change the page size value at the bottom of the page, then click **Change Page Size**.
- Step 3** Click the **Resource Records** tab.
- Step 4** Add the RR name, TTL (if not using the default TTL), type, and data as appropriate.

Step 5 By default, RRs are protected, which means that DNS Updates cannot overwrite them (see [Protecting Resource Record Sets, on page 166](#)). To unprotect the RRs, click the **Locked** icon to the left of the record name to change it to the Unlocked icon. Likewise, to protect the record, click the **Unlocked** icon to change it to the **Locked** icon.

Step 6 Click **Add Resource Record**.

CLI Commands

Use **zone name addRR** to add a protected RR of a certain type. You can specify the name as a relative name, if the owner is in the same domain, an absolute name (by supplying the FQDN), or the same name as the zone name (by using the at [`@`] symbol).

For example:

```
nrcmd> zone example.com addRR -sync host101 A 192.168.50.101
```

Use **zone name addDNSRR type data** to add an unprotected RR.

Editing Resource Records

You can edit RRs as an individual record or as an RR set:

- **Individual RRs**—Click the Edit icon next to the record name to open the Edit RR in Zone page.
- **RR sets**—Click the name of the record to open the Edit RR Set in Zone page.

For a description of the fields to enter data, see [Adding Resource Record to Zone, on page 164](#).

Removing Resource Records from Zone

You can remove RRs from a zone.

Local and Regional Web UI

On the Resource Records tab for the Zone page:

- To remove an entire record name set, click the **Delete** icon next to the record set name in the list, then confirm the deletion.
- To remove individual records from the set, click the name of the record set to open the edit page, click the **Delete** icon next to the individual record in the list, then confirm the deletion.

CLI Commands

The CLI includes two removal commands, depending on the type of RR to remove:

- Use **zone name removeRR** to remove any RR. You must specify the owner. If you omit the data, Cisco Prime Network Registrar removes all records of the specified type for the specified owner. Similarly, if you omit the type, Cisco Prime Network Registrar removes all records for the specified owner.
- Use **zone name removeDNSRR** to remove unprotected RRs only.

Managing Resource Records for Host

You can manage the RRs for a host by configuring the host record rather than the individual RRs. When you define a host, the DNS server automatically creates an Address (A) RR for IPv4, or an AAAA RR for IPv6, for it. If the reverse zone for the host exists, the server can also create the associated Pointer (PTR) RR for it.

See [Managing Hosts, on page 175](#) for details.

Protecting Resource Record Sets

When an RR is protected, DNS Updates cannot modify the record. Most administratively created RRs are protected. However, RRs created by DNS Updates must be unprotected to allow the server to modify them. You can set this protection status for each RR set on the List/Add DNS Server RRs for Zone page.

Note that only the primary DNS server can recognize this protection status; secondary servers do not recognize the protection status of their RRs.



Caution Zone scavenging can remove RRs that are unprotected. See the *"Scavenging Dynamic Records"* section in *Cisco Prime Network Registrar 11.0 DHCP User Guide* for details.

Local and Regional Web UI

To protect an existing RR, do the following:

-
- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
 - Step 2** In the Forward Zones pane, click the zone name to open the Edit Zone page.
 - Step 3** Click the **Resource Records** tab.
 - Step 4** On the Resource Records tab, click the Resource Record name in the list of Resource Records to edit the resource record.
 - Step 5** Click **Protect Set** button to unprotect the selected RR set.
 - Step 6** Click **Save** to save the resource record attribute modification.
-

Unprotecting Resource Record Sets

You can also unprotect an RR. To unprotect an RR while adding, click the **Locked** icon next to the Resource Record name field. The icon changes to the **Unlocked** icon.

Local and Regional Web UI

To unprotect an existing RR, do the following:

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

- Step 2** In the Forward Zones pane, click the zone name to open the Edit Zone page.
- Step 3** Click the **Resource Records** tab.
- Step 4** On the Resource Records tab, click the Resource Record name in the list of Resource Records to edit the resource record.
- Step 5** Click **Unprotect Set** button to unprotect the selected RR set.
- Step 6** Click **Save** to save the resource record attribute modification.

Note The icon to the left of the RR set name indicates the status of the Resource Record, whether it is protected or unprotected.

CLI Commands

To protect the RR sets, use `zone name protect-name rrset-name`. To unprotect the zone, use `unprotect-name rrset-name`. For example:

```
nrcmd> zone example.com protect-name boston
100 Ok
protected boston
```

```
nrcmd> zone example.com unprotect-name boston
100 Ok
unprotected boston
```

Searching Server-Wide for Records and Addresses

With Cisco Prime Network Registrar, you can search for RRs and IP addresses server-wide. The search is a filter mechanism whereby you can specify a combination of RR and address attributes to target one or more RRs or addresses configured for the network. The search function is available at the local cluster only.

You can search RRs by:

- IP address
- Protection state
- Name prefix
- Type
- Zone

Local Advanced Web UI

To search resource records by IP address, do the following:

- Step 1** From the **Operate** menu, choose **DNS RRs By IP Address** under the **Reports** submenu to open the IP Address Search page.
- Step 2** To search by IP address, enter an IP address, then click **Search**.

Note In an IP address search, the DNS server does not search all forward zones for RRs that have the specified address in the data field. Instead, the server looks up the matching PTR record in the reverse zone and returns all the respective RRs in the forward zone.

Local Advanced Web UI

To search resource records, do the following:

-
- Step 1** From the **Operate** menu, choose **DNS Resource Records** under the **Reports** submenu to open the DNS Resource Record Search page.
- Step 2** Choose a filter attribute from the drop-down list.
- Step 3** Choose a filter type from the drop-down list depending on the filter attribute you chose:
- **RR Protection State**—RR Protection Status, either locked or unlocked.
 - **RR Name Prefix**—RR Name Prefix.
 - **RR Type**—RR Type.
 - **Zone**—Zone List, Regular expression, or Zone Flags.
- Step 4** Enter or select a Value, based on the Type selected. To clear the filter, click **Clear Filter**.
- Step 5** Click **Add Element** to add the search element to the filter elements list. The Filter Elements heading changes to identify the filter attribute and value used for the filter. If you add more than one element, the heading identifies the ANDed values of the elements. For example, if you add an element for a name prefix search for user, then add another element for an RR type search for A records, the filter element heading will identify the search as ****RR Name Prefix = user AND RR Type = A**.
- Step 6** You can add as many elements as you like (remembering that the search results are an intersection of the filter elements). View the filter elements list by clicking the plus sign (+).
- Step 7** Click **Search**.
- Step 8** Check the table of resulting RRs from the search, which shows for each RR its zone, hostname, TTL, type, and associated data. If necessary, change the page size to see more entries at one time (you might still need to page forward and back). The RRs are sorted in DNSSEC order.
- Tip** If the search results are less than expected due to the ANDing of the filter elements, look at the filter list for any element that might be compromising the search, delete it by clicking the Delete icon next to it, then redo the search.
-

CLI Commands

Use **dns findRR** to find RRs across the zones. The command syntax is of two kinds:

```
nrcmd> dns findRR -name fqdn | domainaddr
```

```
nrcmd> dns findRR [-namePrefix nameprefix] [-rrTypes RRtypelist] [-protected| -unprotected]
[-zoneType
forward| reverse| primary|secondary| ALL]
```

You can search by domain or its address, or enter the beginning characters of the RR name (the name prefix). If you search by RR name prefix, you can narrow the search by a list of RR types, protection status, or zone type. The output clearly indicates the zone for each found entry. For example:

```
nrcmd> dns findRR -namePrefix user -rrTypes A

userhost101.example.com IN A 192.168.50.101
```



```
userhost102.example.com IN A 192.169.50.102
userhost103.boston.example.com IN A 192.168.50.103
```

Filtering Resource Records

You might want to filter records to display only one type of record, such as an A (or IPv6 AAAA) or PTR record. (See also [Searching Server-Wide for Records and Addresses, on page 167.](#))

Local Basic or Advanced and Regional Web UI

You can filter RRs right from the Edit Zone page. Look for the Name and Type fields just below the **Add Resource Record** button.

By default, RRs are sorted alphabetically by name, starting with the top-of-zone records (marked with the @ symbol), and secondarily sorted by type, then data. You can also sort them by:

- **Protected state**—You can click All, Unprotected, or Protected.
- **Name prefix**—Starting characters in the name. Note that the * character is not a wildcard. For example, entering **al** returns alberta, allen.wrench, and allie, whereas entering **al*** returns al* and al*ert.
- **RR type**—Click one of the RR types in the drop-down list, such as A (or IPv6 AAAA) or TXT.

When the selection is complete, click **Filter List**. This returns just the filtered entries in the table below the fields. To return to the full, unfiltered list, click **Clear Filter**.

CLI Commands

Use **zone zonename findRR** to search on RR name prefixes, RR types, or protection status:

```
nrcmd> zone zonename findRR [-namePrefix nameprefix] [-rrTypes RRtypelist] [-protected|
-unprotected]
```

Advertising Services to Network Using Service Location (SRV) Records

The service location (SRV) RR is used to advertise services to the network. This RR is defined in the RFC 2782, “A DNS RR for specifying the location of services (DNS SRV).” The SRV can have an associated A or AAAA record. Windows domain controller is one service that uses the SRV records.

The RFC defines the format of the SRV record (DNS type code 33) as:

```
_service._protocol.name ttl class SRV priority weight port target
```

There should always be an A record associated with the SRV record target so that the client can resolve the service back to a host. In the Microsoft Windows implementation of SRV records, the records might look like this:

```
myserver.example.com A 201.165.201.1
_ldap._tcp.example.com SRV 0 0 389 myserver.example.com
_kdc._tcp.example.com SRV 0 0 88 myserver.example.com
_ldap._tcp.dc._msdcs.example.com SRV 0 0 88 myserver.example.com
```

An underscore (_) always precedes the service and protocol names. In the example, **_kdc** is the Key Distribution Center. The priority and weight help a client choose between target servers providing the same service (the

weight differentiating those with equal priorities). If the priority and weight are all set to zero, the client orders the servers randomly.



Note For a description of how Windows clients interoperate with DNS and DHCP servers, including scavenging dynamic RRs, see the *"Configuring DNS Update for Windows Clients"* section in *Cisco Prime Network Registrar 11.0 DHCP User Guide*.

Name Resolution in a Namespace Using NAPTR Resource Records

Cisco Prime Network Registrar supports Naming Authority Pointer (NAPTR) RRs. These records help with name resolution in a particular namespace and are processed to get to a resolution service. Because NAPTR records are a proposed standard, RFC 3403, Cisco Prime Network Registrar only validates their numeric record fields. However, the proposed standard requires a value for each field, even if it is null (""), and there are no preset values.

When using a NAPTR record to locate a Session Initiation Protocol (SIP) proxy, see the proposed standard, RFC 2916 or RFC 3263. In RFC 2916, the ENUM working group of the Internet Engineering Task Force specifies NAPTR records to map E.164 addresses to Universal Resource Identifiers (URIs). Using the NAPTR record resolves a name in the E.164 international public telecommunication namespace to a URI, instead of providing the name of a service to use as a resolver. The U flag was added to the NAPTR record for this purpose.

For example, to specify a SIP proxy for the phone number +4689761234, add a NAPTR record at the name 4.3.2.1.6.7.9.8.6.4.e164.arpa. with this content:

```
100 10 "u" "sip+E2U" "/^.*$/sip:info@example.com/" .
```

This sets these fields of the NAPTR record:

```
order = 100
preference = 10
flags = "u"
service = "sip+E2U"
regexp = "/^.*$/sip:info@example.com/"
replacement = .
```

After you configure these fields, the DNS client dealing with phone number +4689761234 can now find an SIP service URI by replacing the number with sip:info@tele2.se. The E.164 zone mostly uses the NAPTR record for wholesale replacement of the input telephone number. Section 3.2.3 of RFC 2916 includes an example of one transformation to a Lightweight Directory Access Protocol (LDAP) query that preserves some of the digits. The E.164 zone does not map to service location (SRV) records because it wants to obtain a SIP URL that is more humanly readable to the left of the at (@) symbol.

Local Basic or Advanced and Regional Web UI

- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
- Step 2** Click the **Resource Records** tab.

- Step 3** Enter the owner of the record in the **Name** field.
- Step 4** Enter the **TTL** (if necessary).
- Step 5** Select **NAPTR** from the **Type** drop-down list.
- Step 6** Enter the data as a string embedded in quotes and separated by spaces:
- Order
 - Preference
 - Flags
 - Service
 - Regular expression
 - Replacement string

Example:

```
"100 10 u sip+E2U /^.*$/sip:info@tele2.se/ ."
```

- Step 7** Click **Add Resource Record**.

CLI Commands

Use `zone name addRR` to add a protected resource record to a zone.

DNS Certification Authority Authorization (CAA) Resource Record

DNS Certification Authority Authorization (CAA) is an Internet security policy mechanism which allows domain owners to declare which certificate authorities are allowed to issue a certificate for a domain. CAA is a standard that brings an extra security confirmation for your web domains. The DNS CAA record is specified in RFC 6844.

The CAA record (DNS type code 257) consists of the following:

- **Flag**—An unsigned integer between 0-255.
- **Tag**—The RFC currently defines 3 available tags:
 - **issue**—Explicitly authorizes a single certificate authority to issue a certificate (any type) for the hostname.
 - **issuewild**—Explicitly authorizes a single certificate authority to issue a wildcard certificate (and only wildcard) for the hostname.
 - **iodef**—Specifies a URL to which a certificate authority may report policy violations.
- **Value**—A character-string.



Note The CAA record consists of a flags byte and a tag-value pair referred to as a ‘property’. Multiple properties may be associated with the same domain name by publishing multiple CAA RRs at that domain name.

Examples of CAA records:

```
example.com. CAA 0 issue "letsencrypt.org"
example.com. CAA 0 issuewild "comodoca.com"
```

In Cisco Prime Network Registrar, you can add, maintain, and query for CAA RR type using web UI and CLI commands. Add a CAA DNS record for each Certificate Authority (CA) that you plan to use for your domain.

The rdata part of CAA is *flag tag value*.

where:

- *flag*—A byte size. Currently, bit 0 and bit 7 are used, and other bits are reserved for future use (supported values: 0, 1, and 128).
- *tag*—A non-zero sequence of US-ASCII letters and numbers. The tag length must be at least 1 and no more than 15.
- *value*—A character-string.

Local and Regional Web UI

To add a CAA RR type on the DNS server, do the following:

-
- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
 - Step 2** Click the **Resource Records** tab.
 - Step 3** Enter the owner of the record in the **Name** field.
 - Step 4** Enter the **TTL**.
 - Step 5** Select **CAA** from the **Type** drop-down list.
 - Step 6** Enter the data as a string in the **Data** field by following the correct syntax.

Example:

```
0 issue "letsencrypt.org"
```

- Step 7** Click **Add Resource Record**.
-

CLI Commands

Use the **addRR**, **removeRR**, and **modifyRR** commands to add, delete, and modify CAA records. For example:

```
nrcmd> zone example.com addRR test1 CAA 0 issue comodoca.com
nrcmd> zone example.com removeRR test1
nrcmd> zone example.com modifyRR test1 CAA 0 issue comodoca.com rdata="0 issue
new-comodoca.com" ttl=86400
```

Uniform Resource Identifier (URI) Resource Records

Cisco Prime Network Registrar supports Uniform Resource Identifier (URI) resource records. URI is a string of characters used to identify a resource on the internet either by location or by name, or both. To guarantee uniformity, all URIs follow a predefined set of syntax rules, but also maintain extensibility through a separately defined hierarchical naming scheme (for example, `http://`). In DNS, a URI record (RFC 7553) is a means for

publishing mappings from hostnames to URIs. The clients use the URI records for applications where the relevant protocol/service to be used is known.

In Cisco Prime Network Registrar, you can add, maintain, and query for URI RR type using web UI and CLI commands. This helps to get an explicit URI of the actual connection that is to be made, by providing protocol/service and domain names as the input. You can also synchronize the zone with the URI RR with the HA partner and then query either partners for the URI RR.

Querying for URI RRs is not replacing querying for NAPTR RRs. Instead, the URI RR type provides a complementary mechanism to be used, when one already knows what service field is interesting. With it, one can directly query for the specific subset of the large RRSets returned when querying for NAPTR RRs.

The URI record (DNS type code 256) is expressed in the following format:

```
_service._proto.name. TTL class URI priority weight target
```

where:

- *service*—The symbolic name of the desired service.
- *proto*—The transport protocol of the desired service; this is usually either TCP or UDP.
- *name*—The domain name for which this record is valid, ending in a dot.
- *TTL*—Standard DNS time to live field.
- *class*—Standard DNS class field (this is always IN).
- *priority*—The priority of the target URI in this RR. Its range is 0-65535. Lower the value means, it is more preferred.
- *weight*—A relative weight for records with the same priority. Its range is 0-65535. Higher value means, it is more preferred.
- *target*—The URI of the target, enclosed in double-quotes. The length of this field must be greater than zero.

Example of a URI record:

```
_ftp._tcp IN URI 10 1 "ftp://ftpl.example.com/public"
```

Local and Regional Web UI

To add a URI RR type on the Authoritative DNS Server, do the following:

-
- Step 1** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
 - Step 2** Click the **Resource Records** tab.
 - Step 3** Enter the owner of the record in the **Name** field.
 - Step 4** Enter the **TTL**.
 - Step 5** Select **URI** from the **Type** drop-down list.
 - Step 6** Enter the data as a string in the **Data** field by following the correct syntax.

Example:

```
10 1 "ftp://ftpl.example.com/public"
```

Step 7 Click **Add Resource Record**.**CLI Commands**

Use the **addRR**, **removeRR**, and **modifyRR** commands to add, delete, and modify URI records. For example:

```
nrcmd> zone example.com addRR _ftp._tcp URI 10 1 "ftp://ftp1.example.com/public"
nrcmd> zone example.com removeRR _ftp._tcp URI 10 1 "ftp://ftp1.example.com/public"
nrcmd> zone example.com modifyRR _ftp._tcp URI 10 1 "ftp://ftp1.example.com/public"
rdata="11 1 ftp://ftp1.example.com/public"
```



CHAPTER 13

Managing Hosts

This chapter explains how to configure hosts in DNS zones. Before you proceed with the concepts in this chapter, read [Managing Zones, on page 129](#) which explains how to set up the basic properties of a primary and secondary DNS server and its zones.

- [Adding Hosts in Zones, on page 175](#)
- [Adding Additional RRs for the Host, on page 176](#)
- [Editing Hosts, on page 176](#)
- [Removing Hosts, on page 177](#)

Adding Hosts in Zones

You can manage the resource records for a host by configuring the host rather than the individual RRs. When you define a host, the DNS server automatically creates an Address (A) RR in IPv4, or an AAAA RR in IPv6, for each address you specify. If you specify one or more aliases for the host, the server also creates a Canonical Name (CNAME) RR for each alias. You can also have the server create a Pointer (PTR) RR for the host in the reverse zone for the host, if the reverse zone exists.

Local Basic or Advanced Web UI

- Step 1** From the **Design** menu, choose **Hosts** under the **Auth DNS** submenu to open the List/Add Hosts for Zone page.
- Tip** You can sort by hostname, IP address, IPv6 address (if appropriate), or alias by clicking the corresponding column heading on the List/Add Host for Zone page. However, for zones with a large number of hosts (more than 50,000), restrict the sort to the hostname. Sorting based on IP address or alias can take significantly longer, and could fail if you exceed the memory capacity of the CCM server.
- Step 2** Enter the name of the host and its IPv4 or IPv6 address or comma-separated addresses.
- Step 3** If the host has alias names, enter a comma-separated list.
- Step 4** If you want to create a corresponding Pointer (PTR) RR for the host and you know that the reverse zone for the host exists, check the **Create PTR Records?** check box.
- Step 5** Click **Add Host**.
- Step 6** To confirm, from the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu to open the List/Add Forward Zones page.
- Step 7** Click the **Resource Records** tab to view RRs for the selected zone.

Note If you want to view the list of hosts for a particular zone, click the **Hosts** tab.

CLI Commands

To create A RRs, alias RRs, and PTR RRs for existing reverse zones in a single operation, use **zone name addHost** *hostname address alias* for each host. To list the created zones, use **zone name listHosts**.

Adding Additional RRs for the Host

You add additional RRs for the host based on the dns edit mode you chose, either staged or synchronous. For details, see [Adding Resource Record to Zone, on page 164](#).

Reload the DNS server if you want these RRs to become active server RRs.

Local Basic or Advanced Web UI

For example, to add additional CNAME RRs, add the alias hostname in the Name field under the Resource Records tab of the List/Add Forward Zones page, choose **CNAME** from the Type drop-down list, add the canonical name of the host in the Data field, then click **Add Resource Record**. Note that the DNS specification does not allow a CNAME RR with the same name as that of another RR.

For an MX RR, add the origin hostname in the Name field; choose **MX** from the Type drop-down list; add the integer preference value, a space, and the domain name of the mail exchanger for the origin host in the Data field; then click **Add Resource Record**. These entries should appear in the list at the bottom of the page.

CLI Commands

To create a CNAME record, use **zone name addRR** *alias CNAME canonical* for protected RRs or **zone name addDNSRR** *alias CNAME canonical* for unprotected RRs.

To create an MX record, use **zone name addRR** *hostname MX preference mxname* for protected RRs or **zone name addDNSRR** *hostname MX preference mxname* for unprotected RRs.

Editing Hosts

Editing a host involves:

- Adding additional addresses or aliases
- Modifying its Resource Records (RRs)

Local Basic or Advanced Web UI

Step 1 From the **Design** menu, choose **Hosts** under the **Auth DNS** submenu to open the List/Add Hosts for Zone page.

If you have multiple zones configured, select the zone from the list of zones in the Hosts pane on the left.

Step 2 Click the hostname to add additional IP addresses or aliases, and click **Save**.

Step 3 To modify the RRs, click the **Edit RRs** button to open the Edit View RR List page.

CLI Commands

To edit the host, you must remove its RRs by using **zone name removeRR name type data** or **zone name removeDNSRR name type data**, and then re-enter the RRs using **zone name addRR name ttl class type data** or **zone name addDNSRR name ttl type data**.

Removing Hosts

Removing a host removes all A, CNAME, and PTR RRs for that host.

Local Basic or Advanced Web UI

On the List/Add Hosts in Zone page (see [Editing Hosts, on page 176](#) for the possible ways to get there), click the **Delete** icon next to the host you want to remove, then confirm the deletion.

CLI Commands

Remove the host by using **zone name removeHost**, then re-add it by using **zone name addHost**.



CHAPTER 14

Authoritative DNS Metrics

Following authoritative DNS metric elements are available in the dashboard. For the complete list of Authoritative DNS server statistics, see the *"DNS Statistics" section under the "Server Statistics" appendix of Cisco Prime Network Registrar 11.0 Administration Guide*.

- [DNS General Indicators, on page 179](#)
- [DNS Inbound Zone Transfers, on page 180](#)
- [DNS Network Errors, on page 180](#)
- [DNS Outbound Zone Transfers, on page 181](#)
- [DNS Queries Per Second, on page 181](#)
- [DNS Related Servers Errors , on page 181](#)

DNS General Indicators

The DNS General Indicators dashboard element rendered as a table shows the server state, its last and startup reload time, the number of zones per server, and the total resource record (RR) count. The table is available if you choose **DNS Metrics: DNS General Indicators** in the Chart Selections page.

The resulting table shows:

- **Server State**—Up or Down (based on whether statistics are available), and how long the server has been in this state.
- **Last Reload**—How long since the last server reload.
- **Total Zones**—Number of configured zones.
- **Total RRs**—Number of resource records.

How to Interpret the Data

The data in this chart shows general server health and operational duration. The objective is to make decisions about the server, such as whether it might be time for another reload, perhaps warranted by the number of configured zones.

Troubleshooting Based on the Results

If the server state is Down, all the DNS chart indicators show a red status box, so no data will be available. In the case of a server that is down, restart the server. The number of zones indicated might also require some evaluation and possible reconfiguration.

DNS Inbound Zone Transfers

The DNS Inbound Zone Transfers dashboard element rendered as an area chart tracks the rate of change in full and incremental inbound zone transfer responses, and any associated errors. The chart is available if you choose **DNS Metrics: DNS Inbound Zone Transfers** in the Chart Selections page.

The resulting area chart plots the following trends:

- **Full Response**—Number of full inbound zone transfers (AXFRs in).
- **Incremental Responses**—Number of incremental inbound zone transfers (IXFRs in).
- **Authorization Errors**—Number of refused responses (xfer-in-auth-errors).
- **Failed Attempts**—Number of failures other than refusals (xfer-failed-attempts).
- **Exceed Max Transfers In**—Number of times that the concurrent inbound transfers reach the maximum limit.

How to Interpret the Data

This chart is useful in gauging if inbound zone transfers to a secondary DNS server are occurring as predicted and if there are any authentication or failed transfer attempts in the process. The most significant indicator is the trend in the number of inbound zonended transfers denied for lack of permission, for not being authorized for the zone, or for other reasons.

Troubleshooting Based on the Results

Check the primary and secondary server configurations if there are errors or exceeded limits in the inbound zone transfers.

DNS Network Errors

The DNS Network Errors dashboard element rendered as an area chart tracks the rate of change in DNS server network errors. The chart is available if you choose **DNS Metrics: DNS Network Errors** in the Chart Selections page.

The resulting area chart plots the following trends:

- **Query Error Packets/Query Responses**—Ratio of query error packets over responses. Responses consist of:
 - Authoritative
 - Authoritative no-such-name
 - Authoritative no-such-data
 - Nonauthoritative
 - Nonauthoritative no-such-data
 - Requests refused
- **Non Error Dropped Packets/Query Responses**—Ratio of nonerror dropped packets (queries dropped) over responses.
- **Update Errors/Updates**—Ratio of DNS Update errors over total updates.

How to Interpret the Data

This chart indicates query and response errors as an indication of the health of the server.

Troubleshooting Based on the Results

Check the DNS server network configuration if errors are increasing.

DNS Outbound Zone Transfers

The DNS Outbound Zone Transfers dashboard element rendered as an area chart tracks the rate of change in full and incremental outbound zone transfer responses, and any associated errors. The chart is available if you choose **DNS Metrics: DNS Outbound Zone Transfers** in the Chart Selections page.

The resulting area chart plots the following trends:

- **Full Responses**—Number of full outbound zone transfers (AXFRs out).
- **Incremental Responses**—Number of incremental outbound zone transfers (IXFRs out).
- **Authorization Errors**—Number of unauthorized (refused) zone transfer requests.
- **Exceed Max Transfers Out**—Number of failed outbound transfers that exceed the maximum limit.
- **Other Errors**—Number of other outbound transfer errors that are not authorization errors.

How to Interpret the Data

This chart is useful in gauging if outbound zone transfers to a secondary DNS server are occurring as predicted and if there are any authorizations or failed transfer attempts in the process. The most significant indicator is the trend in the number of outbound zone transfers denied for lack of permission or for not being authorized for the zone.

Troubleshooting Based on the Results

Check the primary and secondary server configurations if there are errors or exceeded limits in the outbound zone transfers.

DNS Queries Per Second

The DNS Queries Per Second dashboard element rendered as an area chart displays queries per second for the Authoritative DNS server. This chart is available if you choose **DNS Metrics: DNS Queries Per Second** in the Chart Selections page.

DNS Related Servers Errors

The DNS Related Servers Errors dashboard element rendered as an area chart tracks the rate of change in DNS related server errors. The chart is available if you choose **DNS Metrics: DNS Related Servers Errors** in the Chart Selections page.

The resulting area chart plots the following trends:

- **Referral Timeouts/Referrals**—Ratio of referral timeouts over referrals.

- **Failed Responses/Total Incoming Zone Transfer Requests**—Ratio of failed responses over incoming zone transfer requests.
- **TSIG Errors/TSIG Attempts**—Ratio of transaction signature (TSIG) errors (bad times, keys, or signatures) over total TSIG attempts (successfully received packets).

How to Interpret the Data

This chart indicates the health of connections and data transfers with related DNS servers. All three chart lines can have diagnostic significance.

Troubleshooting Based on the Results

Check the configurations and connectivity of the related servers in HA DNS relationships if errors are increasing.



APPENDIX A

Resource Records

This chapter lists all the resource record types supported in Cisco Prime Network Registrar.

- [Resource Records, on page 183](#)

Resource Records

Resource records comprise the data within a DNS zone. There is no fixed limit to the number of resource records a zone can own. In general, there can be zero, one, or more resource records of a given type. However, there are constraints on the number of certain types of records a zone can have.

All resource records have these required entries:

- **Name**—Name (host) that owns the record, such as example.com.
- **Class (not required for all formats)**—DNS supports only the IN (Internet) class of record.
- **TTL**—Amount of time to store the record in cache, in seconds. If you do not include a TTL, Cisco Prime Network Registrar uses the zone default TTL, defined in the SOA resource record.
- **Type**—Type of the record, such as A, NS, SOA, MX, and so on. There are many types that various RFCs define, although ten or fewer are in common use.
- **Record data**—Data types whose format and meaning varies with record type.

The following table lists all the resource record types Cisco Prime Network Registrar supports. It provides the field syntax and the field descriptions, as well as how the fields are represented in the Cisco Prime Network Registrar GUI.

Table 46: Resource Records

Record	No.	Name	Syntax and Description	RFC
A	1	Host Address—Name-to-address mapping for the zone	<p><i>name ttl class A address</i></p> <p>Web UI: Add or Edit Host for Zone page: Hostname, IP Address or Resource Records for Zone page: Name, TTL, Type, Data</p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host123 3600 IN A 192.168.40.123</pre>	1035

Record	No.	Name	Syntax and Description	RFC
A6	38	IPv6 Address— (Obsolete; use AAAA records instead)	<p><i>name ttl class A6 address</i></p> <p>In the data, the suffix address is an IPv6 address encoded in network order (high-order octet first). There must be exactly enough octets in this field to contain a number of bits equal to 128 minus prefix length, with 0 to 7 leading pad bits to make this field an integral number of octets. Pad bits, if present, must be set to zero when loading a zone file and ignored on reception. For example:</p> <p>2001:0:734c:c0::</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=A6, Data=<i>prefixlength suffixaddr prefixname</i>, with data in the form:</p> <p>CLI Command:</p> <pre>0 2345:00c1:ca11:0001:1234:5678:9abc:def0</pre> <pre>nrcmd> zone example.com addRR host456 A6 0 1345:c1:ca11:1:1234:5678:9abc:def0</pre>	6563
AAAA	28	IPv6 Address	<p><i>name ttl class AAAA address</i></p> <p>Data is the IPv6 address format of eight sets of four hexadecimal digits, separated by colons. The first set of four digits is the high-order 16 bits of the address. You can omit leading zeros in sets and omit a value in a set if the value of the set is zero.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=AAAA, Data=<i>address</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host456 AAAA 1345:c1:ca11:1:1234:5678:9abc:def0</pre>	3596
AFSDB	18	Andrew File System (AFS) Data Base	<p><i>name ttl class AFSDB subtype hostname</i></p> <p>Subtype is either 1—AFS cell database server, or 2—DCE authentication name server. Hostname is the domain name of host that has a server for the cell named by the owner.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=AFSDB, Data=<i>subtype hostname</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host4 AFSDB 1 AFSDBhost.example.com.</pre>	1183

Record	No.	Name	Syntax and Description	RFC
AXFR	252	Authoritative Zone Transfer	<p>Transfer entire zone file from the primary name server to secondary name servers. AXFR records are not used in ordinary zone files. Rather, they are used on a secondary DNS server to replicate the zone file from a primary DNS server.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=AXFR, Data=Auth Zone Transfer</p>	1995
CAA	257	Certification Authority Authorization	<p><i>name ttl class CAA flag tag value</i></p> <p>Data contains <i>flag</i>, <i>tag</i>, and <i>value</i>. Where:</p> <ul style="list-style-type: none"> • <i>flag</i>—A byte size. Currently, bit 0 and bit 7 are used, and other bits are reserved for future use (supported values: 0, 1, and 128). • <i>tag</i>—A non-zero sequence of US-ASCII letters and numbers. The tag length must be at least 1 and no more than 15. • <i>value</i>—A character-string. <p>Web UI: Resource Records for Zone page: Name, TTL, Type=CAA, Data= <i>flag tag value</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR test1 CAA 0 issue comodoca.com</pre>	6844
CNAME	5	Canonical Name— Aliases or nicknames	<p><i>alias ttl class CNAME canonicalname</i></p> <p>You cannot have any other resource records associated with a CNAME. Aliases are useful when you want the outside world to know a single, easily remembered name. You can also use aliases when a host changes its name. In that case, ensure that you have a CNAME pointer so that when people use the original name, it can be resolved to the newer one.</p> <p>Web UI: Resource Records for Zone page: Name=<i>alias</i>, TTL, Type=CNAME, Type, Data=<i>canonicalname</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host456 CNAME host1234</pre>	1035
DHCID	49	Dynamic Host Configuration Identifier— (RFC 4701)	<p><i>name ttl class DHCID data</i></p> <p>The DNS server uses this RR to allow DHCP clients and servers to update DNS automatically. This RR is not user-configurable. The data is the result of a one-way hash computation of the client message and the domain name. Sample RR output for an IPv6 address:</p> <pre>chi6.example.com IN DHCID (AAIBY2/AuCccgoJbaxcQc9TUapptP6910jxfNuVAA2kjEA=)</pre>	4701

Record	No.	Name	Syntax and Description	RFC
HINFO	13	Host Info— Hardware and software information for the host	<p><i>name ttl class HINFO cpu os</i></p> <p>Data is the hardware (CPU) and operating system.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=HINFO, Data=<i>cpu os</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host5 HINFO CPU1 OS2</pre>	1035
ISDN	20	Integrated Services Digital Network (ISDN) Address	<p><i>name ttl class ISDN ISDNnumber [subaddr]</i></p> <p>Data is the ISDN number of the owner and Direct Dial In, if any, and an optional ISDN subaddress string</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=ISDN, Data=<i>ISDNnumber [subaddr]</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host6 ISDN ISDN88888</pre>	1183
IXFR	251	Incremental Zone Transfer	<p>Incremental transfer (IXFR) is an efficient means to transfer changes in zones from IXFR servers to IXFR clients. As proposed it is more efficient mechanism as it transfers only the changed portion(s) of a zone. The goal of these mechanism is to enable a set of DNS name servers to remain coherently authoritative for a given zone.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=IXFR</p>	1995
MB	7	Mailbox Domain Name	<p><i>name ttl class MB mbox</i></p> <p>Data is the domain name of the host with the specified mailbox.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=MB, Data=<i>mbox</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host7 MB mailbox.example.com.</pre>	1035
MD	3	Mail Destination- (Obsolete; use MX instead)	A mail destination (OBSOLETE - use MX)	1035
MF	4	Mail Forwarder- (Obsolete; use MX instead)	A mail forwarder (OBSOLETE - use MX)	1035

Record	No.	Name	Syntax and Description	RFC
MG	8	Mail Group Member	<p><i>name ttl class MG mgroup</i></p> <p>Data is the domain name of the mailbox group (mailing list).</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=MG, Data=<i>mgroup</i></p> <p>CLI Command:</p>	1035
MINFO	14	Mailbox Info	<p><i>name ttl class MINFO respmbx errormsg</i></p> <p>Data is the mailbox responsible for the mailing list, and the mailbox to receive error messages.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=MINFO, Data=<i>respmbx errormsg</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host7 MINFO resp.example.com. error.example.com.</pre>	1035
MR	9	Mail Rename	<p><i>name ttl class MR newmbx</i></p> <p>Data is the mailbox name to rename the owner mailbox.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=MR, Data=<i>newmbx</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host7 MR renamemb.example.com.</pre>	1035
MX	15	Mail Exchanger— Where to deliver the mail for a domain name	<p><i>name ttl class MX pref mxname</i></p> <p>Data is the preference value (16-bit integer for the preference for the record, with lower values having preference), and the domain name of the mail exchanger for the owner.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=MX, Data=<i>pref mxname</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host8 MX 10 exchanger.example.com.</pre>	1035

Record	No.	Name	Syntax and Description	RFC
NAPTR	35	Naming Authority Pointer— Produces a new domain label or Uniform Resource Identifier (URI). You can then use DNS to look up services for many resource names that are not in domain name syntax.	<p><i>name ttl class NAPTR order pref flags serv regexp replace</i></p> <ul style="list-style-type: none"> <i>order</i>—16-bit integer for the order in which to process the NAPTR records to ensure the correct ordering of rules, with low numbers processed before high numbers. <i>pref</i>—16-bit unsigned integer for the order in which to process NAPTR records with equal <i>order</i> values, with low numbers processed before high numbers. <i>flags</i>—Character-string containing flags to control aspects of rewriting and interpreting fields, single characters from the set [A-Z0-9] (not case-sensitive); the S, A and U flags denote a terminal lookup, the P flag says that the remainder of the application-side algorithm should be carried out protocol-specific. <i>serv</i>—Valid protocols or services. <i>regexp</i>—String containing a substitution expression applied to the original string held by the client to construct the next domain name to look up. (For common regex usage, see the "Common Regex Values" table in <i>Cisco Prime Network Registrar 11.0 Administration Guide</i>). <i>replace</i>—Next FQDN to query for NAPTR, SRV, or address records, depending on the value of the <i>flags</i> field. <p>Web UI: Resource Records for Zone page: Name, State, TTL, Type=NAPTR, Data=<i>order pref flags service regexp replace</i></p> <p>CLI Command:</p> <pre>nrcmd> zone 8.6.4.e164.arpa addRR 4.3.2.1.6.7.9 naptr 100 10 u sip+E2U /^.*\$/sip:info@tele2.se/ .</pre>	2915
NS	2	Name Server— Authoritative server for the zone	<p><i>name ttl class NS nameserver</i></p> <p>Machines that provide name service must not reside in the owner domain. For each domain, you must have at least one NS record. NS records for a domain must exist in both the zone that delegates the domain and in the domain itself. NS record names must have an equivalent A record (they cannot point to an alias).</p> <p>Web UI: Add or Edit Zone page Nameservers: NS TTL, Add Nameserver</p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR @ NS DNSserv2.example.com.</pre>	1035

Record	No.	Name	Syntax and Description	RFC
NSAP	22	Network Service Access Point (NSAP) Address	<p><i>name ttl class NASP NSAPaddr</i></p> <p>Data is the <i>NSAPaddr</i>—Octet values assigned by the assigning authority, a character string of the type used in TXT and HINFO records (see RFC 1706).</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=NSAP, Data=<i>NSAPaddr</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host10 NSAP 39840f80005a0000000001e13708002010726e00</pre>	1706
NSEC	47	Next Secure record	<p>Part of DNSSEC—used to prove a name does not exist. Uses the same format as the (obsolete) NXT record.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=NSEC, Data=<i>Next Secure record</i></p>	
OPT	41	DNS EDNS(0) Options	<p>This is a "pseudo DNS record type" needed to support EDNS. An OPT pseudo-RR (sometimes called a meta-RR) MAY be added to the additional data section of a request. If an OPT record is present in a received request, compliant responders MUST include an OPT record in their respective responses.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=OPT</p>	
PTR	12	Pointer—Reverse mapping	<p><i>name ttl class PTR dname</i></p> <p>Data is the domain name of host having the reverse record indicated by the owner. PTR records are used for reverse mapping, specifically in the in-addr.arpa zones for translation of addresses to names. PTRs use official names, not aliases. The name in a PTR record is the local IP address portion of the reverse name.</p> <p>Web UI: Resource Records for Zone page: Name, State, TTL, Type=PTR, Data=<i>dname</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR 45.40.168.192.in-addr.arpa. PTR host1234</pre>	1035

Record	No.	Name	Syntax and Description	RFC
RP	17	Responsible Person	<p><i>name ttl class RP mbox txt host</i></p> <p>Data is the domain name of the mailbox for the responsible person, and the domain name of host where TXT records exist.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=RP, Data=<i>mbox txt host</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host7 RP resp.example.com. text.example.com.</pre>	1183
RT	21	Route Through	<p><i>name ttl class RT pref intermediate host</i></p> <p>Data is the <i>pref</i>—16-bit integer for preference to give to this record among others of the same owner, and <i>intermediate host</i>—domain name of the host serving as intermediate to reach the owner.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=RT, Data=<i>pref intermediate host</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host7 RT 10 routthru.example.com.</pre>	1183
SOA	6	Start of Authority— Every zone must have a single SOA record	<p><i>name ttl class SOA primeserver host admin (serial refresh retry expire minimum)</i></p> <p>Web UI: Add or Edit Zone page SOA Attributes: Serial Number, SOA TTL, Nameserver, Contact E-Mail, Secondary Refresh, Secondary Retry, Secondary Expire, Minimum TTL</p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR @ 172800 IN SOA ns hostadmin 1 10800 3600 604800 86400</pre>	1035
SPF	99	Sender Policy Framework	<p>Sender Policy Framework (SPF) record is a type of Domain Name Service (DNS) TXT record that identifies which mail servers are permitted to send email on behalf of your domain. The purpose of an SPF record is to detect and prevent spammers from sending messages with forged From addresses on your domain.</p> <p>SPF records are defined as a single string of text.</p>	7208

Record	No.	Name	Syntax and Description	RFC
SRV	33	Service Location	<p><i>name ttl class SRV priority weight port target</i></p> <ul style="list-style-type: none"> • <i>priority</i> —16-bit priority to give the record among the owner SRV records. • <i>weight</i> —16-bit load to give the record at the same priority level. • <i>port</i> —16-bit port on which to run the service. • <i>target</i> —Domain name of host running on the specified port. <p>Administrators can use several servers for a single domain, move services between hosts with little difficulty, and designate some hosts as primary servers for a service and others as backups. Clients ask for a specific service or protocol for a domain and receive the names of any available servers.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=SRV, Data=<i>priority weight port target</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host2 SRV 10 1 60 host7.example.com.</pre>	2782
TSIG	250	Transaction Signature	<p>Key name, which must be unique on client and server. Can be used to authenticate dynamic updates as coming from an approved client, or to authenticate responses as coming from an approved recursive name server similar to DNSSEC.</p>	2854
TXT	16	Text	<p><i>name ttl class TXT textstring</i></p> <p>Data is one or more text character strings that can contain any type of information.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=TXT, Data=<i>textstring</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host2 TXT "this message"</pre>	1035

Record	No.	Name	Syntax and Description	RFC
URI	256	Uniform Resource Identifier	<p><i>name ttl class URI priority weight target</i></p> <p>Data contains <i>priority</i>, <i>weight</i>, and <i>target</i>. Where:</p> <ul style="list-style-type: none"> • <i>priority</i>—The priority of the target URI in this RR. Its range is 0-65535. Lower the value means, it is more preferred. • <i>weight</i>—A relative weight for records with the same priority. Its range is 0-65535. Higher value means, it is more preferred. • <i>target</i>—The URI of the target, enclosed in double-quotes. The length of this field must be greater than zero. <p>Web UI: Resource Records for Zone page: Name, TTL, Type=URI, Data=<i>priority weight target</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR _ftp._tcp URI 10 1 "ftp://ftp1.example.com/public"</pre>	7553
WKS	11	Well Known Services	<p><i>name ttl class WKS addr protocol servicelist</i></p> <ul style="list-style-type: none"> • <i>addr</i> —32-bit IP address. • <i>protocol</i> —8-bit IP protocol number, which can be TCP or UDP. • <i>servicelist</i> —Variable-length bit map in 8-bit multiples of services, which can be TIME, TELNET, FTP, or SMTP. <p>Web UI: Resource Records for Zone page: Name, TTL, Type=WKS, Data=<i>addr protocol servicelist</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host8 WKS 192.168.40.56 TCP TELNET</pre>	1035
X25	19	X.25 Address	<p><i>name ttl class X25 PSDNaddr</i></p> <p>Data is the character string of the Public Switch Data Network (PSDN) address in the X.121 numbering plan associated with the owner.</p> <p>Web UI: Resource Records for Zone page: Name, TTL, Type=X25, Data=<i>PSDNaddr</i></p> <p>CLI Command:</p> <pre>nrcmd> zone example.com addRR host9 IN X25 311061700956</pre>	1183



APPENDIX **B**

DNS Anycast with Cisco Prime Network Registrar

Anycast is a network and routing mechanism that enables a packet from a single client to go to one of many servers offering the same service. All the servers in the Anycast group are configured with the same Anycast IP address and the packet is routed from the client to the closest server by the best path as determined by routing algorithms. Anycast routing enables several important capabilities such as seamless redundancy, load balancing, and horizontal scaling by grouping multiple servers as a single service. Anycast DNS is simply an implementation of Anycast for DNS services. Anycast is used in conjunction with a routing protocol, such as Border Gateway Protocol (BGP) to advertise the availability of the service to an adjoining router, which makes the Anycast DNS to work effectively.

This chapter provides the knowledge and tools to configure Cisco Prime Network Registrar DNS services using Anycast.

- [Basic Requirements for DNS Anycast, on page 193](#)
- [Anycast Routing, on page 194](#)
- [Script, on page 195](#)
- [Router Configuration, on page 195](#)
- [Sample Anycast Configuration Using BGP, on page 195](#)
- [Network Router Configuration, on page 196](#)
- [Configure FRRouting on DNS Servers, on page 197](#)
- [Configure Quagga on DNS Servers, on page 199](#)
- [Run Diagnostics on Router, on page 200](#)
- [Monitor BGP Traffic Logs, on page 201](#)
- [Configure DNS Zones, on page 202](#)

Basic Requirements for DNS Anycast

The following is a list of requirements and recommendations for supporting Anycast DNS:

- Clients should be configured to resolve DNS queries via the Caching DNS server's Anycast address(es).
- Nameservers should advertise their Anycast address in NS and A RRs.
- Nameservers should listen to DNS queries on the Anycast IP addresses.
- Nameservers should be configured with at least one Anycast IP address on a loopback interface.

- Additionally, the server should be configured with a management IP, which can be either a physical or an additional loopback interface.
- At least one physical IP must be defined on the DNS server for the exchange of routing information, as well as, system access and maintenance in the absence of the routes to the Anycast IP address(es).
- Nameservers should be configured to use the physical or management IP addresses for zone transfers, zone updates, and/or query source to ensure that these updates go to the intended server.
- Nameservers should Inject Anycast IP address(es) into the routed network using routing protocols such as RIP, OSPF, or BGP.

Anycast Routing

Anycast can be manually configured, it is best implemented using routing protocols such as BGP or OSPF, which announces the Anycast destination address to its gateway router. Using a routing protocol to announce availability of the DNS service ensures that routers do not send DNS queries into a blackhole if the service goes down. As the Cisco Prime Network Registrar DNS application does not have routing capabilities, some code, external to the DNS application must be added to the DNS environment (physical server or virtual machine). The prominent and open source products are FRRouting (FRR) for RHEL/CentOS 8.x and Quagga for RHEL/CentOS 7.x.

FRRouting



Note With RHEL/CentOS 8.x, use FRR.

FRR is an IP routing protocol suite for Linux and Unix platforms which includes protocol daemons for BGP, IS-IS, LDP, OSPF, PIM, and RIP.

FRR is forked from Quagga which is another routing protocol suite for Linux. FRR includes the fundamentals that made Quagga so popular as well as many enhancements that greatly improve on that foundation.

FRR does not ship with Cisco Prime Network Registrar. For more information about FRR, see the FRR documentation.

Quagga



Note With RHEL/CentOS 7.x, use Quagga.

Quagga is a routing software suite, providing implementations of OSPFv2, OSPFv3, RIP v1 and v2, RIPng and BGP-4 for Unix platforms, Linux, Solaris, and NetBSD. The solution described in this chapter uses BGP.

The Quagga architecture consists of a core daemon, zebra, which acts as an abstraction layer to the underlying Linux kernel and presents the Zserv API over an Unix or TCP stream to Quagga clients. It is these Zserv clients, which typically implement a routing protocol and communicate routing updates to the zebra daemon.

Quagga daemons are configurable via a network accessible CLI (called **vty**). The CLI follows a style similar to that of other routing software. There is an additional tool included with Quagga called **vtysh**, which acts as a single cohesive front-end to all the daemons, allowing one to administer nearly all aspects of the various Quagga daemons in one place.

Quagga does not ship with Cisco Prime Network Registrar. For more information about Quagga, see the Quagga documentation.

Script

A sample python script is included with Cisco Prime Network Registrar installation and is located at:

- FRR:

```
/opt/nwreg2/local/examples/dns/python/dns_anycast_bgp_frr.py
```

- Quagga:

```
/opt/nwreg2/local/examples/dns/python/dns_anycast_bgp.py
```

The script starts and stops FRR/Quagga, and monitors the DNS service by sending DNS queries to ensure that it is operational. When FRR/Quagga is started, its FRR/Quagga daemon sends an Anycast advertisement to the connected router making the DNS service available over the Anycast address. If the DNS server does not respond to queries from the script, the script will stop the FRR/Quagga daemon. Stopping FRR/Quagga breaks the TCP connection and the router will stop receiving BGP keep-alive messages. The router will then remove the DNS service from its Anycast group, and start sending DNS queries to the next closest and available DNS service. If the DNS server responds to queries from the script, the script checks to see if the FRR/Quagga daemons are running. If the daemons are not running, then the script starts the daemons.

It is recommended to copy the sample script to a different location, set up a cron job to periodically run the script to check the status of the DNS server (recommended interval is 5 minutes) and start or stop the BGP daemon accordingly. An example of a cron job is outside the scope of this solution.

Router Configuration

Your configuration will probably be different based on your network requirements and variations in addressing schemes.

Sample Anycast Configuration Using BGP

This section describes the basic setup and configuration of Anycast using BGP on a Cisco router and FRR/Quagga host-based routing software. The purpose of this section is not to instruct administrators on the configuration of routers and BGP, but to show a configuration that was successfully tested in the Cisco Prime Network Registrar labs. Note that your network requirements may be different.

BGP is a standardized exterior gateway protocol designed to exchange routing and reachability information among Autonomous Systems (AS) on the Internet. This configuration uses a single AS this recipe is not intended to be solution deployed across Autonomous Systems.

Perform the following steps on the hosts DNS-1 and DNS-2:

For FRR:**Install FRR Routing Software**

Install FRR on the same system that is running Cisco Prime Network Registrar. This will install FRR package like the following:

```
frr-7.0-5.el8.x86_64
```

For Quagga:**Install Quagga Routing Software**

Install Quagga on the same system that is running Cisco Prime Network Registrar. This will install Quagga package like the following:

```
quagga-0.99.15-7.el6_3.2.x86_64
```

Create a Loopback Interface

Create a loopback interface alias on the system. Configure the anycast IP address on this loopback interface.

On RHEL, the interface configuration files are located at `/etc/sysconfig/network-scripts`. Create a file in that directory named `ifcfg-lo:0` with the following contents:

```
DEVICE=lo:0
IPADDR=10.10.10.1
NETMASK=255.255.255.255
BOOTPROTO=none
ONBOOT=yes
```

Bring up the new loopback interface using the `ifup lo:0` command.

Network Router Configuration

This router configuration is used in the validation of this DNS Anycast solution. It is provided as a reference to assist in the development of DNS Anycast solution. While it is a complete configuration for this specific solution, it is only intended to be a reference for developing your solution.

```
csr1000v# sh run
Building configuration...
!
interface Loopback0
 ip address 2.2.2.2 255.255.255.255
!
interface GigabitEthernet1
 ip address 10.78.29.77 255.255.255.0 (Router)
 negotiation auto
!
interface GigabitEthernet2
 ip address 10.0.2.1 255.255.255.0 (Client)
 negotiation auto
!
interface GigabitEthernet4 (DNS-2)
 platform ring rx 256
 ip address 10.0.3.1 255.255.255.0
 negotiation auto
!
interface GigabitEthernet5 (DNS-3)
 platform ring rx 256
 ip address 10.0.5.1 255.255.255.0
```

```

negotiation auto
!
router ospf 1
router-id 2.2.2.2(is the loopback IP address)
redistribute bgp 65500 subnets
network 2.2.2.2 0.0.0.0 area 1
network 10.0.6.0 0.0.0.255 area 1
network 10.0.0.0 0.0.255.255 area 1
!
router bgp 65500
bgp log-neighbor-changes
neighbor IBGP peer-group
neighbor IBGP update-source Loopback0
neighbor ANY peer-group
neighbor 1.1.1.1 remote-as 65500
neighbor 1.1.1.1 peer-group IBGP
neighbor 1.1.1.1 update-source Loopback0
neighbor 10.0.3.2 remote-as65500
!(This should be the bgp AS in Quagga for DNS-2)
neighbor 10.0.3.2 peer-group ANY
neighbor 10.0.5.2 remote-as 65500
!(This should be the bgp AS in Quagga for DNS-3)
neighbor 10.0.5.2 peer-group ANY
!
address-family ipv4
redistribute ospf 1
neighbor IBGP next-hop-self
neighbor ANY next-hop-self
neighbor 1.1.1.1 activate
neighbor 10.0.3.2 activate
neighbor 10.0.5.2 activate
exit-address-family
!
virtual-service csr_mgmt
ip shared host-interface GigabitEthernet1
activate
!
ip default-gateway 10.78.28.1
ip forward-protocol nd
!
no ip http server
ip http secure-server
ip route 0.0.0.0 0.0.0.0 10.78.28.1
ip route 10.78.28.0 255.255.254.0 GigabitEthernet1 10.78.28.1
!
ip prefix-list anycast-ip seq 5 permit 10.10.10.1/32
!
control-plane
!
line con 0
stopbits 1
line vty 0 4
login local
!
!
end

```

Configure FRRouting on DNS Servers

Configure the FRR configuration files on both the servers. The following example is for DNS-1. DNS-2 also needs to be configured similarly. The configuration files are located in `/etc/frr`.

There are number of example configuration files in **/etc/frr**: one for each routing protocol that FRR supports; one for zebra, the main process. For enabling Anycast using BGP, we need to configure **zebra.conf**, **bgpd.conf**, and **daemons** file.

Enable zebra and bgpd in Daemons File

```
# cat /etc/frr/daemons
# This file tells the frr package which daemons to start.
watchfrr_enable=yes
watchfrr_options="-r '/usr/lib/frr/frr restart %s' -s '/usr/lib/frr/frr start %s' -k
'/usr/lib/frr/frr stop %s'"
#
zebra=yes
bgpd=yes
ospfd=no
```

FRR Zebra Configuration

```
# cat /etc/frr/zebra.conf
hostname DNS-1
!
password zebra
enable password zebra
!
interface eth0
 ip address 10.0.3.2/24
!
interface lo
 ip address 10.10.10.1/32
!
line vty
!
```



Note Repeat the steps for any other Anycast servers that are part of the group.

FRR BGP Configuration

```
# cat /etc/frr/bgpd.conf
! -- bgp --
!
! BGPd sample configuration file
!
!
hostname DNS-1
password zebra
log stdout
!
router bgp 65500
bgp router-id 10.78.29.79
bgp log-neighbor-changes
network 10.10.10.1/32
timers bgp 4 16
neighbor 10.0.3.1 remote-as 65500
neighbor 10.0.3.1 next-hop-self
neighbor 10.0.3.1 prefix-list DEFAULT in
```

```

neighbor 10.0.3.1 prefix-list ANYCAST out
!
address-family ipv4
network 10.0.3.1/24
neighbor 10.0.3.1 activate
exit-address-family
!
ip prefix-list ANYCAST seq 5 permit 10.10.10.1/32
ip prefix-list DEFAULT seq 5 permit 0.0.0.0/0
line vty
!

```

Start FRR Service

Start the FRR service using the following command:

```
systemctl start frr.service
```

Create Additional IP address on the Loopback Interface

To create the additional IP address on the loopback interface for anycast using FRR, refer the Red Hat documentation.

Restart FRR Service

Restart the FRR service using the following command:

```
systemctl restart frr.service
```

Configure Quagga on DNS Servers

Configure the Quagga configuration files on both the servers. The following example is for DNS-1. DNS-2 also needs to be configured similarly. The configuration files are located in **/etc/Quagga**.

There are number of example configuration files in **/etc/Quagga**: one for each routing protocol that Quagga supports; one for zebra, the main process. For enabling Anycast using BGP, we need to configure **zebra.conf** and **bgpd.conf**.

Quagga Zebra Configuration

```

# cat /etc/quagga/zebra.conf
hostname DNS-1
!
password zebra
enable password zebra
!
interface eth0
 ip address 10.0.3.2/24
!
interface lo
!
line vty
!

```



Note Repeat the steps for any other Anycast servers that are part of the group.

Quagga BGP Configuration

```
# cat /etc/quagga/bgpd.conf
! -- bgp --
!
! BGPd sample configuration file
!
!
hostname DNS-1
password zebra
log stdout
!
router bgp 65500
  bgp router-id 10.78.29.79
  bgp log-neighbor-changes
  network 10.10.10.1/32
  timers bgp 4 16
  neighbor 10.0.3.1 remote-as 65500
  neighbor 10.0.3.1 next-hop-self
  neighbor 10.0.3.1 prefix-list DEFAULT in
  neighbor 10.0.3.1 prefix-list ANYCAST out
!
address-family ipv4
  network 10.0.3.1/24
  neighbor 10.0.3.1 activate
exit-address-family
!
ip prefix-list ANYCAST seq 5 permit 10.10.10.1/32
ip prefix-list DEFAULT seq 5 permit 0.0.0.0/0
line vty
!
```

Start BGP daemon

Start the BGP daemon using the following command:

```
systemctl start bgpd
```

Run Diagnostics on Router

Run diagnostics on the router to make sure that the Anycast is set up properly.

The **sh ip bgp summary** command output shows that router-1 has opened a BGP session with the two neighbors. The value **State/PfxRcd** indicates that the TCP session is up and the routers and hosts are exchanging routes. This field should be a numeric value showing how many route prefixes have been received from the remote neighbor. The example value is 1. At this point, the BGP connection with the DNS servers is in Established state.

The **sh ip bgp summary**:

```
BGP router identifier 2.2.2.2, local AS number 65500
BGP table version is 86, main routing table version 86
```



```

1 network entries using 248 bytes of memory
2 path entries using 240 bytes of memory
1/1 BGP path/bestpath attribute entries using 248 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 736 total bytes of memory
BGP activity 16/15 prefixes, 61/59 paths, scan interval 60 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
1.1.1.1	4	65500	0	0	1	0	0	4w0d	Idle
10.0.3.2	4	65500	137919	129519	86	0	0	1w0d	1
10.0.5.2	4	65500	137923	129519	86	0	0	1w0d	1

The **show ip bgp neighbors** command shows information about the neighbors in detail.

The **show ip route** command should have an entry for the Anycast address and the host via which it is currently routed.

#sh ip route

```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
a - application route
+ - replicated route, % - next hop override, p - overrides from PfR
B 10.10.10.1/32 [200/0] via 10.0.3.2, 00:00:10

```

Monitor BGP Traffic Logs

To monitor the BGP traffic logs on the hosts DNS-1 and DNS-2, use the **telnet localhost bgpd** command.

FRR:

```

Trying ::1...
telnet: connect to address ::1: Connection refused
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.

```

```

Hello, this is FRRouting (version 7.0).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

```

```
User Access Verification
```

```

Password:
dns-anycast-1> enable
dns-anycast-1# terminal monitor
dns-anycast-1# conf t
dns-anycast-1(config)# debug bgp keepalives
dns-anycast-1(config)# 2020/10/27 02:56:22 BGP: : 10.0.3.1 KEEPALIVE rcvd

dns-anycast-1(config)# 2020/10/27 02:56:23 BGP: : 10.0.3.1 sending KEEPALIVE
2020/10/27 02:56:27 BGP: : 10.0.3.1 KEEPALIVE rcvd
2020/10/27 02:56:28 BGP: : 10.0.3.1 sending KEEPALIVE

```

```

2020/10/27 02:56:32 BGP: : 10.0.3.1 KEEPALIVE rcvd
2020/10/27 02:56:33 BGP: : 10.0.3.1 sending KEEPALIVE
2020/10/27 02:56:37 BGP: : 10.0.3.1 KEEPALIVE rcvd
2020/10/27 02:56:38 BGP: : 10.0.3.1 sending KEEPALIVE
2020/10/27 02:56:42 BGP: : 10.0.3.1 KEEPALIVE rcvd
2020/10/27 02:56:43 BGP: : 10.0.3.1 sending KEEPALIVE
2020/10/27 02:56:47 BGP: : 10.0.3.1 KEEPALIVE rcvd
2020/10/27 02:56:48 BGP: : 10.0.3.1 sending KEEPALIVE

```

Quagga:

```

Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
Hello, this is Quagga (version 0.99.15).
Copyright 1996-2005 Kunihiro Ishiguro, et al.
User Access Verification
Password:
DNS-1> enable
DNS-1# terminal monitor
DNS-1# 2016/07/13 15:49:20 BGP: 10.0.5.1 send message type 4, length (incl. header) 19
2016/07/13 15:49:21 BGP: 10.0.5.1 rcv message type 4, length (excl. header) 0
2016/07/13 15:49:25 BGP: 10.0.5.1 send message type 4, length (incl. header) 19
2016/07/13 15:49:27 BGP: 10.0.5.1 rcv message type 4, length (excl. header) 0

```

Configure DNS Zones

While this is the conclusion of setting up the Anycast functionality, administrators will need to complete the configuration of the DNS servers. See [Managing Zones, on page 129](#).

For more information, refer the following links:

- <http://www.pacnog.org/pacnog6/IXP/Anycast-v10.pdf>
- <http://www.nongnu.org/Quagga>
- <https://frrouting.org/>
- <https://cumulusnetworks.com/learn/frrouting/>
- <https://bgpgeek.com/installing-frr/>
- <https://access.redhat.com/solutions/4967711>
- <https://access.redhat.com/solutions/4538371>
- <http://www.linuxjournal.com/magazine/ipv4-anycast-linux-and-Quagga>
- <http://ddiguru.com/blog/125-anycast-dns-part-5-using-bgp>



Note

The above links reference external websites and Cisco is not responsible for keeping them up-to-date. They are provided for reference only. If you find that the content is outdated or if you cannot access the links, please contact the website owner for updated information.



APPENDIX **C**

DNS Security and Attack Prevention

A DNS attack is any attack targeting the availability or stability of a network's DNS service. There are many different ways in which the DNS can be attacked, such as DNS cache poisoning, DDoS, DNS spoofing, and so on. This chapter explains the features available in Cisco Prime Network Registrar which help in preventing the DNS security related threats and attacks.

- [Prevention of DNS Attacks in Cisco Prime Network Registrar, on page 203](#)

Prevention of DNS Attacks in Cisco Prime Network Registrar

Following features in Cisco Prime Network Registrar help to prevent the DNS security related threats and attacks:

Cache Poisoning

A cache poisoning attack can change an existing entry in the DNS cache as well as insert a new invalid record into the DNS cache. This attack causes a hostname to point to the wrong IP address. For more information on handling cache poisoning attacks, see [Detecting and Preventing DNS Cache Poisoning, on page 43](#).

- **Dynamic allocation of UDP ports**

The Caching DNS server uses a large number of UDP port numbers. The large number of port numbers reduce the risk of cache poisoning via Birthday Attacks. For more information, see [Dynamic Allocation of UDP Ports, on page 40](#).

- **Randomization of DNS transaction ID and source port**

The DNS transaction ID and source port number used to validate DNS responses are not sufficiently randomized and can easily be predicted, which allows an attacker to create forged responses to DNS queries. The DNS server will consider such responses as valid.

- **Randomized query names**

Domain randomization allows a DNS server to send upstream queries for resolution with a randomly generated query name. A valid name server responds with the query name unchanged and therefore this technique can be used to ensure that the response was valid.

Cisco Prime Network Registrar supports randomizing upstream queries, but there are some name servers that do not maintain the randomized case. Therefore, if you enable case randomization, you may block out valid name servers. The *randomize-query-case-exclusion* attribute allows you to create an exclusion list, so that you can continue to use case randomization, but exclude name servers that do not maintain

the case but still respond with a valid answer. For more information, see [Specifying Resolver Settings, on page 40](#).

DDoS Attacks

In a distributed denial-of-service attack (DDoS attack), the incoming traffic flooding the targeted server, service, or network originates from many different sources. This effectively makes it impossible to stop the attack simply by blocking a single source.

- **Rate limiting**

Rate limiting helps the DNS server from being overwhelmed by a small number of clients. It also protects against upstream query attacks against Authoritative DNS servers. This feature helps to mitigate some of the DDoS attacks and prevents the server from being overwhelmed by a small number of clients. It allows you to limit the malevolent traffic. For more information, see [Managing Caching Rate Limiting, on page 53](#).

- **Smart cache**

Whenever Authoritative DNS servers face an outage or are offline for other reasons, this could cause issues with being able to reach Internet services that are likely not impacted. Smart caching allows the Caching DNS server to continue to serve the expired data (last known answer) when it cannot reach the authoritative name servers. The Caching DNS server will still continue to contact the authoritative name servers and when the name servers are once again functional, the Caching DNS server will update its expired data. Smart Caching is useful to mitigate network outages and possible DDoS attacks that make the authoritative name servers unavailable. For more information, see [Enabling Smart Caching, on page 38](#).

- **DNS amplification attack prevention**

A DNS amplification attack is a popular form of DDoS attack that relies on the use of publicly accessible open DNS servers to flood a target system with DNS response traffic. The primary technique consists of an attacker sending a DNS name lookup request to an open DNS server with the source address spoofed to be the target's address. When the DNS server sends the DNS record response, it is sent instead to the target. Attackers typically submit a request for as much zone information as possible to maximize the amplification effect. In most attacks of this type, the spoofed queries sent by the attacker are of the type, "ANY," which returns all known information about a DNS zone in a single request. Because the size of the response is considerably larger than the request, the attacker is able to increase the amount of traffic directed at the target.

- **Allow ANY Query ACL**

In Cisco Prime Network Registrar, the *allow-any-query-acl* attribute on the Manage Servers page helps in minimizing the size of the response. This attribute is present in both Authoritative and Caching DNS server pages, and the default value is "none".

- **Minimal responses**

Cisco Prime Network Registrar supports *minimal-responses* in which authority and additional sections are omitted in the response. This reduces the query response size and defers Denial Of Service to some extent. As of Cisco Prime Network Registrar 11.0, *minimal-responses* is enabled on the Caching DNS server by default and is disabled on the Authoritative DNS server by default.

Data Authentication and Authorization

- **DNSSEC**

DNSSEC provides origin authority, data integrity, and authenticated denial of existence. With DNSSEC, the DNS protocol is much less susceptible to certain types of attacks, particularly DNS spoofing attacks. Cisco Prime Network Registrar supports DNSSEC in both Authoritative and Caching DNS servers.

For more information on DNSSEC support in the Authoritative DNS server, see [Managing Authoritative DNSSEC, on page 99](#).

For more information on DNSSEC support in the Caching DNS server, see [Managing DNSSEC, on page 52](#).

- **DNS firewall**

Caching DNS firewall controls the domain names, IP addresses, and name servers that are allowed to function on the network. The DNS firewall rules can also be set up for specially designated zones on the Authoritative DNS server using RPZ. The RPZ and RR data combined with DNS resolver effectively creates a DNS firewall to prevent misuse of the DNS server. For more information, see [Managing DNS Firewall, on page 115](#).

- **Cisco Umbrella**

Cisco Umbrella provides the first line of defense against threats on the Internet, such as phishing and malware. By setting up the Caching DNS to use Umbrella for resolution, you can allow the Cisco cloud service of Umbrella to provide the latest responses for the requested domain/host. For more information, see [Configuring Caching DNS to Use Umbrella, on page 57](#).

- **Secure DNS server activity with ACLs**

You can restrict clients to query only certain zones based on an ACL.

- **Restricting Zone Queries**—The Authoritative DNS server attribute *restrict-query-acl* limits device queries that the server must honor. The Caching DNS server attributes *acl-query* and *acl-do-not-query* specify IP addresses or subnets that are queried and not queried respectively.
- **Restricting Zone Transfer Requests**—The *restrict-xfer-acl* attribute filters the zone transfer request to the known secondary servers.
- **Restricting DDNS Updates**—The *update-acl* attribute filters DDNS packet from the known DHCP servers.
- **Blocking Malicious Client**—The *acl-blocklist* attribute blocks requests from clients listed in this access control list. This list can contain hosts, network addresses, and/or other ACLs. Request from clients matching this ACL will be dropped.

- **Secure zone transfers and DNS updates using TSIG or GSS-TSIG**

Zone transfer in secure mode supports both HMAC-MD5 based TSIG and GSS-TSIG. You can add an optional TSIG key or GSS-TSIG keys (see the *"Transaction Security"* or *"GSS-TSIG"* sections in *Cisco Prime Network Registrar 11.0 DHCP User Guide*) to the primary server address by hyphenating the entry in the format *address-key*.

- **Secure queries with DoT**

DNS over TLS (DoT) is a security protocol for encrypting and wrapping DNS queries and answers via the TLS protocol. It improves privacy and security between clients and resolvers. It uses TCP as the basic connection protocol and layers over TLS encryption and authentication.

For more information on TLS settings in the Authoritative DNS server, see the [Specifying TLS Settings](#) section in the "Managing Authoritative DNS Server" chapter.

For more information on TLS settings in the Caching DNS server, see the [Specifying TLS Settings](#) section in the "Managing Caching DNS Server" chapter.



INDEX

- A**
- A records [165, 183](#)
 - adding [165](#)
 - resource records [183](#)
 - A6 records [183](#)
 - AAAA records [183](#)
 - About EDNS0 [6](#)
 - absolute domain names [164](#)
 - AFSDB records [183](#)
 - area chart [13](#)
 - Authoritative DNS [67](#)
 - packet logging [67](#)
 - authoritative name servers [3](#)
- B**
- BIND files [137](#)
 - format [137](#)
 - zones [137](#)
 - importing [137](#)
- C**
- CAA [171](#)
 - CAA Record [171, 183](#)
 - cache [43](#)
 - cache, flushing [43](#)
 - caching-only servers [4](#)
 - cdns [47](#)
 - add forwarders [47](#)
 - list forwarders [47](#)
 - CDNS [23, 38, 51](#)
 - DNS64 [51](#)
 - packet logging [23](#)
 - smart cache [38](#)
 - cdns command (CLI) [22, 40, 50](#)
 - addException [50](#)
 - addRootHint [40](#)
 - listExceptions [50](#)
 - removeException [50](#)
 - set [40](#)
 - msg-cache-size [40](#)
 - neg-cache-size [40](#)
 - rrset-cache-size [40](#)
 - cdns command (CLI) (*continued*)
 - show [22](#)
 - cdns commands [45](#)
 - CDNS domain redirect [120](#)
 - overview [120](#)
 - cdns64 command (CLI) [51](#)
 - create [51](#)
 - enable [51](#)
 - CNAME records [183](#)
 - column chart [13](#)
 - configuring DNS servers [96, 130](#)
 - loopback zones [130](#)
 - NOTIFY [96](#)
 - enabling [96](#)
- D**
- dashboard [59–62, 179–181](#)
 - Caching DNS General Indicators chart [59](#)
 - DNS [179–181](#)
 - dashboard [179–181](#)
 - general indicators chart [179](#)
 - inbound zone transfers chart [180](#)
 - network errors chart [180](#)
 - outbound zone transfers chart [181](#)
 - related servers errors chart [181](#)
 - DNS Caching Activity chart [60](#)
 - DNS Caching Server Recursion Rate Limit chart [60](#)
 - DNS Incoming Queries chart [60](#)
 - DNS queries responses chart [61](#)
 - DNS Queries Type chart [62](#)
 - DNS Recursive Query Time chart [62](#)
 - DHCID records [183](#)
 - distributions [147](#)
 - zone [147](#)
 - DNS [2–4, 37, 39–40, 43, 45–46, 49, 65, 92, 98–99, 106, 109, 142, 144, 146, 179–181, 183](#)
 - address format [2](#)
 - cache, flushing [43](#)
 - caching-only servers, See caching-only servers [4](#)
 - dashboard [179–181](#)
 - general indicators chart [179](#)
 - inbound zone transfers chart [180](#)
 - network errors chart [180](#)
 - outbound zone transfers chart [181](#)

DNS (*continued*)

- dashboard (*continued*)
 - related servers errors chart [181](#)
- dns command (CLI) [109](#)
 - set [109](#)
 - log-settings [109](#)
- domain names [2](#)
 - space [2](#)
- exception handling [49](#)
- external ports [106](#)
- flushing DNS cache [43](#)
- glue records [144, 146](#)
 - invalid glue records [144](#)
 - removing [146](#)
- localhost [109](#)
- maximum [37](#)
 - cache TTL property [37](#)
- name-to-address resolution [183](#)
- options [40, 106](#)
 - maximum memory cache size [40](#)
- ports [106](#)
- root name servers [39](#)
- secondary servers, See secondary name servers [3](#)
- servers [45–46, 65, 98–99](#)
 - commands [45, 98](#)
 - network interfaces, configuring [46, 99](#)
- servers logging [109](#)
- Top Names [92](#)
- troubleshooting [109](#)
- zones [142](#)

dns command (CLI) [43, 48, 66, 95–97, 103, 106, 109, 141, 168](#)

- addForwarder [48](#)
- enable [43, 95–97](#)
 - ixfr-enable [96](#)
 - notify [97](#)
 - round-robin [43, 95](#)
- findRR [168](#)
- get [43, 95](#)
 - round-robin [43, 95](#)
- getZoneCount [141](#)
- listForwarders [48](#)
- removeForwarder [48](#)
- set [106, 109](#)
 - activity-summary-interval [109](#)
 - log-settings [109](#)
 - mem-cache-size [106](#)
 - notify-min-interval [106](#)
 - notify-send-stagger [106](#)
 - notify-wait [106](#)
- show [66, 103](#)

dns command CLI) [127](#)

- getStats [127](#)
- ha [127](#)

DNS commands [98](#)

DNS ENUM [151](#)

- manage defaults [151](#)

DNS ENUM (*continued*)

- overview [151](#)

DNS ENUM command (CLI) [151](#)

- add [151](#)
- remove [151](#)
- set default [151](#)

DNS ENUM domain [151, 153–154](#)

- add [151](#)
- pull [154](#)
- push [153](#)

DNS ENUM domain command (CLI) [152](#)

- create [152](#)
- delete [152](#)

DNS ENUM number [152, 154–155](#)

- add [152](#)
- pull [155](#)
- push [154](#)

DNS ENUM number command (CLI) [153](#)

- add [153](#)

DNS views [157–161](#)

- configure [157](#)
- key points [158](#)
- manage [159](#)
- pull [161](#)
- push [160](#)
- reorder [160](#)
- synchronize [160](#)

DNS views command (CLI) [160](#)

- reorder [160](#)

DNS64 [51](#)

- Managing [51](#)

DNSSEC [52](#)

- Managing [52](#)

domain names [2](#)

- space [2](#)
- tree structure [2](#)

domains [3](#)

- registering [3](#)

E

exception [49](#)

- See resolution exception [49](#)

F

forwarding DNS servers [47–48](#)

- listing [48](#)

H

HA DNS [125–126](#)

- backup server, setting [125](#)
- dns command (CLI) [126](#)
- setPartnerDown [126](#)

- HA DNS (*continued*)
 - enabling [125](#)
 - ha-dns-pair command (CLI) [126](#)
 - create [126](#)
 - main server, setting [125](#)
 - synchronizing server pairs [126](#)
 - ha-dns-pair command (CLI) [125–126](#)
 - set [125](#)
 - ha-dns-backup-server [125](#)
 - ha-dns-main-server [125](#)
 - sync [126](#)
 - Handling Malicious DNS Clients and Unresponsive Nameservers [106](#)
 - Host Info records [183](#)
 - See HINFO records [183](#)
 - hosts [142, 175–176](#)
 - adding to zones [175](#)
 - dynamic [142](#)
 - editing [176](#)
- I**
- import command (CLI) [137](#)
 - in-addr.arpa domain [5](#)
 - incremental zone transfers [96](#)
 - enabling [96](#)
 - IP addresses [1](#)
 - See addresses, IP [1](#)
 - ISDN records [183](#)
- L**
- LDAP [170](#)
 - line chart [13](#)
 - localhost [130](#)
 - logging [109](#)
 - NOTIFY [109](#)
 - logging transactions [109](#)
- M**
- MB records [183](#)
 - MG records [183](#)
 - MINFO records [183](#)
 - MR records [183](#)
 - MX records [183](#)
- N**
- name server [4](#)
 - DNS client/server model [4](#)
 - DNS primary servers, See primary name servers [4](#)
 - name-to-address resolution [4](#)
 - Name Server records [183](#)
 - See NS records [183](#)
 - name servers [3–4](#)
 - domain [4](#)
 - primary, See primary name servers [3](#)
 - secondary [3](#)
 - name servers, DNS [3](#)
 - types [4](#)
 - name-to-address resolution [183](#)
 - Naming Authority Pointer records [183](#)
 - See NAPTR records [183](#)
 - NAPTR records [170, 183](#)
 - network interfaces [46, 99](#)
 - DNS server [46, 99](#)
 - network number [3](#)
 - NOTIFY [109](#)
 - logging transactions [109](#)
 - NS records [183](#)
 - NSAP records [183](#)
 - nslookup utility [109](#)
- P**
- packet logging [23, 67](#)
 - Pointer (Reverse Mapping) record [183](#)
 - See PTR records [183](#)
 - primary name servers [3, 130, 133, 183](#)
 - configuring [130](#)
 - SOA records [183](#)
 - zones [133](#)
 - primary server, setting [133](#)
- R**
- relative domain names [165](#)
 - remote-dns command (CLI) [106](#)
 - create [106](#)
 - resolution exception [49](#)
 - resource records [146, 163, 165–166, 183](#)
 - A [183](#)
 - A6 [183](#)
 - AAAA [183](#)
 - adding in CLI [165](#)
 - AFSDB [183](#)
 - CAA [183](#)
 - CNAME [183](#)
 - configuring [163](#)
 - DHCID [183](#)
 - editing [146, 163](#)
 - subzone information [146](#)
 - HINFO [183](#)
 - ISDN [183](#)
 - MB [183](#)
 - MG [183](#)
 - MINFO [183](#)
 - MR [183](#)
 - MX [183](#)

resource records (*continued*)

- NAPTR [183](#)
- NS [183](#)
- NSAP [183](#)
- protecting [166](#)
- PTR [183](#)
- RP [183](#)
- RT [183](#)
- SOA [183](#)
- SRV [183](#)
- TXT [183](#)
- types [183](#)
- URI [183](#)
- WKS [183](#)

Resource Records [171–172](#)

- CAA [171](#)
- URI [172](#)

reverse [5, 139, 183](#)

- domains [5](#)
- mapping records [183](#)
- zones [5, 139](#)
 - configuring [139](#)

RFCs [96, 137, 142, 169–172, 183](#)

- 1035 [137](#)
- 1995 [96](#)
- 1996 [96](#)
- 2136 [142](#)
- 2782 [169, 183](#)
- 2915 [183](#)
- 2916 [170](#)
- 3263 [170](#)
- 3403 [170](#)
- 4701 [183](#)
- 6844 [171](#)
- 7553 [172](#)

round-robin [94](#)

- enabling [94](#)

RP records [183](#)

RT records [183](#)

S

scatter chart [13](#)

secondary [4, 105, 142–143](#)

- DNS [142](#)
 - zones [142](#)
- expiration time [105](#)
- name servers, DNS [4](#)
 - defined [4](#)
- refresh time [143](#)
- retry time [105](#)
- SOA records [105](#)
- time zones [105](#)
- zones [142](#)
- Zones [105](#)

session command (CLI) [133](#)

- set [133](#)
 - dns-edit-mode [133](#)

Session Initiation Protocol (SIP) proxies [170](#)

smart cache [38](#)

SOA records [104, 133, 183](#)

- defined [133](#)
- TTL property [104](#)
 - zone [104](#)

SRV records [183](#)

staged and synchronous modes [132](#)

subzones [144–146](#)

- adding [144](#)
- delegating [145](#)
- name server [144](#)
- naming [144](#)
- removing [146](#)

synchronous [132](#)

- dns edit mode [132](#)
- edits [132](#)
- staged dns edit mode [132](#)

T

templates [130](#)

- zone [130](#)

Text records [183](#)

- TXT records [183](#)

Time to Live property [104](#)

- See TTL property [104](#)

TLS settings [34, 92](#)

TTL property [37, 40, 104–105, 183](#)

- default [104–105](#)
 - responses [104](#)
- DNS [37](#)
 - maximum [37](#)
 - cache TTL property [37](#)
 - maximum DNS options [40](#)

TXT records [183](#)

U

URI [172](#)

URI Record [172, 183](#)

W

Well Known Services record [183](#)

- See WKS records [183](#)

WKS records [183](#)

Z

zone [132–133](#)

- edit mode, setting [133](#)

- zone (*continued*)
 - templates [132](#)
 - cloning [132](#)
 - zone-template command (CLI) [132](#)
 - create [132](#)
- zone command (CLI) [97, 105–106, 133, 135–136, 144, 165, 169, 176–177](#)
 - addDNSRR [165, 176](#)
 - addHost [176](#)
 - addRR [135, 165, 176](#)
 - staged or -sync [165](#)
 - A [135](#)
 - applyTemplate [135](#)
 - create [135](#)
 - primary [135](#)
 - template, using [135](#)
 - enable [97](#)
 - notify [97](#)
 - findRR [169](#)
 - forceXfer [144](#)
 - get [133, 135](#)
 - serial [133, 135](#)
 - listHosts [135, 176](#)
 - listRR [136](#)
 - removeDNSRR [165, 177](#)
 - removeHost [177](#)
 - removeRR [165, 177](#)
 - restrict-xfer [144](#)
 - set [97, 105–106, 135](#)
 - defttl [105](#)
 - expire [106](#)
 - nameservers [135](#)
 - notify-set [97](#)
 - refresh [105](#)
 - retry [105](#)
 - show [135](#)
- zone distributions [147, 150](#)
 - managing [147](#)
 - synchronizing [150](#)
 - zone-dist command (CLI) [150](#)
 - addSecondary [150](#)
 - create [150](#)
- zone templates [130, 132, 150](#)
 - applying to zones [132](#)
 - creating [130](#)
 - zone distributions, associating [150](#)
- zone transfers [143](#)
 - enabling [143](#)
- zone transfers (*continued*)
 - forcing [143](#)
 - forcing all [143](#)
- zone-dist command (CLI) [150](#)
 - sync [150](#)
- zone-template command (CLI) [132, 150](#)
 - apply-to [132](#)
 - create [132](#)
 - clone [132](#)
 - set [150](#)
 - dist-map [150](#)
- zones [3–5, 104, 133, 135, 137, 141–143, 146, 150, 164–166, 169, 175–177](#)
 - adding [164](#)
 - authoritative name servers [133](#)
 - adding [133](#)
 - defined [3](#)
 - dns command (CLI) [141](#)
 - getZoneCount [141](#)
 - DNS update [142](#)
 - domains [3, 164](#)
 - difference from zones [3](#)
 - enabling DNS update [142](#)
 - host tables, editing [176](#)
 - hosts [175, 177](#)
 - removing [177](#)
 - importing [137](#)
 - names, creating [133](#)
 - point of delegation [3](#)
 - removing [146](#)
 - removing resource records [165](#)
 - resource records [164–166, 169](#)
 - filtering [169](#)
 - protecting [166](#)
 - reverse, See reverse zones [5](#)
 - serial numbers [133](#)
 - subzones [146](#)
 - editing [146](#)
 - removing [146](#)
 - template, adding from [135](#)
 - transfers, See zone transfers [4](#)
 - TTL property, setting [104](#)
 - zone command (CLI) [150](#)
 - set [150](#)
 - dist-map [150](#)
 - zone transfers, defined [4](#)
 - zone transfers, enabling [143](#)

