



Cisco Prime Network Registrar 9.0 Administration Guide

First Published: 2016-12-22

Last Modified: 2016-12-22

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.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

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: [www.cisco.com go trademarks](http://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)

© 2016 Cisco Systems, Inc. All rights reserved.



CONTENTS

PART I

Getting Started 15

CHAPTER 1

Introduction to Cisco Prime Network Registrar 1

- Target Users 1
- Regional and Local Clusters 1
- Deployment Scenarios 2
 - Related Topics 2
 - Small-to-Medium-Size LANs 3
 - Large Enterprise and Service Provider Networks 3
- Configuration and Performance Guidelines 5
 - Related Topics 5
 - General Configuration Guidelines 5
 - Special Configuration Cases 6
 - General Performance Guidelines 6
 - Interoperability with Earlier Releases 6

CHAPTER 2

Cisco Prime Network Registrar User Interfaces 9

- Management Components 9
- Introduction to the Web-Based User Interfaces 10
 - Related Topics 10
 - Supported Web Browsers 11
 - Access Security 11
 - Logging In to the Web UIs 11
 - Multiple Users 12
 - Changing Passwords 12
 - Navigating the Web UIs 13

Waiting for Page Resolution Before Proceeding	13
Committing Changes in the Web UIs	14
Role and Attribute Visibility Settings	14
Displaying and Modifying Attributes	14
Grouping and Sorting Attributes	14
Modifying Attributes	14
Displaying Attribute Help	15
Left Navigation Pane	15
Help Pages	15
Logging Out	16
Local Cluster Web UI	16
Related Topics	16
Local Basic Main Menu Page	16
Local Advanced Main Menu Page	17
Setting Local User Preferences	18
Configuring Clusters in the Local Web UI	19
Regional Cluster Web UI	20
Related Topics	20
Command Line Interface	20
Global Search in Prime Network Registrar	21

CHAPTER 3

Server Status Dashboard	23
Opening the Dashboard	23
Display Types	24
General Status Indicators	24
Graphic Indicators for Levels of Alert	25
Magnifying and Converting Charts	25
Legends	25
Tables	25
Line Charts	25
Stacked Area Charts	27
Other Chart Types	27
Getting Help for the Dashboard Elements	28
Customizing the Display	28

Refreshing Displays	28
Setting the Polling Interval	28
Displaying Charts as Tables	29
Exporting to CSV Format	29
Displaying or Hiding Chart Legends	29
Selecting Dashboard Elements to Include	30
Configuring Server Chart Types	30
Host Metrics	32
System Metrics	32
JVM Memory Utilization	33
<hr/>	
PART II	Local and Regional Administration
<hr/>	
CHAPTER 4	Managing Administrators
Administrators, Groups, Roles, and Tenants	37
Related Topics	37
How Administrators Relate to Groups, Roles, and Tenants	38
Administrator Types	38
Roles, Subroles, and Constraints	39
Groups	42
External Authentication Servers	42
Configuring an RADIUS External Authentication Server	43
Configuring an AD External Authentication Server	44
Managing Tenants	46
Adding a Tenant	47
Editing a Tenant	47
Managing Tenant Data	48
Assigning a Local Cluster to a Single Tenant	49
Pushing and Pulling Tenant Data	49
Assigning Tenants When Using External Authentication	50
Using cnr_exim With Tenant Data	50
Managing Administrators	51
Adding Administrators	51
Editing Administrators	52

Deleting Administrators	52
Managing Passwords	52
Managing Groups	52
Adding Groups	52
Editing Groups	53
Deleting Groups	53
Managing Roles	53
Adding Roles	53
Editing Roles	54
Deleting Roles	54
CLI Commands	54
Granular Administration	54
Local Advanced and Regional Web UI	54
Related Topics	55
Scope-Level Constraints	55
Prefix-Level Constraints	56
Link-Level Constraints	57
Centrally Managing Administrators	58
Related Topics	58
Pushing and Pulling Administrators	59
Pushing Administrators to Local Clusters	59
Pushing Administrators Automatically to Local Clusters	59
Pulling Administrators from the Replica Database	60
Pushing and Pulling External Authentication Servers	61
Pushing and Pulling Groups	63
Pushing Groups to Local Clusters	63
Pulling Groups from the Replica Database	64
Pushing and Pulling Roles	64
Pushing Roles to Local Clusters	64
Pulling Roles from the Replica Database	65
Pushing and Pulling Tenants	66
Pushing Tenants to Local Clusters	66
Pulling Tenants from the Replica Database	66

CHAPTER 5**Managing Owners and Regions 69**

Managing Owners 69

Local Advanced and Regional Advanced Web UI 69

CLI Commands 69

Managing Regions 70

Local Advanced and Regional Advanced Web UI 70

CLI Commands 70

Centrally Managing Owners and Regions 70

Related Topics 70

Pushing and Pulling Owners or Regions 71

Pushing Owners or Regions to Local Clusters 71

Pulling Owners and Regions from the Replica Database 71

CHAPTER 6**Managing the Central Configuration 73**

Central Configuration Tasks 73

Default Ports for Cisco Prime Network Registrar Services 74

Firewall Considerations 75

Licensing 75

Regional Web UI 75

Adding License 76

CLI Commands 76

Registering a Local Cluster that is Behind a NAT 77

CLI Commands 77

License History 78

Configuring Server Clusters 78

Related Topics 78

Adding Local Clusters 79

Editing Local Clusters 80

Connecting to Local Clusters 80

Synchronizing with Local Clusters 81

Replicating Local Cluster Data 81

Viewing Replica Data 82

Purging Replica Data 82

Deactivating, Reactivating, and Recovering Data for Clusters	82
Central Configuration Management Server	84
Managing CCM Server	84
Editing CCM Server Properties	84
Trivial File Transfer	84
Related Topics	85
Viewing and Editing the TFTP Server	85
Managing the TFTP Server Network Interfaces	85
Simple Network Management	86
Related Topics	87
Setting Up the SNMP Server	87
How Notification Works	88
Handling SNMP Notification Events	91
Handling Deactivated Scopes or Prefixes	93
Editing Trap Configuration	93
Deleting Trap Configuration	93
Server Up/Down Traps	94
Handling SNMP Queries	94
Integrating Cisco Prime Network Registrar SNMP into System SNMP	95
Bring Your Own Device Web Server	96
Managing BYOD Web Server	96
Editing BYOD Web Server Properties	96
Setting Up BYOD Theme and Content	97
Adding and Previewing BYOD Themes	97
Adding and Previewing BYOD Content	97
Polling Process	98
Polling Subnet Utilization and Lease History Data	98
Adjusting the Polling Intervals	99
Enabling Subnet Utilization Collection	99
Enabling Lease History Collection	100
Managing DHCP Scope Templates	100
Related Topics	101
Pushing Scope Templates to Local Clusters	101
Pulling Scope Templates from Replica Data	101

Managing DHCP Policies	102
Related Topics	102
Pushing Policies to Local Clusters	102
Pulling Policies from Replica Data	103
Managing DHCP Client-Classes	103
Related Topics	103
Pushing Client-Classes to Local Clusters	104
Pulling Client-Classes from Replica Data	104
Managing Virtual Private Networks	105
Related Topics	105
Pushing VPNs to Local Clusters	105
Pulling VPNs from Replica Data	105
Managing DHCP Failover Pairs	106
Regional Web UI	106
Managing Lease Reservations	106
Related Topics	107
DHCPv4 Reservations	107
DHCPv6 Reservations	107
Monitoring Resource Limit Alarms	108
Configuring Resource Limit Alarm Thresholds	109
Setting Resource Limit Alarms Polling Interval	109
Viewing Resource Limit Alarms	110
Local Cluster Management Tutorial	111
Related Topics	111
Administrator Responsibilities and Tasks	111
Create the Administrators	112
Create the Address Infrastructure	112
Create the Zone Infrastructure	113
Create the Forward Zones	113
Create the Reverse Zones	114
Create the Initial Hosts	114
Create a Host Administrator Role with Constraints	115
Create a Group to Assign to the Host Administrator	116
Test the Host Address Range	117

Regional Cluster Management Tutorial	117
Related Topics	118
Administrator Responsibilities and Tasks	118
Create the Regional Cluster Administrator	118
Create the Central Configuration Administrator	119
Create the Local Clusters	119
Add a Router and Modify an Interface	120
Add Zone Management to the Configuration Administrator	121
Create a Zone for the Local Cluster	121
Pull Zone Data and Create a Zone Distribution	122
Create a Subnet and Pull Address Space	122
Push a DHCP Policy	123
Create a Scope Template	123
Create and Synchronize the Failover Pair	124

CHAPTER 7**Managing Routers and Router Interfaces 125**

Adding Routers	125
Local Advanced and Regional Web UI	125
CLI Commands	125
Editing Routers	126
Local Advanced and Regional Web UI	126
CLI Commands	126
Viewing and Editing the Router Interfaces	126
Local Advanced and Regional Web UI	126
CLI Commands	126
Related Topics	126
Changeable Router Interface Attributes	126
Bundling Interfaces	127
Pushing and Reclaiming Subnets for Routers	127

CHAPTER 8**Maintaining Servers and Databases 129**

Managing Servers	129
Local Basic or Advanced and Regional Web UI	130
CLI Commands	131

Scheduling Recurring Tasks	131
Local Basic or Advanced Web UI	132
Logs	132
Log Files	132
Logging Server Events	134
Logging Format and Settings	134
Searching the Logs	135
View Change Log	135
Dynamic Update on Server Log Settings	136
Running Data Consistency Rules	137
Local Basic or Advanced and Regional Web UI	137
CLI Tool	138
Monitoring and Reporting Server Status	140
Related Topics	140
Server States	141
Displaying Health	141
Server Health Status	142
Displaying Statistics	142
DNS Statistics	143
CDNS Statistics	146
DHCP Statistics	149
TFTP Statistics	151
Displaying IP Address Usage	153
Displaying Related Servers	153
Monitoring Remote Servers Using Persistent Events	153
DNS Zone Distribution Servers	155
DHCP Failover Servers	155
Displaying Leases	156
Troubleshooting DHCP and DNS Servers	156
Related Topics	156
Immediate Troubleshooting Actions	156
Modifying the cnr.conf File	157
Troubleshooting Server Failures	159
Linux Troubleshooting Tools	160

Using the TAC Tool	160
Troubleshooting and Optimizing the TFTP Server	161
Related Topics	161
Tracing TFTP Server Activity	161
Optimizing TFTP Message Logging	161
Enabling TFTP File Caching	162
<hr/>	
CHAPTER 9	Backup and Recovery 163
Backing Up Databases	163
Related Topics	163
Syntax and Location	164
Backup Strategy	164
Manual Backup (Using <code>cnr_shadow_backup</code> utility)	164
Setting Automatic Backup Time	165
Performing Manual Backups	165
Using Third-Party Backup Programs with <code>cnr_shadow_backup</code>	165
Backing Up CNRDB Data	166
Backing Up All CNRDBs Using <code>tar</code> or Similar Tools	166
Database Recovery Strategy	167
Recovering CNRDB Data from Backups	168
Recovering All CNRDBs Using <code>tar</code> or Similar Tools	169
Recovering Single CNRDB from <code>tar</code> or Similar Tools	169
Virus Scanning While Running Cisco Prime Network Registrar	170
Troubleshooting Databases	170
Related Topics	170
Using the <code>cnr_exim</code> Data Import and Export Tool	170
Using the <code>cnrdb_recover</code> Utility	172
Using the <code>cnrdb_verify</code> Utility	173
Using the <code>cnrdb_checkpoint</code> Utility	174
Using the <code>cnrdb_util</code> Utility	174
Restoring DHCP Data from a Failover Server	177
<hr/>	
CHAPTER 10	Managing Reports 179
ARIN Reports and Allocation Reports	179

Managing ARIN Reports	179
Related Topics	180
Managing Point of Contact and Organization Reports	180
Creating a Point of Contact Report	181
Registering a Point of Contact	181
Editing a Point of Contact Report	181
Creating an Organization Report	182
Registering an Organization	182
Editing an Organization Report	183
Managing IPv4 Address Space Utilization Reports	183
Regional Web UI	184
Managing Shared WHOIS Project Allocation and Assignment Reports	184
Managing BYOD Reports	184
Registered Devices	185
Registered Devices Report	185
Scopes/Prefix	185
Scope/Prefix Report	185
<hr/>	
PART III	Virtual Appliance 187
<hr/>	
CHAPTER 11	Introduction to Cisco Prime Network Registrar Virtual Appliance 189
	How the Cisco Prime Network Registrar Virtual Appliance Works 189
	Invoking Cisco Prime Network Registrar on the Virtual Appliance 190
	Monitoring Disk Space Availability on VMware 190
	Monitoring Disk Space Availability in Use by the Virtual Appliance 190
	Increasing the Size of the Disk on VMware 190
	Increasing the Size of the Disk on a KVM Hypervisor 191
	Troubleshooting 192
	 Glossary 193



PART I

Getting Started

- [Introduction to Cisco Prime Network Registrar, on page 1](#)
- [Cisco Prime Network Registrar User Interfaces, on page 9](#)
- [Server Status Dashboard, on page 23](#)



CHAPTER 1

Introduction to Cisco Prime Network Registrar

Cisco Prime Network Registrar is a full featured, scalable Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and Trivial File Transfer Protocol (TFTP) implementation for medium to large IP networks. It provides the key benefits of stabilizing the IP infrastructure and automating networking services, such as configuring clients and provisioning cable modems. This provides a foundation for policy-based networking.

Service provider and enterprise users can better manage their networks to integrate with other network infrastructure software and business applications.

- [Target Users, on page 1](#)
- [Regional and Local Clusters, on page 1](#)
- [Deployment Scenarios, on page 2](#)
- [Configuration and Performance Guidelines, on page 5](#)

Target Users

Cisco Prime Network Registrar is designed for these users:

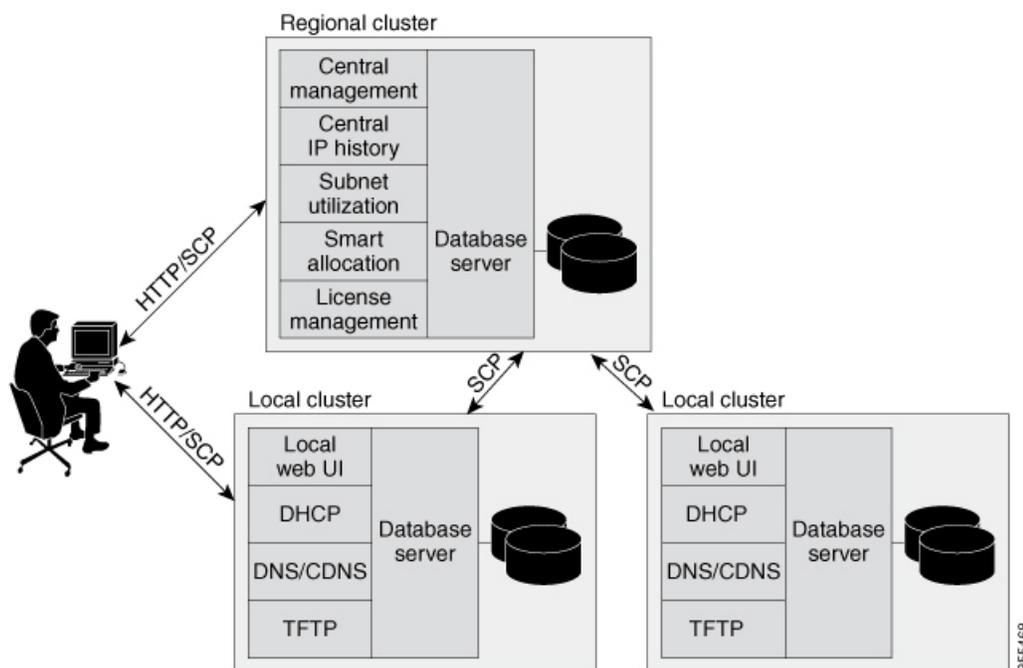
- **Internet service providers (ISPs)**—Helps ISPs drive the cost of operating networks that provide leased line, dialup, and DSL (Point-to-Point over Ethernet and DHCP) access to customers.
- **Multiple service operators (MSOs)**—Helps MSOs provide subscribers with Internet access using cable or wireless technologies. MSOs can benefit from services and tools providing reliable and manageable DHCP and DNS services that meet the Data Over Cable Service Interface Specification (DOCSIS). Cisco Prime Network Registrar provides policy-based, robust, and scalable DNS and DHCP services that form the basis for a complete cable modem provisioning system.
- **Enterprises**—Helps meet the needs of single- and multisite enterprises (small-to-large businesses) to administer and control network functions. Cisco Prime Network Registrar automates the tasks of assigning IP addresses and configuring the Transport Control Protocol/Internet Protocol (TCP/IP) software for individual network devices. Forward-looking enterprise users can benefit from class-of-service and other features that help integrate with new or existing network management applications, such as user registration.

Regional and Local Clusters

The regional cluster acts as an aggregate management system for up to a hundred local clusters. Address and server administrators interact at the regional and local clusters through the regional and local web-based user

interfaces (web UIs), and local cluster administrators can continue to use the command line interface (CLI) at the local cluster. The regional cluster consists of a Central Configuration Management (CCM) server, Tomcat web server, servlet engine, and server agent (see [Management Components, on page 9](#)). The license management is now done at the regional cluster and hence the local server has to be registered to a regional server to avail the necessary services. See the "Overview" chapter in *Cisco Prime Network Registrar Installation Guide* for more details.

Figure 1: Cisco Prime Network Registrar User Interfaces and Server Clusters



A typical deployment is one regional cluster at a customer network operation center (NOC), the central point of network operations for an organization. Each division of the organization includes a local address management server cluster responsible for managing a part of the network. The System Configuration Protocol (SCP) communicates the configuration changes between the servers.

Deployment Scenarios

The Cisco Prime Network Registrar regional cluster web UI provides a single point to manage any number of local clusters hosting DNS, CDNS, DHCP, or TFTP servers. The regional and local clusters also provide administrator management so that you can assign administrative roles to users logged in to the application.

This section describes two basic administrative scenarios and the hardware and software deployments for two different types of installations—a small-to-medium local area network (LAN), and a large-enterprise or service-provider network with three geographic locations.

Related Topics

[Small-to-Medium-Size LANs, on page 3](#)

[Large Enterprise and Service Provider Networks, on page 3](#)

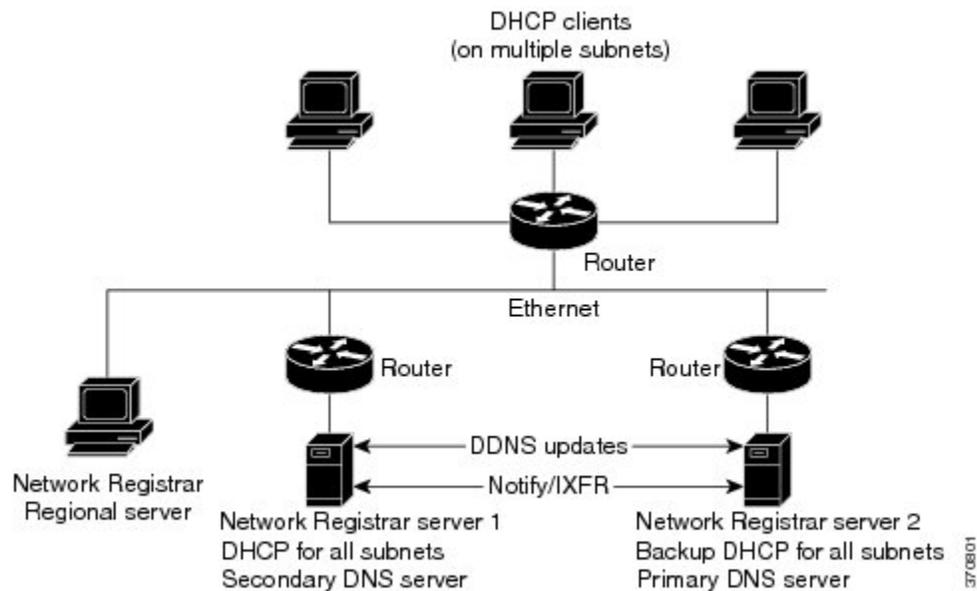
Small-to-Medium-Size LANs

In this scenario, low-end Windows or Linux servers are acceptable. The image below shows a configuration that would be adequate for this network.



Note Regional server is MUST in deployment for small and medium sized LANs.

Figure 2: Small-to-Medium LAN Configuration

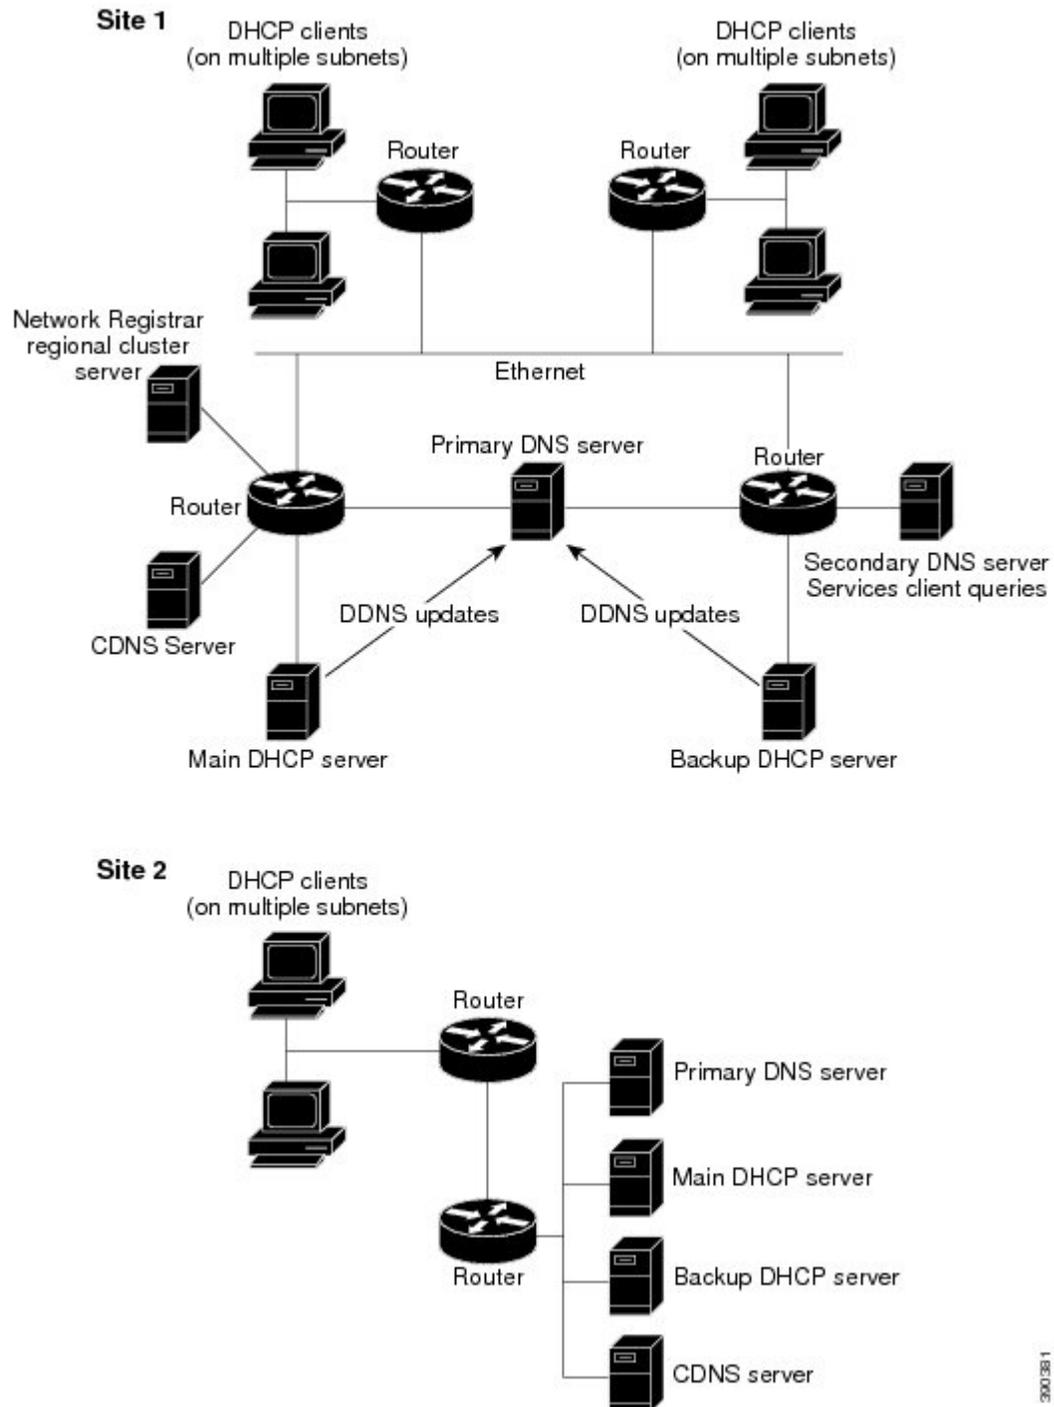


Large Enterprise and Service Provider Networks

In a large enterprise or service provider network serving over 500,000 DHCP clients, use mid-range Windows or Linux servers. Put DNS and DHCP servers on different systems. The image below shows the hardware that would be adequate for this network.

When supporting geographically dispersed clients, locate DHCP servers at remote locations to avoid disrupting local services if wide-area connections fail. Install the Cisco Prime Network Registrar regional cluster to centrally manage the distributed clusters.

Figure 3: Large Enterprise or Service Provider Network Configuration



390381

Configuration and Performance Guidelines

Cisco Prime Network Registrar is an integrated DHCP, DNS, and TFTP server cluster capable of running on a Windows or Linux workstation or server.

Because of the wide range of network topologies for which you can deploy Cisco Prime Network Registrar, you should first consider the following guidelines. These guidelines are very general and cover most cases. Specific or challenging implementations could require additional hardware or servers.

Related Topics

[General Configuration Guidelines, on page 5](#)

[Special Configuration Cases, on page 6](#)

[General Performance Guidelines, on page 6](#)

General Configuration Guidelines

The following suggestions apply to most Cisco Prime Network Registrar deployments:

- Configure a separate DHCP server to run in remote segments of the wide area network (WAN).
Ensure that the DHCP client can consistently send a packet to the server in under a second. The DHCP protocol dictates that the client receive a response to a DHCPDISCOVER or DHCPREQUEST packet within four seconds of transmission. Many clients (notably early releases of the Microsoft DHCP stack) actually implement a two-second timeout.
- In large deployments, separate the secondary DHCP server from the primary DNS server used for dynamic DNS updates.
Because lease requests and dynamic DNS updates are persisted to disk, server performance is impacted when using a common disk system. So that the DNS server is not adversely affected, run it on a different cluster than the DHCP server.
- Include a time server in your configuration to deal with time differences between the local and regional clusters so that aggregated data at the regional server appears in a consistent way. See the [Polling Subnet Utilization and Lease History Data, on page 98](#).
- Set DHCP lease times in policies to four to ten days.
To prevent leases from expiring when the DHCP client is turned off (overnight or over long weekends), set the DHCP lease time longer than the longest period of expected downtime, such as seven days. See *"Managing Leases" section in Cisco Prime Network Registrar 9.0 DHCP User Guide*.
- Locate backup DNS servers on separate network segments.
DNS servers are redundant by nature. However, to minimize client impact during a network failure, ensure that primary and secondary DNS servers are on separate network segments.
- If there are high dynamic DNS update rates in the network, configure separate DNS servers for forward and reverse zones.
- Use NOTIFY/IXFR.

Secondary DNS servers can receive their data from the primary DNS server in two ways: through a full zone transfer (AXFR) or an incremental zone transfer (NOTIFY/IXFR, as described in RFCs 1995 and

1996). Use NOTIFY/IXFR in environments where the name space is relatively dynamic. This reduces the number of records transferred from the primary to the secondary server. See the *"Enabling Incremental Zone Transfers (IXFR)"* section in *Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*.

Special Configuration Cases

The following suggestions apply to some special configurations:

- When using dynamic DNS updates for large deployments or very dynamic networks, divide primary and secondary DNS and DHCP servers across multiple clusters.

Dynamic DNS updates generate an additional load on all Cisco Prime Network Registrar servers as new DHCP lease requests trigger dynamic DNS updates to primary servers that update secondary servers through zone transfers.

- During network reconfiguration, set DHCP lease renewal times to a small value.

Do this several days before making changes in network infrastructure (such as to gateway router and DNS server addresses). A renewal time of eight hours ensures that all DHCP clients receive a changed DHCP option parameter within one working day. See the *"Managing Leases"* section in *Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*

General Performance Guidelines

For Cisco Prime Network Registrar, the general guideline is to invest in the highest performance disk I/O subsystem available, then memory, and finally the processors. DHCP and Authoritative DNS (especially if using DNS updates) will be most impacted by disk latency, then memory and network performance, and finally CPU (these applications are not CPU intensive).

- The best way to reduce latency and improve performance is to provide high performance disks (SSD are recommended over traditional hard disks). High performance disk controllers are also recommended. This is especially important for DHCP and Authoritative DNS servers that handle Dynamic Updates.
- Providing lots of memory is also important as it reduces disk read requirements if the file system cache can be used. The recommendation here is to assure that a system has sufficient free memory that is twice the size of the CPNRdatabases. It is difficult to give exact requirements here as it depends on many variables.
- Network performance is also an important consideration and 1 GB or better Ethernet controllers are recommended.
- As most Cisco Prime Network Registrar uses are not CPU intensive, the CPU performance tends to be least important.

Interoperability with Earlier Releases

The following table shows the interoperability of Cisco Prime Network Registrar features on the regional CCM server with versions of the local cluster.

Table 1: CCM Regional Feature Interoperability with Server Versions

Feature	Local Cluster Version			
	8.1	8.2	8.3	9.0
Push and pull:				
Address space	x	x	x	x
IPv6 address space	x	x	x	x
Scope templates, policies, client-classes	x	x	x	x
IPv6 prefix and link templates	x	x	x	x
Zone data and templates	x	x	x	x
Groups, owners, regions	x	x	x	x
Resource records (RRs)	x	x	x	x
Local cluster restoration	x	x	x	x
Host administration	x	x	x	x
Extended host administration	x	x	x	x
Administrators and roles	x	x	x	x
Zone Views		x	x	x
Administrator:				
Single sign-on	x	x	x	x
Password change	x	x	x	x
IP history reporting:				
Lease history	x	x	x	x
Detailed lease history	x	x	x	x
Utilization reporting:				
Subnet utilization history	x	x	x	x
Subnet and scope utilization	x	x	x	x
IPv6 prefix utilization	x	x	x	x



CHAPTER 2

Cisco Prime Network Registrar User Interfaces

Cisco Prime Network Registrar provides a regional and a local web UI and a regional and local CLI to manage the CDNS, DNS, DHCP, TFTP, and CCM servers:

- **Web UI for the regional cluster to access local cluster servers**—See [Regional Cluster Web UI](#), on page 20.
- **Web UI for the local cluster**—See [Local Cluster Web UI](#), on page 16.
- **CLI for the local clusters**—Open the `CLIContent.html` file in the installation `/docs` directory (see [Command Line Interface](#), on page 20).
- **CCM servers that provide the infrastructure to support these interfaces**— See [Central Configuration Management Server](#), on page 84.

This chapter describes the Cisco Prime Network Registrar user interfaces and the services that the CCM servers provide. Read this chapter before starting to configure the Cisco Prime Network Registrar servers so that you become familiar with each user interface capability.

- [Management Components](#), on page 9
- [Introduction to the Web-Based User Interfaces](#), on page 10
- [Local Cluster Web UI](#), on page 16
- [Regional Cluster Web UI](#), on page 20
- [Command Line Interface](#), on page 20
- [Global Search in Prime Network Registrar](#), on page 21

Management Components

Cisco Prime Network Registrar contains two management components:

- Regional component, consisting of:
 - Web UI
 - CLI
 - CCM Server
 - Bring your own device (BYOD)
 - Simple Network Management Protocol (SNMP) server
- Local component, consisting of:

- Web UI
- CLI
- CCM server
- Authoritative Domain Name System (DNS) server
- Caching / Recursive Domain Name System (CDNS) server
- Dynamic Host Configuration Protocol (DHCP) server
- Trivial File Transport Protocol (TFTP) server
- SNMP server
- Management of local address space, zones, scopes, DHCPv6 prefixes and links, and users



Note 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. However, Cisco recommends hybrid mode for smaller sized deployments only. For larger deployments, Cisco recommends separating Caching and Authoritative DNS on separate physical machines or VMs.

License management is done from the regional cluster when Cisco Prime Network Registrar is installed. You must install the regional server first and load all licenses in the regional server. When you install the local cluster, it registers with regional to obtain its license.

The regional CCM server provides central management of local clusters, with an aggregated view of DHCP address space and DNS zones. It provides management of the distributed address space, zones, scopes, DHCPv6 prefixes and links, and users.

The local CCM server provides management of the local address space, zones, scopes, DHCPv6 prefixes and links, and users.

The remainder of this chapter describes the TFTP and SNMP protocols. The CCM server, web UIs, and CLI are described in [Cisco Prime Network Registrar User Interfaces, on page 9](#). The DNS, CDNS, and DHCP servers are described in their respective sections.

Introduction to the Web-Based User Interfaces

The web UI provides granular access to configuration data through user roles and constraints. The UI provides quick access to common functions. The web UI granularity is described in the following sections.

Related Topics

[Supported Web Browsers, on page 11](#)

[Access Security, on page 11](#)

[Logging In to the Web UIs, on page 11](#)

[Multiple Users, on page 12](#)

[Changing Passwords, on page 12](#)

[Navigating the Web UIs, on page 13](#)

[Waiting for Page Resolution Before Proceeding, on page 13](#)

[Committing Changes in the Web UIs, on page 14](#)

[Role and Attribute Visibility Settings, on page 14](#)

[Displaying and Modifying Attributes, on page 14](#)

[Help Pages, on page 15](#)

[Logging Out, on page 16](#)

Supported Web Browsers

The web UI has been tested on Microsoft Internet Explorer 9, Mozilla Firefox 21 and later, and Google Chrome 53. Internet Explorer 8 is not supported.

Access Security

At Cisco Prime Network Registrar installation, you can choose to configure HTTPS to support secure client access to the web UIs. You must specify the HTTPS port number and provide the keystore at that time. With HTTPS security in effect, the web UI Login page indicates that the “Page is SSL¹ Secure.”



Note Do not use a dollar sign (\$) symbol as part of a keystore password.

Logging In to the Web UIs

You can log into the Cisco Prime Network Registrar local or regional cluster web UIs either by HTTPS secure or HTTP nonsecure login. After installing Cisco Prime Network Registrar, open one of the supported web browsers and specify the login location URL in the browser address or netsite field. Login is convenient and provides some memory features to increase login speed.

You can log in using a nonsecure login in two ways:

- On Windows, from the Start menu, choose **Start > All Programs > Network Registrar 9.0 > Network Registrar 9.0 {local | regional} Web UI**. This opens the local or regional cluster web UI from your default web browser.



Note Open the regional Web UI first and add the licenses for the required services.

- Open the web browser and go to the web site. For example, if default ports were used during the installation, the URLs would be **http://hostname:8080** for the local cluster web UI, and **http://hostname:8090** for the regional cluster web UI.

This opens the New Product Installation page if no valid license is added at the time of installation. You have to browse and add the valid license. If the license key is acceptable, the Cisco Prime Network Registrar login page is displayed.

¹ This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).



Note You can add the licenses only in the regional server. The local has to be registered to the regional at the time of installation to run the desired licensed services.

In the local server, confirm the regional server IP address and port number and also the services you want to run at the time of your first login. Click **Register** to confirm registration. If the regional server is configured with the required licenses, you will be displayed the login page.

Enter the superuser username and password created at the time of installation to log into the Web UI. The password is case-sensitive (See [Managing Passwords, on page 52](#)). If you already added the valid license and superuser and configured a password at the time of installation, then you can log into the web UI using that username and password.



Note There is no default username or password for login.



Note To prepare for an HTTPS-secured login, see *Cisco Prime Network Registrar Installation Guide*.

Depending on how your browser is set up, you might be able to abbreviate the account name or choose it from a drop-down list while setting the username.

To log in, click **Login**.

Multiple Users

The Cisco Prime Network Registrar user interfaces support multiple, concurrent users. If two users try to access the same object record or data, a **Modified object** error will occur for the second user. If you receive this error while editing user data, do the following:

- **In the web UI**—Cancel the edits and refresh the list. Changes made by the first user will be reflected in the list. Redo the edits, if necessary.
- **In the CLI**—Use the **session cache refresh** command to clear the current edits, before viewing the changes and making further edits. Make changes, if you feel that it is necessary even after the other user's changes.

Changing Passwords

Whenever you edit a password on a web UI page, it is displayed as a string of eight dots. The actual password value is never sent to the web browser. So, if you change the password, the field is automatically cleared. You must enter the new password value completely, exactly as you want it to be.



Note The password should not be more than 255 characters long.

For details on changing administrator passwords at the local and regional cluster, see [Managing Passwords, on page 52](#).

Navigating the Web UIs

The web UI provides a hierarchy of pages based on the functionality you desire and the thread you are following as part of your administration tasks. The page hierarchy prevents you from getting lost easily.

**Caution**

Do not use the Back button of the browser. Always use the navigation menu, or the **Cancel** button on the page to return to a previous page. Using the browser Back button can cause erratic behavior or can cause failures.

A single sign-on feature is available to connect between the regional and local cluster web UIs. The regional cluster web UI pages include the Connect button in the List/Add Remote clusters page, which you can click to connect to the local cluster associated with the icon. If you have single sign-on privileges to the local cluster, the connection takes you to the related local server management page (or a related page for related server configurations). If you do not have these privileges, the connection takes you to the login page for the local cluster. To return to the regional cluster, local cluster pages have the Return button on the main toolbar.

The Search bar in the navigation menu provides an easy way to search for menus. The Pin icon in the top right corner of the navigation menu helps to pin/unpin the menu.

Starting from release 9.0, Cisco Prime Network Registrar provides a facility to save the frequently used pages/menus as favorites, which helps in accessing them easily. To configure the page/menu as favorite, after navigating to the desired menu, click the Favorite icon (star icon (*)) next to the navigation path), provide the appropriate name, and then click **OK**. The pages/menus which are configured as favorites appear under the Favorites section of the global navigation. You can delete the menus from the favorites list by clicking the Delete icon next to them. Configuration Summary page is listed under the Favorites section by default.

**Note**

Click the double arrow icon (⌂) in any page to view the hidden options/functionalities.

**Note**

Navigation menu items can vary based on if you have the role privileges for IPv4 or IPv6. For example, the **Design** menu can be **DHCPv4** and **DHCPv6** if you have the ipv6-management subrole of the addrblock-admin role assigned.

Waiting for Page Resolution Before Proceeding

Operations performed in the web UI, such as resynchronizing or replicating data from server clusters, are synchronous in that they do not return control to the browser until the operation is completed. These operations display confirmation messages in blue text. Also, both the Netscape and IE browsers display a wait cursor while the operation is in progress.

**Tip**

Wait for each operation in the web UI to finish before you begin a new operation. If the browser becomes impaired, close the browser, reopen it, then log in again. Some operations like zone distributions can take significant amount of time, so you may have to wait till the operation completes.

Committing Changes in the Web UIs

You do not actually commit the page entries you make until you click **Save** on the page. You can delete items using the Delete icon. To prevent unwanted deletions, a Confirm Delete dialog box appears in many cases so that you have a chance to confirm or cancel the deletion.

Role and Attribute Visibility Settings

Click the **Settings** drop-down list on the top of the main page to modify user preferences, session settings, user permissions, or debug settings.

- To view the user groups and roles for the administrator, select the **User Preferences** option. Superuser is a special kind of administrator. (For details how to set up these administrator roles, see [Create the Administrators](#), on page 112.)
- Select **Session Settings** to open the Session Settings dialog, select the mode from the **Session Web UI Mode** drop-down list, and click **Modify Session Settings**. You can also click the drop-down arrow of the Mode icon () to view the list of modes. Select the required mode from the list:
 - **Basic**—Basic user mode (the preset choice).
 - **Advanced**—Advanced user mode that exposes the normal attributes.
 - **Expert**—Expert user mode that exposes a set of attributes that are relevant for fine-tuning or troubleshooting the configuration. In most cases, you would accept the default values for these expert attributes and not change them without guidance from the Cisco Technical Assistance Center (TAC). Each Expert mode attribute is marked with a Warning icon on the configuration pages. Each page is clearly marked as being in Expert mode.

Displaying and Modifying Attributes

Many of the web UI pages, such as those for servers, zones, and scopes, include attribute settings that correspond to those you can set using the CLI. (The CLI name equivalents appear under the attribute name.) The attributes are categorized into groups by their function, with the more prominent attributes listed first and the ones less often configured nearer the bottom of the page.

Grouping and Sorting Attributes

On many Advanced mode web UI pages, you can toggle between showing attributes in groups and in alphabetical order. These pages generally open by default in group view so that you can see the attributes in their respective categories. However, in the case of large numbers of attributes, you might want to see the attributes alphabetized. Click **Show A-Z View** to change the page to show the attributes alphabetically. Click **Show Group View** to change the page to show the attributes in groups. You can also expand or collapse the attribute groups in group view by clicking **Expand All** or **Collapse All**. In Expert mode, the Expert mode attributes are alphabetized separately further down the page under the Visibility=3 heading and are all marked with the Warning icon.

Modifying Attributes

You can modify attribute values and unset those for optional attributes. In many cases, these attributes have preset values, which are listed under the Default column on the page. The explicit value overrides the default

one, but the default one is always the fallback. If there is no default value, unsetting the explicit value removes all values for that attribute.

Displaying Attribute Help

For contextual help for an attribute, click the name of the attribute to open a separate popup window.

Left Navigation Pane

The Web UI also provides a navigation pane on the left of the main pages. This navigation pane provides access to objects that are added as part of the various categories. The objects are listed in a tabular format and you can click the object to edit its properties in the main page.

Each object displayed under a category in the pane has a Quick View icon associated with it. The Quick View icon expands to open a dialog box that displays the main details about the object, and provides links (if any) to perform the main actions associated with the object.

By default, the list of objects is displayed in a single column format. However, you can add additional columns in the left pane. To add additional columns for objects, click the gear icon (⚙) above the objects table in the left pane, select the desired column names, and then click **Close**. You can save the column format by clicking the **Save Column Format** button.

There are Quick Filter and Advanced Filter options available to filter the objects as needed. To do a quick search for the objects, you can use the Quick Filter option. Click the Filter icon (▼) or select **Quick Filter** from the **Show** drop-down list located above the objects table and then enter the search string in the search bar. The objects are listed as per your search criteria.

You can also use Advanced Filter to filter the objects. Select **Advanced Filter** from the **Show** drop-down list, set the appropriate filter and condition in the Advanced Filter dialog box, and then click **OK**. Once you click OK, the object list on the left pane is filtered as per the filter specified. To save the filter, click **Save As** in the Advanced Filter dialog box, enter the appropriate name in the Save Filter dialog box, and then click **Save**. The saved filter name appears in the Show drop-down list and you can use this filter on that particular object list at any time. You can also set this filter as the default filter by clicking the **Set Default Filter** button.

The user defined filters can be edited or removed. To do this, select **Manage User Defined Filters** from the **Show** drop-down list, select the required user defined filter from the filter list in the Manage User Defined Filters dialog box, and then click **Edit** or **Remove** as required.

Help Pages

The web UI provides a separate window that displays help text for each page. The Help pages provide:

- A context-sensitive help topic depending on which application page you have open.
- A clickable and hierarchical Contents and Index, and a Favorites setting, as tabs on a left-hand pane that you can show or hide.
- A Search facility that returns a list of topics containing the search string, ordered by frequency of appearance of the search string.
- Forward and backward navigation through the history of Help pages opened.
- A Print function.
- A Glossary.

Logging Out

Log out of the web UI by clicking **Log Out** link. You can find the **Log Out** under the gear icon  at the top right corner of the application page.

Local Cluster Web UI

The local cluster web UI provides concurrent access to Cisco Prime Network Registrar user and protocol server administration and configuration. It provides granular administration across servers with permissions you can set on a per element or feature basis. The local cluster web UI is available in three user modes:

- **Basic Mode**— Provides a more simplified configuration for the more frequently configured objects, such as DHCP scopes and DNS zones (see [Local Basic Main Menu Page, on page 16](#)).
- **Advanced Mode**— Provides the more advanced configuration method familiar to past users of the Cisco Prime Network Registrar web UI, with some enhancements (see [Local Advanced Main Menu Page, on page 17](#)).
- **Expert Mode** (marked with the icon) - For details on Expert mode, see [Role and Attribute Visibility Settings, on page 14](#).

Change to Basic, Advanced, or Expert mode by clicking the drop-down arrow of the Mode icon () on the toolbar at the top right of the page (see [Setting Local User Preferences, on page 18](#)).



Note

If you change the IP address of your local cluster machine, see the Note in [Configuring Clusters in the Local Web UI, on page 19](#).

Related Topics

[Introduction to the Web-Based User Interfaces, on page 10](#)

[Regional Cluster Web UI, on page 20](#)

Local Basic Main Menu Page

The Basic tab activated on the toolbar at the top right corner of the page implies that you are in Basic user mode. Otherwise, click the drop-down arrow of the Mode icon () to view the list of modes and select **Basic**.

You can see the submenu items under the navigation menu by clicking the global navigation icon on the top left corner of the page. To choose a submenu under a navigation menu, place the cursor over the navigation menu item. For example, place the cursor on **Operate** to choose the **Manage Servers**.

Also, you can select any submenu under the required navigation menu and then navigate to the required submenu page from the left pane. For example, place the cursor on **Operate**, choose **Schedule Tasks**. You can see List/Add Scheduled Tasks page along with a left pane that has links to Manage Servers, Manage Clusters, Schedule Tasks, and View Change Log. Click the **Manage Servers** link to view the Manage Servers page.

The Local Basic main menu page provides functions with which you can:

- **Open the dashboard to monitor system health**—Open **Operate** menu and click **Dashboard**. See the "Server Status Dashboard" chapter.

- **Set up a basic configuration by using the Setup interview pages**—Click the **Setup** icon at the top, and select the different tabs in the Setup page. See *Cisco Prime Network Registrar Quick Start Guide* for more details.
- **Administer users, tenants, encryption keys**—Place the cursor on **Administration** menu (for user access options) or **Design** menu (for Security > Keys option). See [Managing Administrators, on page 37](#).
- **Manage the Cisco Prime Network Registrar protocol servers**—Place the cursor on **Operate** menu and select **Manage Servers** or **Schedule Tasks** option. See [Maintaining Servers and Databases, on page 129](#).
- **Manage clusters**—Place the cursor on **Operate** menu and choose **Manage Clusters** option. See [Configuring Server Clusters, on page 78](#).
- **Configure DHCP**—Place the cursor on **Design** menu and select the options under **DHCP Settings**, **DHCPv4** or **DHCPv6**. See the "Configuring Scopes and Networks" section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*.
- **Configure DNS**—Place the cursor on **Design** menu and select the options under **Cache DNS** and **Auth DNS**. Place the cursor on **Deploy** menu and select the options under **DNS** and **DNS Updates**. See the "Managing Zones" section in *Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*.
- **Manage hosts in zones**—From the **Design** menu, choose **Hosts** under the **Auth DNS** submenu. See the "Managing Hosts" section in *Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*.
- **Go to Advanced mode**—Click **Advanced** in the top right corner of the page. See [Local Advanced Main Menu Page, on page 17](#).

Local Advanced Main Menu Page

To switch to Advanced user mode from the Basic user Main Menu page, click the drop-down arrow of the Mode icon () at the top right of the window to view the list of modes and select **Advanced**. Doing so opens another Main Menu page, except that it shows the Advanced user mode functions. To switch back to Basic mode at any time, click next to the Mode icon at the top right of the window and select **Basic**.

The local Advanced mode Main Menu page includes advanced Cisco Prime Network Registrar functions that are in addition to the ones in Basic mode:

- **Open the dashboard to monitor system health**—Open **Operate** menu and click **Dashboard**. See the "Server Status Dashboard" chapter.
- **Administer users, tenants, groups, roles, regions, access control lists (ACLs), and view change logs**—Place the cursor on **Administration** menu (for user access options), **Design** menu (for ACLs) or **Operate** menu (for change logs). See [Managing Administrators, on page 37](#).
- **Manage the Cisco Prime Network Registrar protocol servers**—Place the cursor on **Operate** menu and select **Manage Servers** or **Schedule Tasks** option. See [Maintaining Servers and Databases, on page 129](#).
- **Manage clusters**—Place the cursor on **Operate** menu and choose **Manage Clusters** option. See [Configuring Server Clusters, on page 78](#).
- **Configure Routers**—Place the cursor on **Deploy** menu and select the options under **Router Configuration**. See [Managing Routers and Router Interfaces, on page 125](#).

- **Configure DHCPv4**—Place the cursor on **Design** and select any option under **DHCPv4**. See the *"Configuring Scopes and Networks"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*.
- **Configure DHCPv6**—Place the cursor on **Design** and select any option under **DHCPv6**. See the *"Managing DHCPv6 addresses"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*.
- **Configure DNS**—Place the cursor on **Design** menu and select the options under **Cache DNS** and **Auth DNS**. Place the cursor on **Deploy** menu and select the options under **DNS** and **DNS Updates**. See the *"Managing Zones"* section in *Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*.
- **Manage hosts in zones**—From the **Design** menu, choose **Hosts** under the **Auth DNS** submenu. See the *"Managing Hosts"* section in *Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*.
- **Manage IPv4 address space**—Place the cursor on **Design** and select any option under **DHCPv4**. See the *"Managing Address Space"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*.
- **Configure IPv6 address space**—Place the cursor on **Design** and select any option under **DHCPv6**. See the *"Managing DHCPv6 Addresses"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*.
- **Go to Basic mode**—Click the drop-down arrow of the Mode icon () at the top right corner of the page and choose **Basic**. See [Local Basic Main Menu Page, on page 16](#).

The Advanced user mode page provides additional functions:

- **View the user role and group data for the logged-in user**—See [Role and Attribute Visibility Settings, on page 14](#).
- **Set your preferred session settings**—See [Role and Attribute Visibility Settings, on page 14](#).
- **Set server debugging**—You can set debug flags for the protocol servers. Set these values only under diagnostic conditions when communicating with the Cisco Technical Assistance Center (TAC).
- **Change your login administrator password**—See [Managing Passwords, on page 52](#).

Setting Local User Preferences

You can maintain a short list of web UI settings through subsequent user sessions. The only difference between the Basic and Advanced or Expert mode user preference pages is that Advanced and Expert modes have additional columns listing the data types and defaults.

You can edit the user preferences by going to **User Preferences** under the **Settings** drop-down list. The user preference attributes to set are:

- **Username**—Username string, with a preset value of **admin**. You cannot modify this field.
- **Web UI list page size**—Adjust the page size by the number of displayed lines in a list; the preset value is 10 lines.
- **Web UI mode**—User mode at startup: Basic, Advanced, or Expert (see [Role and Attribute Visibility Settings, on page 14](#)). If unset, the mode defaults to the one set in the CCM server configuration (see [Managing Servers, on page 129](#)).
- **Web UI tree page size**—Adjust the page size when displaying a tree view in the web UI.
- **Web UI log page size**—Adjust the page size on log pages.
- **Views**—Specify the DNS view setting at session startup in the web UI or CLI.

- **VPN**—Specify the VPN setting at session startup in the web UI or CLI.
- **Alarm poll interval**—Adjust the alarm poll interval; that is, how often Network Registrar polls the alarm data from server.
- **Homepage**—Set a page from favorites list as the homepage for the application. By default, Configuration Summary page is set as the homepage. Starting from release 9.0, you can set a page of your choice as the homepage for the application. To do this, add the desired page to the Favorites list (see [Navigating the Web UIs, on page 13](#)), select the page name from the Homepage drop-down list, and then click **Modify User Preferences**. You can click the Home icon (🏠) on the top left corner of the web UI to go to the homepage.
- **Date format**—Set the date-time format for date-time values in the web UI. A format can be selected from the default list or entered in text form as <date-pattern> <time-pattern>.

Supported patterns are:

- Year as "yy", "yyyy"
- Month as "M", "MM", "MMM", "MMMM"
- Day as "d", "dd"
- Hour as "h", "hh", "H", "HH"
- Minute as "mm"
- Second as "s", "ss"
- Delimiters as ":", "-", "/"

You can unset the page size and web UI mode values by checking the check box in the *Unset?* column, next to the attribute. After making the user preference settings, click **Modify User Preferences**.

Configuring Clusters in the Local Web UI

You can define other local Cisco Prime Network Registrar clusters in the local web UI. The local cluster on the current machine is called the **localhost** cluster. To set up other clusters, choose **Manage Clusters** from **Operate** menu to open the List/Add Clusters page. Note that the **localhost** cluster has the IP address and SCP port of the local machine.

Click the **Add Cluster** icon in the left pane to open the Add Cluster page. At a minimum, you must enter the name and address (IPv4 and/or IPv6) of the remote local cluster. You should also enter the admin name and password, along with possibly the SCP port (if not 1234) of the remote cluster. Click **Add Cluster**. To edit a cluster, click the cluster name in the Clusters pane on the left to open the Edit Cluster page. If you want to use secure access mode, select use-ssl as disabled, optional, or required (optional is the preset value; you need the security library installed if you choose required). Make the changes and then click **Save**.



Note If you change the IP address of your local cluster machine, you must modify the **localhost** cluster to change the address in the ipaddr field. Avoid setting the value to the loopback address (127.0.0.1); if you do, you must also set the actual IP addresses of main and backup servers for DHCP failover and High-Availability (HA) DNS configurations.

Regional Cluster Web UI

The regional cluster web UI provides concurrent access to regional and central administration tasks. It provides granular administration across servers with permissions you can set on a per element or feature basis. After you log into the application, the Home page appears. Regional cluster administration is described in [Managing the Central Configuration](#), on page 73.

Related Topics

[Introduction to the Web-Based User Interfaces](#), on page 10

[Local Cluster Web UI](#), on page 16

Command Line Interface

Using the Cisco Prime Network Registrar CLI (the **nrcmd** program), you can control your local cluster server operations. You can set all configurable options, as well as start and stop the servers.



Note The CLI provides concurrent access, by at most 14 simultaneous users and processes per cluster.



Tip See the **CLIContents.html** file in the **/docs** subdirectory of your installation directory for details.

The **nrcmd** program for the CLI is located on:

- **Windows**—In the *install-path* \bin directory.
- **Linux**—In the *install-path* /usrbin directory.

On a local cluster, once you are in the appropriate directory, use the following command at the prompt:

```
nrcmd [-C cluster[:port]] [-N user] [-P password] [-h] [-r] [-v] [-b < script | command]
```

- **-C**—Cluster name, preset value **localhost**. Specify the port number with the cluster name while invoking **nrcmd** to connect to another cluster. See the preceding example.

The port number is optional if the cluster uses the default SCP port—1234 for local and 1244 for regional. Ensure that you include the port number if the port used is not the default one.

- **-N**—Username. You have to enter the username that you created when first logged into the Web UI.
- **-P**—User password. You have to enter the password that you created for the username.
- The local cluster (**-L**) is implied; use **-R** to open the regional cluster CLI.
- **-b < script** Process script file of **nrcmd** commands.
- **-h** Print this help text.
- **-r** Login as a read-only user.
- **-R** Connect to regional.
- **-v** (or **-vv**) Report the program version and exit.

- -V Specify the session visibility



Note Cluster defaults to localhost if not specified.



Tip For additional command options, see the **CLIGuide.html** file in /docs.



Note If you change the IP address of your local cluster machine, you must modify the **localhost** cluster to change the address in the *ipaddress* attribute. Do not set the value to 127.0.0.1.

You can also send the output to a file using:

```
nrcmd> session log filename
```

For example:

To send the leases on the DHCP server to a file (leases.txt), use the following commands:

```
nrcmd> session log leases.txt
nrcmd> lease list
```



Note To close a previously opened file, use session log (no filename). This stops writing the output to any file.

To disconnect from the cluster, use **exit**:

```
nrcmd> exit
```



Tip The CLI operates on a coordinated basis with multiple user logins. If you receive a cluster lock message, determine who has the lock and discuss the issue with that person. (See [Multiple Users, on page 12.](#))

Global Search in Prime Network Registrar

The Local and Regional Web UI in Prime Network Registrar also provides a global search functionality for the IP addresses or DNS names available in the local clusters. The search interface element is available at the top right corner of the main page.



Note To view the search interface element and run the search for IP addresses and DNS names, Cisco Prime Network Registrar must be licensed with DHCP or DNS, and the DHCP or DNS services must be enabled for the local cluster (in the List/Add Remote Clusters page in Regional Web UI).

The following table shows the typical search results under different scenarios.

Table 2: Typical Search Results

You search for...	With active licenses and services for...	Search Results
An IPv4 address	Only DHCP	The closest matching scope, scope lease or scope reservation
An IPv4 address or a DNS FQDN	Only DNS	The related Zone or Resource Record
An IPv6 address	Only DHCP	The closest matching prefix, prefix lease or prefix reservation
An IPv6 address or a DNS FQDN	Only DNS	The related Zone or Resource Record
An IPv4 address, an IPv6 address or a DNS FQDN	Both DHCP and DNS	All of the above, based on the type of address



CHAPTER 3

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 23](#)
- [Display Types, on page 24](#)
- [Customizing the Display, on page 28](#)
- [Selecting Dashboard Elements to Include, on page 30](#)
- [Host Metrics, on page 32](#)

Opening the Dashboard

Starting from Cisco Prime Network Registrar 9.0, 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

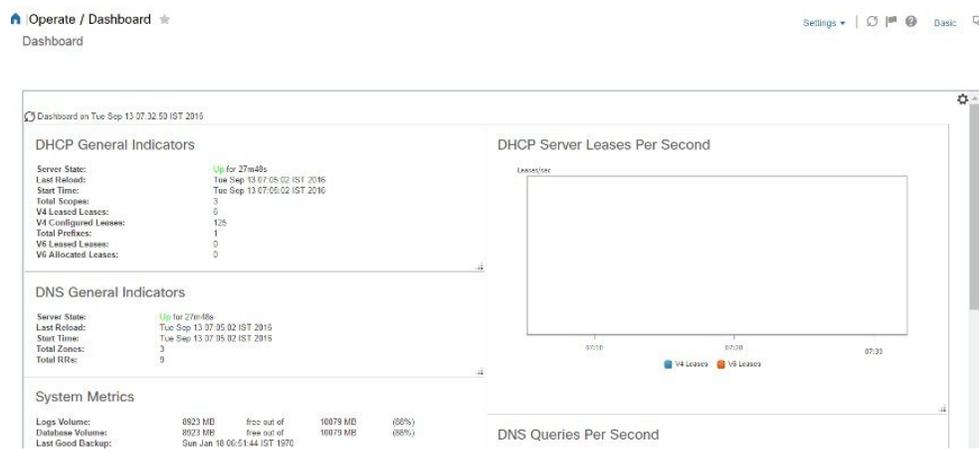
Provided you have DHCP and DNS privileges through administrator roles assigned to you, the preset display of the dashboard consists of the following tables (See the table below for an example):

- **System Metrics**—See [System Metrics](#), on page 32.
- **DHCP General Indicators**—See the *"DHCP General Indicators"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*.
- **DNS General Indicators**—See the *"DNS General Indicators"* section in *Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*.



Tip These are just the preset selections. See [Selecting Dashboard Elements to Include](#), on page 30 for other dashboard elements you can select. The dashboard retains your selections from session to session.

Figure 4: Preset Dashboard Elements



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

- **Table**—See [Tables](#), on page 25.
- **Line chart**—See [Line Charts](#), on page 25.
- **Stacked area chart**—See [Stacked Area Charts](#), on page 27.

General Status Indicators

Note the green box next to each dashboard element name in the above image. This box indicates that the server sourcing the information is functioning normally. A yellow box indicates that server operation is less than optimum. A red box 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

If Magnified Chart is the selected Chart Link, you can magnify a chart in a separate window by clicking the chart. In magnified chart view, you can choose an alternative chart type from the one that comes up initially (see [Other Chart Types, on page 27](#)).



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 initially includes a color-coded legend. Removing the legend renders the graphic chart size relatively larger, which can be helpful if you have many charts displayed. You cannot remove legends in magnified views.

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:

- System Metrics
- 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.

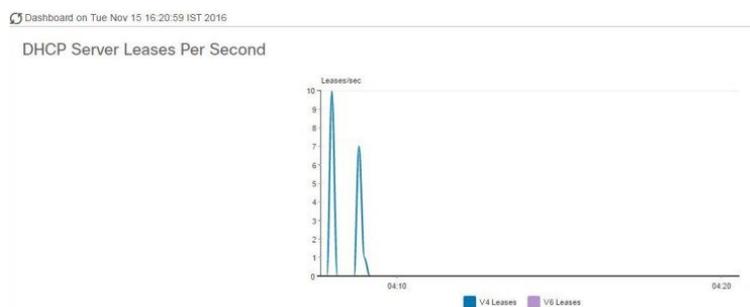
Table 3: Line Chart Types

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 Transfer
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, set the Chart Link to Data Table, then click the chart. The Raw Data table is below the Chart Data table.

Figure 5: Line Chart Example



Stacked Area Charts

Dashboard elements rendered as stacked 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 5: Line Chart Example, on page 26](#) rendered as a stacked area chart.)

Figure 6: Stacked Area Chart Example

Dashboard on Tue Nov 15 16:20:13 IST 2016

DHCP Server Leases Per Second



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 stacked area charts are:

- DHCP Server Request Activity
- DHCP Server Response Activity
- DHCP Response Latency
- DNS Outbound Zone Transfers
- DNS Inbound Zone Transfers

Other Chart Types

The other chart types available for you to choose are:

- **Line**—One of the line charts described in [Table 3: Line Chart Types, on page 26](#).
- **Stacked Area**—Charts described in the [Stacked Area Charts, on page 27](#).
- **Pie**—Shows a single percentage pie chart of the data averaged over the time sampled.
- **Bar**—Multiple related current value metrics plotted side by side as groups of bars that show the actual data sampled.
- **Stacked Bar**—Addition total of the actual samples. This chart shows more distinct data points than the stacked area chart.



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 title of the element.

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 chart title.



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

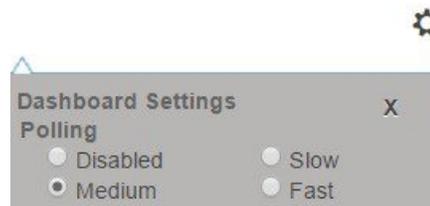
Refreshing Displays

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 7: 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

You can choose to display a graphic chart as a table when you magnify the chart by clicking it. At the middle of the top of the dashboard display are the controls for the chart links (see the image below)

Figure 8: Specifying Chart Conversion to Table Format



Click the **Data Table** radio button. When you click the chart itself, it opens as a table. The preset display format is Magnified Chart. Displaying Charts as Tables

Exporting to CSV Format

You can dump the chart data to a comma-separated value (CSV) file (such as a spreadsheet) when you magnify the chart by clicking it. In the Chart Link controls at the top of the page (see the above image), click the **CSV Export** radio button, then click the chart. A Save As window appears, where you can specify the name and location of the CSV file.

Displaying or Hiding Chart Legends

You can include or exclude the color-coded legends for charts on the main dashboard page. You might want to remove the legends as you become more familiar with the data and track it on a slightly larger chart display. In the upper-right of the dashboard display are the controls for the legend display (see the image below). The preset value is Visible.

Figure 9: Displaying or Hiding Chart Legends and Selecting Chart



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, 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 **Chart Selections** in the Dashboard Settings dialog. Clicking the link opens the Chart Selection page.

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).

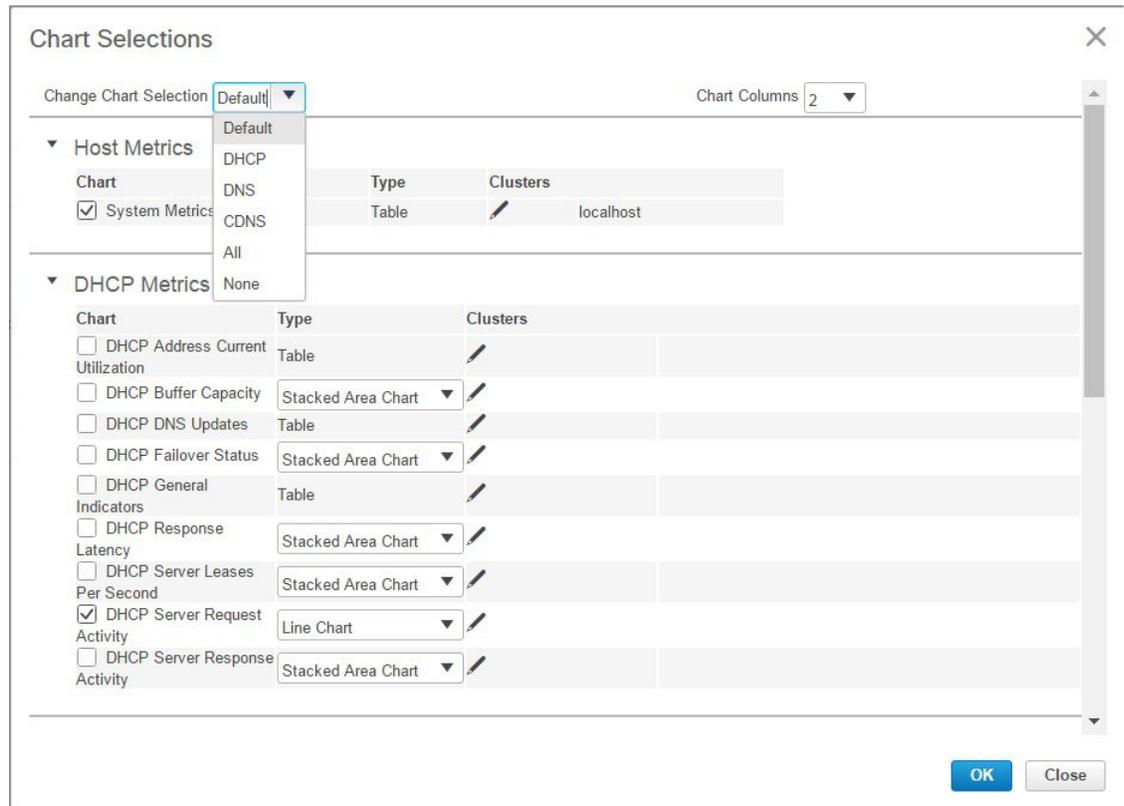


Note You can see either the CDNS or DNS Metrics in the **Dashboard Settings > Chart Selection** page based on the service configured on the server.



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 10: 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 drop-down list, **Change Chart Selection**, at the top of the page. 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 9.0 DHCP User Guide*).
- Select the DNS metrics only, choose **DNS** (see the "Dashboard and Authoritative DNS Metrics" section in *Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*).
- Select the DNS metrics only, choose **CDNS** (see the "Caching DNS Metrics" section in *Cisco Prime Network Registrar 9.0 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.

Host Metrics

Host metrics comprise two charts:

- **System Metrics**—See the [System Metrics, on page 32](#).
- **JVM Memory Utilization** (available in Expert mode only).

System Metrics

The System Metrics dashboard element shows the free space on the disk volumes where the Cisco Prime Network Registrar logs and database directories are located, the date and time of the last server backup, and CPU and memory usage for the various servers. System metrics are available if you choose **Host Metrics: System Metrics** in the Chart Selection list.

The resulting table shows:

- **Logs Volume**—Current free space out of the total space on the disk drive where the logs directory is located, with the equivalent percentage of free space.
- **Database Volume**—Current free space out of the total space on the disk drive where the data directory is located, with the equivalent percentage of free space.
- **Last Good Backup**—Date and time when the last successful shadow database backup occurred (or Not Done if it did not yet occur) since the server agent was last started.
- **CPU Utilization** (in seconds), **Memory Utilization** (in kilobytes), and (in Expert mode only) the **VM Utilization** (in kilobytes) and Process ID (**PID**) for the:
 - Cisco Prime Network Registrar server agent
 - CCM server
 - DNS server
 - DHCP server
 - Web server
 - SNMP server
 - DNS caching server

How to Interpret the Data

The System Metrics data shows how full your disk volumes are getting based on the available free space for the Cisco Prime Network Registrar logs and data volumes. It also shows if you had a last successful backup of the data files and when that occurred. Finally, it shows how much of the available CPU and memory the Cisco Prime Network Registrar servers are using. The difference in the memory and VM utilization values is:

- **Memory Utilization**—Physical memory that a process uses, or roughly equivalent to the Resident Set Size (RSS) value in UNIX **ps** command output, or to the Task Manager Mem Usage value in Windows: the number of pages the process has in real memory minus administrative usage. This value includes only the pages that count toward text, data, or stack space, but not those demand-loaded in or swapped out.
- **VM Utilization**—Virtual memory that a process uses, or roughly equivalent to the SZ value in UNIX **ps** command output, or to the Task Manager VM Size value in Windows: the in-memory pages plus the page files and demand-zero pages, but not usually the memory-mapped files. This value is useful in diagnosing how large a process is and if it continues to grow.

Troubleshooting Based on the Results

If you notice the free disk space decreasing for the logs or data directory, you might want to consider increasing the disk capacity or look at the programs you are running concurrently with Cisco Prime Network Registrar.

JVM Memory Utilization

The Java Virtual Machine (JVM) Memory Utilization dashboard element is available only when you are in Expert mode. It is rendered as a line trend chart that traces the Unused Maximum, Free, and Used bytes of JVM memory. The chart is available if you choose **Host Metrics: JVM Memory Utilization** in the Chart Selection list when you are in Expert mode.

How to Interpret the Data

The JVM Memory Utilization data shows how much memory applies to running the dashboard in your browser. If you see the Used byte data spiking, dashboard elements might be using too much memory.

Troubleshooting Based on the Results

If you see spikes in Used memory data, check your browser settings or adjust the polling interval to poll for data less frequently.



PART II

Local and Regional Administration

- [Managing Administrators, on page 37](#)
- [Managing Owners and Regions, on page 69](#)
- [Managing the Central Configuration, on page 73](#)
- [Managing Routers and Router Interfaces, on page 125](#)
- [Maintaining Servers and Databases, on page 129](#)
- [Backup and Recovery, on page 163](#)
- [Managing Reports, on page 179](#)



CHAPTER 4

Managing Administrators

This chapter explains how to set up network administrators at the local and regional clusters. The chapter also includes local and regional cluster tutorials for many of the administration features.

- [Administrators, Groups, Roles, and Tenants, on page 37](#)
- [External Authentication Servers, on page 42](#)
- [Managing Tenants, on page 46](#)
- [Managing Administrators, on page 51](#)
- [Managing Passwords, on page 52](#)
- [Managing Groups, on page 52](#)
- [Managing Roles, on page 53](#)
- [Granular Administration, on page 54](#)
- [Centrally Managing Administrators, on page 58](#)

Administrators, Groups, Roles, and Tenants

The types of functions that network administrators can perform in Cisco Prime Network Registrar are based on the roles assigned to them. Local and regional administrators can define these roles to provide granularity for the network administration functions. Cisco Prime Network Registrar predefines a set of base roles that segment the administrative functions. From these base roles you can define further constrained roles that are limited to administering particular addresses, zones, and other network objects.

The mechanism to associate administrators with their roles is to place the administrators in groups that include these roles.

The data and configuration that can be viewed by an administrator can also be restricted by tenant. When an administrator is assigned a tenant tag, access is further restricted to configuration objects that are assigned to the tenant or made available for tenant use as read-only core configuration objects.

Related Topics

[How Administrators Relate to Groups, Roles, and Tenants, on page 38](#)

[Administrator Types, on page 38](#)

[Roles, Subroles, and Constraints, on page 39](#)

[Groups, on page 42](#)

[Managing Administrators, on page 51](#)

[Managing Passwords, on page 52](#)

[Managing Groups, on page 52](#)

[Managing Roles, on page 53](#)

[Managing Tenants, on page 46](#)

How Administrators Relate to Groups, Roles, and Tenants

There are four administrator objects in Cisco Prime Network Registrar—administrator, group, role, and tenant:

- **Administrator**—An account that logs in and that, through its association with one or more administrator groups, can perform certain functions based on its assigned role or roles. At the local cluster, these functions are administering the local Central Configuration Management (CCM) server and databases, hosts, zones, address space, and DHCP. At the regional cluster, these functions administer the regional CCM server and databases, central configuration, and regional address space. An administrator must be assigned to at least one group to be effective.

Adding administrators is described in [Managing Administrators, on page 51](#).

- **Group**—A grouping of roles. You must associate one or more groups with an administrator, and a group must be assigned at least one role to be usable. The predefined groups that Cisco Prime Network Registrar provides map each role to a unique group.

Adding groups is described in [Managing Groups, on page 52](#).

- **Role**—Defines the network objects that an administrator can manage and the functions that an administrator can perform. A set of predefined roles are created at installation, and you can define additional constrained roles. Some of the roles include subroles that provide further functional constraints.

Adding roles is described in [Managing Roles, on page 53](#).

- **Tenant**—Identifies a tenant organization or group that is associated with a set of administrators. When you create tenants, the data stored on both regional and local clusters is segmented by tenant. A tenant cannot access the data of another tenant.

Adding tenants is described in [Managing Tenants, on page 46](#).

Administrator Types

There are two basic types of administrators: superusers and specialized administrators:

- **Superuser**—Administrator with unrestricted access to the web UI, CLI, and all features. This administrator type should be restricted to a few individuals. The superuser privileges of an administrator override all its other roles.



Tip You have to create the superuser and password at installation, or when you first log into the web UI.

When a superuser is assigned a tenant tag, unrestricted access is only granted for corresponding tenant data. Data of other tenants cannot be viewed, and core objects are restricted to read-only access.

- **Specialized**—Administrator created by name to fulfill specialized functions, for example, to administer a specific DNS forward or reverse zone, based on the administrator assigned role (and subrole, if applicable). Specialized administrators, like the superuser, require a password, but must also be assigned at least one administrator group that defines the relevant roles. The CLI provides the **admin** command.

For an example of creating a local zone or host administrator, see [Create the Administrators, on page 112](#).

A specialized user that is assigned a tenant tag can only access corresponding tenant or core data that also matches the relevant roles. Core data is further restricted to read-only access.

Roles, Subroles, and Constraints

A license type is associated with each role-subrole combination. A role-subrole is enabled only if that license is available in that cluster.

You can limit an administrator role by applying constraints. For example, you can use the host-admin base role to create a host administrator, named 192.168.50-host-admin, who is constrained to the 192.168.50.0 subnet. The administrator assigned a group that includes this role then logs in with this constraint in effect. Adding roles and subroles is described in [Managing Roles, on page 53](#).

You can further limit the constraints on roles to read-only access. An administrator can be allowed to read any of the data for that role, but not modify it. However, if the constrained data is also associated with a read-write role, the read-write privilege supersedes the read-only constraints.



Tip An example of adding role constraints is in [Create a Host Administrator Role with Constraints, on page 115](#).

The interplay between DNS and host administrator role assignments is such that you can combine an unconstrained dns-admin role with any host-admin role in a group. For example, combining the dns-admin-readonly role and a host-admin role in a group (and naming the group host-rw-dns-ro) provides full host access and read-only access to zones and RRs. However, if you assign a constrained dns-admin role along with a host-admin role to a group and then to an administrator, the constrained dns-admin role takes precedence, and the administrator privileges at login will preclude any host administration.

Certain roles provide subroles with which you can further limit the role functionality. For example, the local ccm-admin or regional-admin, with just the owner-region subrole applied, can manage only owners and regions. By default, all the possible subroles apply when you create a constrained role.

The predefined roles are described in [Table 4: Local Cluster Administrator Predefined and Base Roles , on page 40](#) (local), and [Table 5: Regional Cluster Administrator Predefined and Base Roles , on page 41](#) (regional).

Table 4: Local Cluster Administrator Predefined and Base Roles

Local Role	Subroles and Active Functionality
addrblock-admin	<p>Core functionality: Manage address block, subnets, and reverse DNS zones (also requires dns-admin); and notify of scope activity.</p> <ul style="list-style-type: none"> • <i>ric-management</i>: Push to, and reclaim subnets from, DHCP failover pairs and routers. • <i>ipv6-management</i>: Manage IPv6 prefixes, links, options, leases, and reservations.
ccm-admin	<p>Core functionality: Manage access control lists (ACLs), and encryption keys.</p> <ul style="list-style-type: none"> • <i>authentication</i>: Manage administrators. • <i>authorization</i>: Manage roles and groups. • <i>owner-region</i>: Manage owners and regions. • <i>database</i>: View database change entries and trim the CCM change sets.
cdns-admin	<p>Core functionality: Manage in-memory cache (flush cache and flush cache name).</p> <ul style="list-style-type: none"> • <i>security-management</i>: Manage ACLs and DNSSEC configuration. • <i>server-management</i>: Manage DNSSEC configuration, as well as forwarders, exceptions, DNS64, and scheduled tasks, and stop, start, or reload the server.
cfg-admin	<p>Core functionality: Manage clusters.</p> <ul style="list-style-type: none"> • <i>ccm-management</i>: Manage the CCM server configuration. • <i>dhcp-management</i>: Manage the DHCP server configuration. • <i>dns-management</i>: Manage the DNS server configuration. • <i>cdns-management</i>: Manage Caching DNS server configuration. • <i>ric-management</i>: Manage routers. • <i>snmp-management</i>: Manage the SNMP server configuration. • <i>tftp-management</i>: Manage the TFTP server configuration.

Local Role	Subroles and Active Functionality
dhcp-admin	<p>Core functionality: Manage DHCP scopes and templates, policies, clients, client-classes, options, leases, and reservations.</p> <ul style="list-style-type: none"> • <i>server-management</i>: Manage the DHCP server configuration, failover pairs, LDAP servers, extensions, and statistics. • <i>ipv6-management</i>: Manage IPv6 prefixes, links, options, leases, and reservations.
dns-admin	<p>Core functionality: Manage DNS zones and templates, resource records, secondary servers, and hosts.</p> <ul style="list-style-type: none"> • <i>security-management</i>: Manage DNS update policies, ACLs, and encryption keys. • <i>server-management</i>: Manage DNS server configurations and zone distributions, synchronize zones and HA server pairs, and push update maps. • <i>ipv6-management</i>: Manage IPv6 zones and hosts.
host-admin	<p>Core functionality: Manage DNS hosts. (Note that if an administrator is also assigned a constrained dns-admin role that overrides the host-admin definition, the administrator is not assigned the host-admin role.)</p>

Table 5: Regional Cluster Administrator Predefined and Base Roles

Regional Role	Subroles and Active Functionality
central-cfg-admin	<p>Core functionality: Manage clusters and view replica data.</p> <ul style="list-style-type: none"> • <i>dhcp-management</i>: Manage DHCP scope templates, policies, client-classes, failover pairs, virtual private networks (VPNs), and options; modify subnets; and replicate data. • <i>ric-management</i>: Manage routers and router interfaces, and pull replica router data.
central-dns-admin	<p>Core functionality: Manage DNS zones and templates, hosts, resource records, and secondary servers; and create subzones and reverse zones.</p> <ul style="list-style-type: none"> • <i>security-management</i>: Manage DNS update policies, ACLs, and encryption keys. • <i>server-management</i>: Synchronize DNS zones and HA server pairs, manage zone distributions, pull replica zone data, and push update maps.

Regional Role	Subroles and Active Functionality
central-host-admin	Core functionality: Manage DNS hosts. (Note that if an administrator is also assigned a constrained central-dns-admin role that overrides the central-host-admin definition, the administrator is not assigned the central-host-admin role.)
regional-admin	Core functionality: Manage licenses and encryption keys. <ul style="list-style-type: none"> • <i>authentication</i>: Manage administrators. • <i>authorization</i>: Manage roles and groups. • <i>owner-region</i>: Manage owners and regions. • <i>database</i>: View database change entries and trim the CCM change sets.
regional-addr-admin	Core functionality: Manage address blocks, subnets, and address ranges; generate allocation reports; and pull replica address space data. <ul style="list-style-type: none"> • <i>dhcp-management</i>: Push and reclaim subnets; and add subnets to, and remove subnets from, DHCP failover pairs. • <i>lease-history</i>: Query, poll, and trim lease history data. • <i>subnet-utilization</i>: Query, poll, trim, and compact subnet utilization data.

Groups

Administrator groups are the mechanism used to assign roles to administrators. Hence, a group must consist of one or more administrator roles to be usable. When you first install Cisco Prime Network Registrar, a predefined group is created to correspond to each predefined role.

Roles with the same base role are combined. A group with an unconstrained dhcp-admin role and a constrained dns-admin role, does not change the privileges assigned to the dns-admin role. For example, if one of the roles is assigned unconstrained read-write privileges, the group is assigned unconstrained read-write privileges, even though other roles might be assigned read-only privileges. Therefore, to limit the read-write privileges of a user while allowing read-only access to all data, create a group that includes the unconstrained read-only role along with a constrained read-write role. (See [Roles, Subroles, and Constraints](#), on page 39 for the implementation of host-admin and dns-admin roles combined in a group.)

External Authentication Servers

Cisco Prime Network Registrar includes a RADIUS client component and Active Directory (AD) client component, which are integrated with the authentication and authorization modules of the CCM server. To enable external authentication, you must configure a list of external RADIUS or an AD server at local and regional clusters, and ensure all authorized users are appropriately configured on the respective servers.

When external authentication is enabled, the CCM server handles attempts to log in via the web UI, SDK, or CLI, by issuing a RADIUS request to a RADIUS server or a LDAP request to a AD server that is selected from the configured list. If the corresponding server validates the login request, access is granted, and the CCM server creates an authorized session with the group assignments specified by the RADIUS or the AD server.



Note Any administrators defined in the CCM server's database are ignored when external authentication is enabled. Attempting to log in with these usernames and passwords will fail. To disable external authentication, you must remove or disable all the configured external servers or change the auth-type attribute value to Local.



Tip If all logins fail because the RADIUS servers are inaccessible or misconfigured, use the local.superusers file to create a temporary username and password. See [Managing Administrators, on page 51](#) for more details.

Configuring an RADIUS External Authentication Server

Cisco Prime Network Registrar administrators must be assigned to one or more administrator groups to perform management functions. When using a RADIUS server for external authentication, these are set as a vendor specific attribute for each user. Using the Cisco vendor id (9), create the Cisco Prime Network Registrar groups attribute for each administrator, using the format `cnr:groups=group1, group2, group3`.

For example, to assign an administrator to the built-in groups `dhcp-admin-group` and `dns-admin-group`, enter:

```
cnr:groups=dhcp-admin-group, dns-admin-group
```

To assign superuser access privileges, the reserved group name `superusers` is used. To provide superuser privileges to an administrator, enter:

```
cnr:groups=superusers
```

The superuser privileges override all other groups



Note You cannot add, delete, or modify external user names and their passwords or groups using Cisco Prime Network Registrar. You must use the RADIUS server to perform this configuration.

Adding an RADIUS External Configuration Server

To add an external configuration server, do the following:

Local Advanced and Regional Web UI

Step 1 From the **Administration** menu, choose **Radius** under the External Authentication submenu. The List/Add Radius Server page is displayed.

Step 2 Click the **Add Radius** icon in the Radius pane, enter the name, IPv4 and/or IPv6 address of the server you want to configure as the external authentication server, and you can set the key attribute which will be used for communicating

with this server in the Add External Authentication Server dialog box, and click **Add External Authentication Server**. The CCM server uses the key to set the key-secret attribute which is the secret key shared by client and the server.

- Step 3** To enable the external authentication server, check **enabled** check box of the ext-auth attribute in the Edit Authentication Server page, and then click **Save**.
- Step 4** Change the auth-type attribute to RADIUS in the Manage Server page, click **Save**, and then restart Cisco Prime Network Registrar.

CLI Commands

To create an external authentication server, use **auth-server name create** <address | ip6address> [*attribute=value ...*] (see the **auth-server** command in the CLIGuide.html file in the /docs directory for syntax and attribute descriptions).

Deleting an RADIUS External Authentication Server

Local Advanced and Regional Web UI

To delete an RADIUS external authentication server, select the server in the Radius pane, click the **Delete Radius** icon, and then confirm the deletion. You can also cancel the deletion by clicking the Close button.

Configuring an AD External Authentication Server

Cisco Prime Network Registrar administrators must be assigned to one or more administrator groups to perform management functions. When using an AD server for external authentication, these are set as a vendor specific attribute for each user. Using the Cisco vendor id (9), create the Cisco Prime Network Registrar groups attribute for each administrator, using the format **cnr:groups=group1, group2, group3**.

For example, to assign an administrator to the built-in groups **dhcp-admin-group** and **dns-admin-group**, enter:

```
cnr:groups=dhcp-admin-group,dns-admin-group
```

To assign superuser access privileges, the reserved group name **superusers** is used. To provide superuser privileges to an administrator, enter:

```
cnr:groups=superusers
```

The superuser privileges override all other groups.

A group needs to be created to access CPNR and add the users to that group. Select an user attribute and provide the group information in the format **cnr:group1,group2,..**

To configure an Active Directory (AD) external authentication server:

- Step 1** In AD server, create a new group, for example **CPNR**, with the group scope *Domain Local*.
- Step 2** Select a user and click **Add** to a group.
- Step 3** In Enter the Object Names window, select **CPNR** and click **OK**.
- Step 4** In AD Server Object windows, select **CPNR** for the *ad-group-name* attribute and **info** for the *ad-user-attr-map* attribute.

Note You cannot add, delete, or modify external user names and their passwords or groups using Cisco Prime Network Registrar. You must use the AD server to perform this configuration.

Configuring Kerbero's Realm and KDC

For the Cisco Prime Network Registrar to communicate with the AD server, the Kerbero's Realm and KDC servers are required. To configure the Kerbero's Realm and KDC servers in Windows and Linux platforms follow the below examples.

If the Cisco Prime Network Registrar is running on Windows platform (ksetup), define a KDC entry for a realm by running the following command:

```
ksetup /AddKdc <RealmName> [KdcName]
```

For example, `ksetup /AddKdc ECNR.COM tm-chn-ecnr-ad.ecnr.com`

To verify, run the following command:

```
ksetup /dumpstate
```

The result should be similar to the message below:

```
default realm = partnet.cisco.com (NT Domain)
ECNR.COM:
  kdc = tm-chn-ecnr-ad.ecnr.com
  Realm Flags = 0x0No Realm Flags
No user mappings defined.
```

If the Prime Network Registrar is running on Linux platform, the changes need to be configured in **krb5.conf** (*/etc/krb5.conf*) file, as shown below:

```
default = FILE:/var/log/krb5libs.log
kdc = FILE:/var/log/krb5kdc.log
admin_server = FILE:/var/log/kadmind.log
[libdefaults]
  ticket_lifetime = 1d
  default_realm = ECNR.COM
  default_tkt_enctypes = rc4-hmac
  default_tgs_enctypes = rc4-hmac
  dns_lookup_realm = false
  dns_lookup_kdc = false
  forwardable = true
[realms]
  ECNR.COM = {
    kdc = <kdc server host name>
    admin_server = <kdc server host name>
  }
[domain_realm]
  .ecnr.com = ECNR.COM
  ecnr.com = ECNR.COM
```

Adding an AD External Configuration Server

To add an external configuration server, do the following:

Local Advanced and Regional Web UI

Step 1 From the **Administration** menu, choose **Active Directory** under the **External Authentication** submenu. The List/Add Active Directory Server page is displayed.

- Step 2** Click the **Add Active Directory Server** icon in the Active Directory pane, enter the name, hostname of the server, domain you want to configure as the external authentication server, and you can set the base domain, LDAP user attribute map, AD group name which will be used for communicating with this server in the Add Active Directory Server dialog box, and click **Add Active Directory Server**.
- Step 3** Change the auth-type attribute to Active Directory in the Manage Server page, click **Save**, and then restart Cisco Prime Network Registrar.

CLI Commands

To create an external authentication server, use **auth-server name create** *<address | ip6address>* [*attribute=value ...*].

Deleting an AD External Authentication Server

Local Advanced and Regional Web UI

To delete an AD external authentication server, select the server in the Active Directory pane, click the **Delete Active Directory Server** icon, and then confirm the deletion. You can also cancel the deletion by clicking the Close button.

Managing Tenants

The multi-tenant architecture of Cisco Prime Network Registrar provides the ability to segment the data stored on both regional and local clusters by tenant. When tenants are defined, data is partitioned by tenant in the embedded databases of each cluster. This provides data security and privacy for each tenant, while allowing cloud or managed service providers the flexibility to consolidate many smaller customer configurations on a set of infrastructure servers, or distribute a larger customer configuration across several dedicated servers.

Any given local cluster may be associated with one or more tenants, but within a local cluster, the address pools and domain names assigned to a given tenant must not overlap.

For larger customers, clusters may be explicitly assigned to a tenant. In this case, all data on the local cluster will be associated with the tenant, and may include customized server settings. Alternatively, infrastructure servers may service many tenants. With this model, the tenants can maintain their own address space and domain names, but share common server settings that would be administered by the service provider. Their use of public or private network addresses needs to be managed by the service provider, to ensure that the tenants are assigned non-overlapping addresses.

The following are the key points you should know while configuring tenants:

- Tenant administrators are linked to their data by a tenant object that defines their tenant tag and identifier.
- Tenant objects should be consistent and unique across all clusters.
- You should not reuse tags or identifiers for different tenants.
- You can configure multiple tenants on a single cluster.
- A tenant administrator cannot create, modify, or remove tenant objects.
- A tenant administrator cannot view or modify the data of another tenant.
- Objects that are not assigned to a tenant are defined as core data, and are visible to all tenants in read-only mode.

Adding a Tenant

To add a tenant, do the following:

Local and Regional Advanced Web UI

- Step 1** From the **Administration** menu, choose **Tenants** under the **User Access** submenu. This opens the List/Add Tenants page.
- Step 2** Click the **Add Tenants** icon in the Tenants pane, enter the tenant tag and tenant ID and click **Add Tenant**. The Name and Description attributes are optional.
- Note** You cannot create more than one tenant with the same tenant ID or tenant tag.
- Step 3** Click **Save**.
- The **Settings** drop-down list on the toolbar at the top right of the page will display the tenant under the **Tenants** submenu. You can use this drop-down list to select a tenant when you have to do tenant specific configurations.
-

CLI Commands

To add a tenant, use **tenant tag create tenant-id [attribute=value]** (see the **tenant** command in the CLIGuide.html file in the /docs directory for syntax and attribute descriptions).

Editing a Tenant

To edit a tenant, do the following:

Local and Regional Advanced Web UI

- Step 1** On the List/Add Tenants page, click the name of the desired tenant in the Tenants pane and the Edit Tenant page appears with the details of the selected tenant.
- Step 2** You can modify the tenant tag, name, or description of the tenant on the Edit Tenant page and click **Save**. The tenant ID cannot be modified.
-

Deleting a Tenant



Warning Deleting the tenant will also delete all data for the tenant.

To delete a tenant, select the name of the desired tenant in the Tenants pane, click the **Delete** icon in the Tenants pane, and then confirm the deletion.



Note A user constrained to a specific tenant cannot delete tenants.

Managing Tenant Data

You can create two types of data for tenants:

- Tenant data, which is assigned to a specified tenant and cannot be viewed by other tenants
- Core data, which is visible to all tenants in read-only mode

Local and Regional Web UI

To create tenant data objects in the Web UI, do the following:

Step 1 To set the data for a desired tenant, click the Settings drop-down list on the toolbar at the top of the page and select the desired tenant under the Tenants submenu.

Step 2 Create the object.

When creating tenant data, most object names are only required to be unique for the specified tenant. For example, tenants *abc* and *xyz* may both use their own scope *test* that is private to their configuration.

Note Administrators (Admin), zones (CCMZone, CCMReverseZone, and CCMSecondaryZone), keys (Key), and clients (ClientEntry) must be unique across all tenants.

Administrator names must be unique to perform initial login authentication and establish whether the user is a tenant. Zone and key classes must be unique because these require a DNS domain name that is expected to be unique across the Internet. Client names must correspond to a unique client identifier that the DHCP server can use to match its incoming requests.

Local and Regional Web UI

To create core data objects in the web UI, do the following:

Step 1 Ensure that you select [**all**] from the Settings drop-down list on toolbar at the top of the page and select the desired tenant under the Tenants submenu.

Step 2 Create the object, leaving the object tenant assignment set to **none**. By default **none** is selected in the Tenant drop-down list. Leave it as it is, so that the object is not constrained to any specific tenant.

Core data can be used to provide common configuration elements such as policies or client classes that you choose to offer to tenants. Tenants can view and reference these objects in their configuration, but cannot change or delete them. Because core data is visible to all tenants, objects names must be unique across all tenants.

CLI Commands

Use **session set tenant=tag** to set the selected tenant. Use **session unset tenant** to clear the tenant selection, if set (see the **session** command in the CLIGuide.html file in the /docs directory for syntax and attribute descriptions).



Note Once created, you cannot change the tenant or core designation for the object. You must delete and recreate the object to change its tenant assignment.



Tip You can use the `cnr_exim` tool to move a set of tenant data from one tenant to another.

Assigning a Local Cluster to a Single Tenant

When assigned to a single tenant, core data on the local cluster is not restricted to read-only access. This means tenants may be given the ability to stop and start servers, modify defaults, and install custom extensions. After the cluster is assigned to a specific tenant, other tenants cannot log into the cluster.



Note If synchronization with the local cluster fails, the cluster will not be assigned to the tenant. Resolve any connectivity issues and use the resynchronization icon to set the local cluster tenant.

Regional Web UI

To assign a local cluster to a single tenant, do the following:

-
- Step 1** Add the tenant in the List/Add Tenant page if you want to assign the cluster to a new tenant (see the [Adding a Tenant, on page 47](#)).
 - Step 2** From the **Operate** menu, Choose **Manage Clusters** under the **Servers** submenu. The List/Add Clusters page is displayed.
 - Step 3** Choose the tenant you added in **Step 1** from the Settings drop-down list on the toolbar at the top of the page and select the desired tenant under the Tenants submenu.
 - Step 4** Click the **Add Manage Clusters** icon in the **Manage Clusters** pane. The Add Cluster dialog box appears.
 - Step 5** Click **Add Cluster** to add the cluster. For information on adding the cluster, see the [Create the Local Clusters, on page 119](#).
- Note** Once a cluster is assigned to a particular tenant, it cannot be changed or unset.
-

Pushing and Pulling Tenant Data

In the regional web UI, list pages include push options that let you distribute objects to a list of local clusters, and pull options that let you merge local cluster objects from the Replica data into the central configuration. These operations can be performed on both tenant and core data, but only one set of data can be pushed or pulled in a single operation.

Use the Settings drop-down list on the toolbar at the top of the page and select the desired tenant under the Tenants submenu to specify the set of data to be pushed or pulled.



Note To maintain a consistent view of tenant data, all related clusters should be configured with the same list of tenants. See [Pushing and Pulling Tenants, on page 66](#) for steps that help you manage tenant lists.

Assigning Tenants When Using External Authentication

When external RADIUS authentication is configured, the groups that are assigned in the RADIUS server configuration establish the access privileges of the user. The implicit group name `ccm-tenant-tag` or `ccm-tenant-id` must be added to the list of groups of tenant user to designate the tenant status. Other assigned groups must be core groups or groups assigned to the same tenant. Invalid groups will be ignored when building user credentials at login.

For example, to assign superuser access for the tenant `abc`, specify the groups attribute as:

```
cnr:groups=superusers,ccm-tenant-abc
```

See [External Authentication Servers, on page 42](#).

Using `cnr_exim` With Tenant Data

The `cnr_exim` tool lets you export tenant data, and optionally re-assign the data to a different tenant on import (See the [Using the `cnr_exim` Data Import and Export Tool, on page 170](#)). You can use these features to:

- Create a standard set of objects for each tenant
- Move tenant data to a new tenant



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

Creating a Standard Set of Tenant Objects

You can use a standard set of tenant objects to provide common objects such as scope and zone templates, policies, and client classes. You can use these instead of core data objects to give tenants the option to customize their settings.

To create a standard set of tenant objects, do the following:

Step 1 Create a template tenant user to use as a placeholder, with `tag=template` and `id=9999`, and create the set of objects to be reused for each tenant.

Step 2 Use the `cnr_exim` tool to export the template configuration:

```
cnr_exim -f template -x -e template.bin
```

Step 3 Use the `cnr_exim` tool to import the template configuration for the tenant `abc` :

```
cnr_exim -f template -g abc -i template.bin
```

Note The template tenant user does not need to be present on the cluster to import the data, which lets you reuse the `template.bin` export file on other clusters. Once you have created the export file, you can also delete the placeholder tenant on the original cluster to remove all associated template data, if desired.

Moving Tenant Data

The ID of a tenant can only be changed by deleting and re-creating the tenant. To retain the data of the tenant when this is required, do the following (assuming the tenant tag for the tenant is `xyz`):

Step 1 Use the `cnr_exim` tool to export the configuration for the tenant `xyz`:

```
cnr_exim -f xyz -x -e xyz.bin
```

Step 2 Delete the tenant `xyz`.

Step 3 Recreate the tenant with the corrected tenant id.

Step 4 Use the `cnr_exim` tool to re-import the configuration:

```
cnr_exim -f xyz -g xyz -i xyz.bin
```

Managing Administrators

When you first log in, Cisco Prime Network Registrar will have one administrator—the superuser account. This superuser can exercise all the functions of the web UI and usually adds the other key administrators. However, `ccm-admin` and `regional-admin` administrators can also add, edit, and delete administrators. Creating an administrator requires:

- Adding its name.
- Adding a password.
- Specifying if the administrator should have superuser privileges (usually assigned on an extremely limited basis).
- If not creating a superuser, specifying the group or groups to which the administrator should belong. These groups should have the appropriate role (and possibly subrole) assignments, thereby setting the proper constraints.



Tip If you accidentally delete all the roles by which you can log into Cisco Prime Network Registrar (those having superuser, `ccm-admin`, or `regional-admin` privileges), you can recover by creating a username/password pair in the `install-path/conf/priv/local.superusers` file. You must create this file, have write access to it, and include a line in it with the format `username password`. Use this username and password for the next login session. Note, however, that using the `local.superusers` file causes reduced security. Therefore, use this file only in emergencies such as when temporarily losing all login access. After you log in, create a superuser account in the usual way, then delete the `local.superusers` file or its contents. You must create a new administrator account for each individual, to track administrative changes.

Adding Administrators

To add an administrator, do the following:

Local and Regional Web UI

Step 1 From the **Administration** menu, choose **Administrators** under the **User Access** submenu. This opens the List/Add Administrators page (see the [Create the Administrators, on page 112](#) for an example).

Step 2 Click the **Add Administrators** icon in the **Administrators** pane, enter the name in the Name field, enter the password in the Password field, retype the password in the Confirm Password field in the Add Admin dialog box, and then click **Add Admin**.

- Step 3** Choose one or more existing groups from the Groups Available list (or whether the administrator should be a superuser) and then click **Save**.
-

Editing Administrators

To edit an administrator, select the administrator in the Administrators pane, modify the name, password, superuser status, or group membership on the Edit Administrator page, and then click **Save**. The active group or groups should be in the Selected list.

Deleting Administrators

To delete an administrator, select the administrator in the Administrators pane, click the **Delete Administrators** icon, and then confirm or cancel the deletion.

Managing Passwords

Passwords are key to administrator access to the web UI and CLI. In the web UI, you enter the password on the Login page. In the CLI, you enter the password when you first invoke the **nrcmd** program. The local or regional CCM administrator or superuser can change any administrator password.

You can prevent exposing a password on entry. In the web UI, logging in or adding a password never exposes it on the page, except as asterisks. In the CLI, you can prevent exposing the password by creating an administrator, omitting the password, then using **admin name enterPassword**, where the prompt displays the password as asterisks. You can do this instead of the usual **admin name set password** command that exposes the password as plain text.

Administrators can change their own passwords on clusters. If you want the password change propagated from the regional server to all local clusters, log into the regional cluster. First ensure that your session `admin-edit-mode` is set to synchronous, and then update your password.



Note The password should not be more than 255 characters long.

Managing Groups

A superuser, `ccm-admin`, or `regional-admin` can create, edit, and delete administrator groups. Creating an administrator group involves:

- Adding its name.
- Adding an optional description.
- Choosing associated roles.

Adding Groups

To add a group, do the following:

Local Advanced and Regional Web UI

- Step 1** From the **Administration** menu, choose **Groups** under the **User Access** submenu. This opens the List/Add Administrator Groups page (see the [Create a Group to Assign to the Host Administrator, on page 116](#) for an example).
- Step 2** Click the **Add Groups** icon in the Groups pane, enter a name and an optional description in the Add CCMAdminGroup dialog box, and then click **Add CCMAdminGroup**.
- Step 3** Choose one or more existing roles from the **Roles Available** list and then click **Save**.
-

Editing Groups

To edit a group, click the name of the group that you want to edit in the Groups pane to open the Edit Administrator Group page. You can modify the name, description, or role membership in this page. You can view the active roles in the Selected list.

Deleting Groups

To delete a group, select the group in the Groups pane, click the **Delete Groups** icon, and then confirm the deletion. Click **Cancel** in the confirmation window to cancel the deletion.

Managing Roles

A superuser, ccm-admin, or regional-admin administrator can create, edit, and delete administrator roles. Creating an administrator role involves:

- Adding its name.
- Choosing a base role.
- Possibly specifying if the role should be unconstrained, or read-only.
- Possibly adding constraints.
- Possibly assigning groups.

Adding Roles

To add a role, do the following:

Local and Regional Advanced Web UI

- Step 1** From the **Administration** menu, choose **Roles** under the **User Access** submenu. This opens the List/Add Administrator Roles page.
- Step 2** Click the **Add Role** icon in the Roles pane, enter a name, and choose a tenant and a base role in the Add Roles dialog box, and then click **Add Role**.
- Step 3** On the List/Add Administrator Roles page, specify any role constraints, subrole restrictions, or group selections, then click **Save**.
-

Editing Roles

To edit a role, select the role in the Roles pane, then modify the name or any constraints, subrole restrictions, or group selections on the Edit Administrator Role page. The active subroles or groups should be in the Selected list. Click **Save**.

Deleting Roles

To delete a role, select the role in the Roles pane, click the **Delete Role** icon, and then confirm the deletion.



Note You cannot delete the default roles.

CLI Commands

To add and edit administrator roles, use **role name create base-role [attribute=value]** (see the **role** command in the CLIGuide.html file in the /docs directory for syntax and attribute descriptions). The base roles have default groups associated with them. To add other groups, set the *groups* attribute (a comma-separated string value).

Granular Administration

Granular administration prevents unauthorized users from accidentally making a change on zones, address blocks, subnets, and router interfaces. It also ensures that only authorized users view or modify specific scopes, prefixes, and links. Granular administration constraints administrators to specific set of scopes, prefixes, and links. A constrained administrator can view or make changes to authorized scope, prefix, and link objects only. The CCM server uses owner and region constraints to authorize and filter IPv4 address space objects, and DNS zone related objects (CCMZone, CCMReverseZone, CCMSecondaryZone, CCMRRSet, and CCMHost). The zones are constrained by owners and regions. Owner or region attributes on the CCMSubnet control access to scopes. Also, owner or region attributes on the Prefix and Link objects control access to prefixes and links.

Local Advanced and Regional Web UI

-
- Step 1** From the **Administration** menu, choose **Roles** to open the List/Add Administrator Roles page.
 - Step 2** Click the **Add Role** icon in the Roles pane, enter a name for the custom role, for example, my-dhcp, choose a tenant, and choose **dhcp-admin** from the Role drop-down list and click **Add Role**.
 - Step 3** Click **True** or **False** radio button as necessary, on the Add DHCP Administrator Role page.
 - Step 4** Choose the required sub roles in the Available field and move them to the Selected field.
 - Step 5** Click **Add Constraint**.
 - a) On the Add Role Constraint page, modify the fields as necessary.
 - b) Click **Add Constraint**. The constraint must have an index number of 1.
 - Step 6** Click **Save**.

The name of the custom role appears on the list of roles in the List/Add Administrator Roles page.

Related Topics

[Scope-Level Constraints, on page 55](#)

[Prefix-Level Constraints, on page 56](#)

[Link-Level Constraints, on page 57](#)

Scope-Level Constraints

A dhcp admin user can view or modify a scope if any of the following conditions is met:

- Owner of the subnet for the scope matches the dhcp-admin owner.
- Region of the subnet for the scope matches the region role constraints.
- Owner or region of the parent address block matches the dhcp-admin owner or region role constraints. Note that the most immediate parent address block that has owner or region defined takes precedence.

The following conditions are also valid:

- If the matching owner or region constraint is marked as read-only, you can only view the scope.
- If a scope has a primary network defined, the primary subnet and its parent address block owner or region constraints override secondary subnets.
- If no parent subnet or address block defines owner or region constraints, then you can access the scope.
- If you are an unconstrained dhcp-admin user, you can have access to all scopes.



Note These hierarchical authorization checks for dhcp-admin owner/region constraints are applicable to scopes, subnets, and parent address blocks. Identical hierarchical authorization checks for addrblock-admin owner/region constraints apply to address blocks and subnets. If you have dhcp-admin and the addrblock-admin privileges, you can access address blocks and subnets, if either of the roles allow access.

Examples of Scope-Level Constraints:

```
Parent CCMAAddrBlock 10.0.0.0/8 has owner 'blue' set.
  Scope 'A' has subnet 10.0.0.0/24 has parent CCMSubnet with owner 'red'.
  Scope 'B' has subnet 10.0.1.0/24 has parent CCMSubnet with no owner set.
  Scope 'C' has subnet 10.10.0.0/24 has parent CCMSubnet with owner 'green' and
primary-subnet 10.0.0.0/24.
  Scope 'D' has subnet 100.10.0.0/24 has parent CCMSubnet with owner unset, and no parent
block.

Scope 'A' owner is 'red'.
Scope 'B' owner is 'blue'.
```

```
Scope 'C' owner is 'red'.
Scope 'D' owner is unset. Only unconstrained users can access this scope.
```

Local Advanced Web UI

To add scopes, do the following:

-
- Step 1** From the **Design** menu, choose **Scopes** under the **DHCPv4** submenu to open the List/Add DHCP Scopes.
 - Step 2** Click the **Add Scopes** icon in the Scopes pane, enter a name, subnet, primary subnet, choose policy, enter a selection-tag-list, and select the scope template in the Add DHCP Scope dialog box.
 - Step 3** Click **Add DHCP Scope**. The List/Add DHCP Scopes page appears.
 - Step 4** Enter values for the fields or attributes as necessary.
 - Step 5** To unset any attribute value, check the check box in the **Unset?** column, then click **Unset Fields** at the bottom of the page.
 - Step 6** Click **Save** to add scope or **Revert** to cancel the changes.

Tip If you add new scope values or edit existing ones, click **Save** to save the scope object.

Prefix-Level Constraints

You can view or modify a prefix, if you have either of the following:

- The ipv6-management subrole of the dhcp-admin, or addrblock-admin role on the local cluster.
- The central-cfg-admin, or regional-addr-admin role on the regional cluster.

You can view or modify a prefix if any of the following conditions is true:

- The owner or region of the parent link matches the owner or region role constraints defined for you.
- The owner or region of this prefix matches the owner or region role constraints defined for you.
- The owner or region of the parent prefix matches the owner or region role constraints defined for you.

You can view or modify a prefix if any of the following conditions is true:

- If the matching owner or region constraint for you is marked as read-only, then you can only view the prefix.
- If the prefix references a parent link, the link owner or region constraints is applicable if the link owner or region constraints set.
- If no parent link or prefix defines any owner or region constraints, then you can access this prefix only if owner or region role constraints are not defined for you.
- If you are an unconstrained user, then you have access to all.

Examples of Prefix-Level constraints:

```
Link 'BLUE' has owner 'blue' set.
Parent Prefix 'GREEN' has owner 'green' set.
Prefix 'A' has owner 'red' set, no parent prefix, and no parent link.
Prefix 'B' has owner 'yellow' set, parent Prefix 'GREEN' and parent link 'BLUE'.
```

```
Prefix 'C' has no owner set, parent prefix 'GREEN', and no parent link.  
Prefix 'C' has no owner set, no parent prefix, and no parent link.
```

```
Prefix 'A' owner is 'red'.  
Prefix 'B' owner is 'blue'.  
Prefix 'C' owner is 'green'.  
Prefix 'D' owner is unset. Only unconstrained users can access this prefix.
```

Local Advanced and Regional Web UI

To view unified v6 address space, do the following:

-
- Step 1** From the **Design** menu, choose **Address Tree** under the **DHCPv6** submenu to open the DHCP v6 Address Tree page.
 - Step 2** View a prefix by adding its name, address, and range, then choosing a DHCP type and possible template (see the *"Viewing IPv6 Address Space"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*).
 - Step 3** Choose the owner from the owner drop-down list.
 - Step 4** Choose the region from the region drop-down list.
 - Step 5** Click **Add Prefix**. The newly added Prefix appears on the DHCP v6 Address Tree page.
-

Local Advanced and Regional Web UI

To list or add DHCP prefixes, 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, enter a name, address, and range for the prefix, then choose the DHCP type and possible template.
 - Step 3** Choose the owner from the owner drop-down list.
 - Step 4** Choose the region from the region drop-down list.
 - Step 5** Click **Add IPv6 Prefix**. The newly added Prefix appears on the List Prefixes page.
-

Link-Level Constraints

You can view or modify a link if:

- You are authorized for the ipv6-management subrole of the dhcp-admin or addrblock-admin role on the local cluster, or the central-cfg-admin or regional-addr-admin role on the regional cluster.
- The owner or region of the link matches the owner or region role constraints defined for you.
- No owner or region is defined for the link, and only if no owner or region role constraints are defined for you.

If you are an unconstrained user, then you have access to all links.

The following is an example of Link Level Constraints:

```
Link 'BLUE' has owner 'blue' set.  
Link 'ORANGE' has owner unset.
```

```
Link 'BLUE' owner is 'blue'.
Link 'ORANGE' owner is unset. Only unconstrained users can access this link.
```

Local Advanced and Regional Web UI

To add links, do the following:

-
- Step 1** From the **Design** menu, choose **Links** under the **DHCPv6** submenu to open the List/Add DHCP v6 Links page.
 - Step 2** Click the **Add Links** icon in the Links pane, enter a name, then choose the link type, and enter a group.
 - Step 3** Click **Add Link**. The newly added DHCPv6 Link appears on the List/Add DHCP v6 Links page.
-

Centrally Managing Administrators

As a regional or local CCM administrator, you can:

- Create and modify local and regional cluster administrators, groups, and roles.
- Push administrators, groups, and roles to local clusters.
- Pull local cluster administrators, groups, and roles to the central cluster.

Each of these functions involves having at least one regional CCM administrator subrole defined. The following table describes the subroles required for these operations.

Table 6: Subroles Required for Central Administrator Management

Central Administrator Management Action	Required Regional Subroles
Create, modify, push, pull, or delete administrators	authentication
Create, modify, push, pull, or delete groups or roles	authorization
Create, modify, push, pull, or delete groups or roles with associated owners or regions	authorization owner-region
Create, modify, push, pull, or delete external authentication servers	authentication
Create, modify, push, pull, or delete tenants	authentication

Related Topics

[Pushing and Pulling Administrators, on page 59](#)

[Pushing and Pulling Groups, on page 63](#)

[Pushing and Pulling Roles, on page 64](#)

[Pushing and Pulling Tenants, on page 66](#)

Pushing and Pulling Administrators

You can push administrators to, and pull administrators from local clusters on the List/Add Administrators page in the regional cluster web UI.

You can create administrators with both local and regional roles at the regional cluster. However, you can push or pull only associated local roles, because local clusters do not recognize regional roles.

Related Topics

[Pushing Administrators to Local Clusters, on page 59](#)

[Pushing Administrators Automatically to Local Clusters , on page 59](#)

[Pulling Administrators from the Replica Database, on page 60](#)

Pushing Administrators to Local Clusters

Pushing administrators to local clusters involves choosing one or more clusters and a push mode.

Regional Basic and Advanced Web UI

-
- Step 1** From the **Administration** menu, choose **Administrators**.
 - Step 2** On the List/Add Administrators Page, click the **Push All** icon in the **Administrators** pane to push all the administrators listed on the page. This opens the Push Data to Local Clusters dialog box.
 - Step 3** Choose a push mode by clicking one of the Data Synchronization Mode radio buttons. If you are pushing all the administrators, you can choose Ensure, Replace, or Exact. If you are pushing a single administrator, you can choose Ensure or Replace. In both cases, Ensure is the default mode. You would choose Replace only if you want to replace the existing administrator data at the local cluster. You would choose Exact only if you want to create an exact copy of the administrator database at the local cluster, thereby deleting all administrators 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**.
 - Step 6** On the View Push Data Report dialog box, view the push details, then click **OK** to return to the List/Add Administrators page.
-

Pushing Administrators Automatically to Local Clusters

You can automatically push the new user name and password changes from the regional cluster to the local cluster. To do this, you must enable the synchronous edit mode in the regional cluster. The edit mode is set for the current Web UI session, or set as default for all users is set in the CCM Server configuration.

When synchronous mode is set, all the subsequent changes to user name and password are synchronized with local clusters. You can modify your password on the regional server, and this change is automatically propagated to local clusters.

If you are an admin user, you can make multiple changes to the user credentials on the regional cluster. All these changes are automatically pushed to local clusters.

Regional Basic and Advanced Web UI

-
- Step 1** From the **Operate** menu, choose **Manage Servers** under Servers submenu to open the Manage Servers page.
- Step 2** Click the **Local CCM Server** link on the Manage Servers pane to open the Edit CCM Server page.
- Step 3** Choose the synchronous radio buttons for the regional edit mode values for admin, dhcp, and dns.
- Step 4** Choose the webui mode value from the webui-mode drop-down list.
- Step 5** Enter the idle-timeout value.
- Step 6** To unset any attribute value, check the check box in the Unset? column, then click **Unset Fields** at the bottom of the page. To unset the attribute value or to change it, click **Save**, or **Cancel** to cancel the changes.
- Note** Enter values for the attributes marked with asterisks because they are required for CCM server operation. You can click the name of any attribute to open a description window for the attribute.
-

Connecting to CLI in Regional Mode

You must connect to the CLI in Regional Mode. The -R flag is required for regional mode. To set the synchronous edit mode:

```
nrcmd-R> session set admin-edit-mode=synchronous
```

Pulling Administrators from the Replica Database

Pulling administrators from the local clusters is mainly useful only in creating an initial list of administrators that can then be pushed to other local clusters. The local administrators are not effective at the regional cluster itself, because these administrators do not have regional roles assigned to them.

When you pull an administrator, you are actually pulling it from the regional cluster replica database. Creating the local cluster initially replicates the data, and periodic polling automatically updates the replication. However, to ensure that the replica data is absolutely current with the local cluster, you can force an update before pulling the data.

Regional Basic and Advanced Web UI

-
- Step 1** From the **Administration** menu, choose **Administrators** under the **User Access** submenu.
- Step 2** On the List/Add Administrators page, click **Pull Data** on the **Administrators** pane. This opens the Select Replica Admin Data to Pull dialog box.
- 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, on page 81](#).)
- Step 4** Choose a replication mode using one of the Mode radio buttons. In most cases, you would leave the default Replace mode enabled, unless you want to preserve any existing administrator properties already defined at the regional cluster by choosing Ensure, or create an exact copy of the administrator database at the local cluster by choosing Exact (not recommended).
- Step 5** Click **Pull Core Administrators** next to the cluster, or expand the cluster name and click **Pull Administrator** to pull an individual administrator in the cluster.
- Step 6** On the Select Replica Admin Data to Pull dialog box, view the change set data, then click **OK**. You return to the List/Add Administrators page with the pulled administrators added to the list.

Note If you do not have a regional cluster and would like to copy administrators, roles, or groups from one local cluster to another, you can export them and then reimport them at the target cluster by using the `cnr_exim` tool (see the [Using the cnr_exim Data Import and Export Tool, on page 170](#)). However, the tool does not preserve the administrator passwords, and you must manually reset them at the target cluster. It is implemented this way to maintain password security. The export command is:

```
cnr_exim -c admin -x -e outputfile.txt
```

Pushing and Pulling External Authentication Servers

You can push all external authentication servers to local cluster or pull the external authentication server data from the local cluster on the List/Add RADIUS Server page or List/Add Active Directory Server page in the regional web UI.

Pushing RADIUS External Authentication Servers

To push external authentication servers to the local cluster, do the following:

Regional Advanced Web UI

-
- Step 1** From the **Administration** menu, choose **Radius** under the **External Authentication** submenu to view the List/Add RADIUS Server page in the regional web UI.
- Step 2** Click **Push All** icon in the Radius pane to push all the external authentication servers listed on the page, or **Push** to push an individual external authentication server. This opens the Push Data to Local Clusters dialog box.
- Step 3** Choose a push mode using one of the Data Synchronization Mode radio buttons.
- If you are pushing all the external authentication servers, you can choose Ensure, Replace, or Exact.
 - If you are pushing a single external authentication server, 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 external authentication server data at the local cluster. Choose Exact only if you want to create an exact copy of the external authentication server data at the local cluster, thereby deleting all external authentication servers that are not defined at the regional cluster.
- Step 4** Click **Push Data to Clusters**.

Pulling RADIUS External Authentication Servers

To pull the external authentication server data from the local cluster, do the following:

Regional Advanced Web UI

-
- Step 1** From the **Administration** menu, choose **Radius** under the **External Authentication** submenu to view the List/Add Radius Server page in the regional web UI.
- Step 2** On the List/Add Radius Server page, click **Pull Data** on the **Radius** pane. This opens the Select Replica External Authentication Server Data to Pull dialog box.

Pushing AD External Authentication Servers

- 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, on page 81.](#))
- Step 4** Choose a replication mode using one of the Mode radio buttons.
- Leave the default Replace mode enabled, unless you want to preserve any existing external authentication server properties at the local cluster by choosing Ensure.
- Note** We do not recommend that you create an exact copy of the external authentication server data at the local cluster by choosing Exact.
- Step 5** Click **Pull All External Authentication Servers** next to the cluster.
- Step 6** On the Report Pull Replica Authentication servers page, view the pull details, then click **Run**.
- On the Run Pull Replica Authentication servers page, view the change set data, then click **OK**. You return to the List/Add Authentication Server page with the pulled external authentication servers added to the list.

Pushing AD External Authentication Servers

To push external authentication servers to the local cluster, do the following:

Regional Advanced Web UI

- Step 1** From the **Administration** menu, choose **Active Directory** under the **External Authentication** submenu to view the List/Add Active Directory Server page in the regional web UI.
- Step 2** Click **Push All** on the **Active Directory** pane to push the external authentication server. This opens the Push Data to Local Clusters dialog box.
- Step 3** Choose a push mode using one of the Data Synchronization Mode radio buttons.
- If you are pushing all the external authentication servers, you can choose Ensure, Replace, or Exact.
 - If you are pushing a single external authentication server, 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 external authentication server data at the local cluster. Choose Exact only if you want to create an exact copy of the external authentication server data at the local cluster, thereby deleting all external authentication servers that are not defined at the regional cluster.
- Step 4** Click **Push Data to Clusters**.

Pulling AD External Authentication Servers

To pull the AD external authentication server data from the local cluster, do the following:

Regional Advanced Web UI

- Step 1** From the **Administration** menu, choose **Active Directory** under the **External Authentication** submenu to view the List/Add Active Directory Server page in the regional web UI.
- Step 2** On the List/Add Active Directory Server page, click **Pull Data** on the **Active Directory** pane. This opens the Select Replica External Authentication Server Data to Pull dialog box.

- 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, on page 81](#)).
- Step 4** Choose a replication mode using one of the Mode radio buttons.
- Leave the default Replace mode enabled, unless you want to preserve any existing external authentication server properties at the local cluster by choosing Ensure.
- Note** We do not recommend that you create an exact copy of the external authentication server data at the local cluster by choosing Exact.
- Step 5** Click **Pull All External Authentication Servers** next to the cluster.
- Step 6** On the Report Pull Replica Authentication servers page, view the pull details, and then click **Run**.
- On the Run Pull Replica Authentication servers page, view the change set data, and then click **OK**. You return to the List/Add Authentication Server page with the pulled external authentication servers added to the list.
-

Pushing and Pulling Groups

Pushing and pulling groups is vital in associating administrators with a consistent set of roles at the local clusters. You can push groups to, and pull groups from, local clusters on the List/Add Administrator Groups page in the regional cluster web UI.

Related Topics

[Pushing Groups to Local Clusters, on page 63](#)

[Pulling Groups from the Replica Database, on page 64](#)

Pushing Groups to Local Clusters

Pushing groups to local clusters involves choosing one or more clusters and a push mode.

Regional Basic and Advanced Web UI

- Step 1** From the **Administration** menu, choose **Groups** under the **User Access** submenu.
- Step 2** On the List/Add Administrator Groups page, click the **Push All** icon on Groups pane to push all the groups listed on the page, or **Push** to push an individual group. This opens the Push Data to Local Clusters dialog box.
- Step 3** Choose a push mode using one of the Data Synchronization Mode radio buttons. If you are pushing all the groups, you can choose Ensure, Replace, or Exact. If you are pushing a single group, you can choose Ensure or Replace. In both cases, Ensure is the default mode. You would choose Replace only if you want to replace the existing group data at the local cluster. You would choose Exact only if you want to create an exact copy of the group data at the local cluster, thereby deleting all groups that are not defined at the regional cluster.
- Step 4** By default, the associated roles and owners are pushed along with the group. Roles are pushed in Replace mode and owners in Ensure mode. To disable pushing the associated roles or owners, uncheck the respective check box.
- Step 5** Choose one or more local clusters in the Available field of the Destination Clusters and move it or them to the Selected field.
- Step 6** Click **Push Data to Clusters**.

- Step 7** On the View Push Group Data Report page, view the push details, then click **OK** to return to the List/Add Administrator Groups page.
-

Pulling Groups from the Replica Database

Pulling administrator groups from the local clusters is mainly useful only in creating an initial list of groups that can then be pushed to other local clusters. The local groups are not useful at the regional cluster itself, because these groups do not have regional roles assigned to them.

When you pull a group, you are actually pulling it from the regional cluster replica database. Creating the local cluster initially replicates the data, and periodic polling automatically updates the replication. However, to ensure that the replica data is absolutely current with the local cluster, you can force an update before pulling the data.

Regional Basic and Advanced Web UI

- Step 1** From the **Administration** menu, choose **Groups** under the **User Access** submenu.
- Step 2** On the List/Add Administrator Groups page, click the **Pull Data** icon on the **Groups** pane. This opens the Select Replica CCMAAdminGroup Data to Pull dialog box.
- 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, on page 81](#).)
- Step 4** Choose a replication mode using one of the Mode radio buttons. In most cases, you would leave the default Replace mode enabled, unless you want to preserve any existing group properties at the local cluster by choosing Ensure, or create an exact copy of the group data at the local cluster by choosing Exact (not recommended).
- Step 5** Click **Pull Core Groups** next to the cluster, or expand the cluster name and click **Pull Group** to pull an individual group in the cluster.
- Step 6** On the Report Pull Replica Groups page, view the pull details, then click **Run**.
- Step 7** On the Run Pull Replica Groups page, view the change set data, then click **OK**. You return to the List/Add Administrator Groups page with the pulled groups added to the list.
-

Pushing and Pulling Roles

You can push roles to, and pull roles from, local clusters on the List/Add Administrator Roles page in the regional cluster web UI. You can also push associated groups and owners, and pull associated owners, depending on your subrole permissions (see [Table 6: Subroles Required for Central Administrator Management, on page 58](#)).

Related Topics

[Pushing Roles to Local Clusters, on page 64](#)

[Pulling Roles from the Replica Database, on page 65](#)

Pushing Roles to Local Clusters

Pushing administrator roles to local clusters involves choosing one or more clusters and a push mode.

Regional Advanced Web UI

-
- Step 1** From the **Administration** menu, choose **Roles** under the **User Access** submenu.
- Step 2** On the List/Add Administrator Roles page, click the **Push All** icon in the Roles pane to push all the roles listed on the page, or **Push** to push an individual role. This opens the Push Data to Local Clusters dialog box.
- Step 3** Choose a push mode using one of the Data Synchronization Mode radio buttons. If you are pushing all the roles, you can choose Ensure, Replace, or Exact. If you are pushing a single role, you can choose Ensure or Replace. In both cases, Ensure is the default mode. You would choose Replace only if you want to replace the existing role data at the local cluster. You would choose Exact only if you want to create an exact copy of the role data at the local cluster, thereby deleting all roles that are not defined at the regional cluster.
- Step 4** By default, the associated groups and owners are pushed along with the role. Groups are pushed in Replace mode and owners in Ensure mode. To disable pushing the associated roles or owners, uncheck the respective check box:
- If you disable pushing associated groups and the group does not exist at the local cluster, a group based on the name of the role is created at the local cluster.
 - If you disable pushing associated owners and the owner does not exist at the local cluster, the role will not be configured with its intended constraints. You must separately push the group to the local cluster, or ensure that the regional administrator assigned the owner-region subrole has pushed the group before pushing the role.
- Step 5** Choose one or more local clusters in the Available field of the Destination Clusters and move it or them to the Selected field.
- Step 6** Click **Push Data to Clusters**.
- Step 7** On the View Push Role Data Report page, view the push details, then click **OK** to return to the List/Add Administrator Roles page.
-

Pulling Roles from the Replica Database

Pulling administrator roles from the local clusters is mainly useful only in creating an initial list of roles that can then be pushed to other local clusters. The local roles are not useful at the regional cluster itself.

When you pull a role, you are actually pulling it from the regional cluster replica database. Creating the local cluster initially replicates the data, and periodic polling automatically updates the replication. However, to ensure that the replica data is absolutely current with the local cluster, you can force an update before pulling the data.

Regional Advanced Web UI

-
- Step 1** From the **Administration** menu, choose **Roles** under the **User Access** submenu.
- Step 2** On the List/Add Administrator Roles page, click the **Pull Data** icon in the **Roles** pane. This opens the Select Replica Administrator Role Data to Pull dialog box.
- 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, on page 81](#).)
- Step 4** Choose a replication mode using one of the Mode radio buttons. In most cases, you would leave the default Replace mode enabled, unless you want to preserve any existing role properties at the local cluster by choosing Ensure, or create an exact copy of the role data at the local cluster by choosing Exact (not recommended).
- Step 5** If you have the owner-region subrole permission, you can decide if you want to pull all the associated owners with the role, which is always in Ensure mode. This choice is enabled by default.

- Step 6** Click **Pull Core Roles** next to the cluster, or expand the cluster name and click **Pull Role** to pull an individual role in the cluster.
- Step 7** On the Report Pull Replica Roles page, view the pull details, then click **Run**.
- Step 8** On the Run Pull Replica Roles page, view the change set data, then click **OK**. You return to the List/Add Administrator Roles page with the pulled roles added to the list.
-

Pushing and Pulling Tenants

You can push all tenants to local cluster or pull the tenants data from the local cluster on the List/Add Tenants Page in the regional web UI.

Pushing Tenants to Local Clusters

To push tenants to the local cluster, do the following:

Regional Basic and Advanced Web UI

To add scopes, do the following:

-
- Step 1** From the **Administration** menu, choose **Tenants** under the **User Access** submenu to view the List/Add Tenants page in the regional web UI.
- Step 2** Click the **Push All** icon in the **Tenants** pane to push all the tenants listed on the page, or **Push** to push an individual tenant. This opens the Push Tenant Data to Local Clusters page.
- Step 3** Choose a push mode using one of the Data Synchronization Mode radio buttons.
- If you are pushing all the tenant, you can choose Ensure, Replace, or Exact.
 - If you are pushing a single tenant, you can choose Ensure or Replace.
- In both cases, Ensure is the default mode.
- Choose Replace only if you want to replace the tenant data at the local cluster. Choose Exact only if you want to create an exact copy of the tenant data at the local cluster, thereby deleting all tenants that are not defined at the regional cluster.
- Step 4** Click **Push Data to Clusters**.
-

Pulling Tenants from the Replica Database

To pull tenants from the replica database, do the following:

Regional Basic and Advanced Web UI

-
- Step 1** From the **Administration** menu, choose **Tenants** under the **User Access** submenu to view the List/Add Tenants page.
- Step 2** On the List/Add Tenants page, click the **Pull Data** icon in the **Tenants** pane. This opens the Select Replica Tenant Data to Pull dialog box.

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, on page 81.](#))

Step 4 Choose a replication mode using one of the Mode radio buttons.

Leave the default Replace mode enabled, unless you want to preserve any existing tenant data at the local cluster by choosing Ensure.

Note We do not recommend that you create an exact copy of the tenant data at the local cluster by choosing Exact.

Step 5 Click **Pull Replica**.

Step 6 On the Select Replica Tenant Data to Pull page, click **Pull all Tenants** to view the pull details, and then click **Run**.
On the Run Pull Replica Tenants page, view the change set data, then click **OK**. You return to the List/Add Tenants page with the pulled tenants added to the list.



CHAPTER 5

Managing Owners and Regions

This chapter explains how to configure owners and regions that can be applied to DHCP address blocks, subnets, prefixes, links, and zones.

- [Managing Owners, on page 69](#)
- [Managing Regions, on page 70](#)
- [Centrally Managing Owners and Regions, on page 70](#)

Managing Owners

You can create owners to associate with address blocks, subnets, prefixes, links, and zones. You can list and add owners on a single page. Creating an owner involves creating a tag name, full name, and contact name.

Local Advanced and Regional Advanced Web UI

- Step 1** From the **Administration** menu, choose **Owners** under the **Settings** submenu to open the List/Add Owners page. The regional cluster includes pull and push functions also.
- Step 2** Click the **Add Owners** icon in the Owners pane on the left. This opens the Add Owner page.
- Step 3** Enter a unique owner tag.
- Step 4** Enter an owner name.
- Step 5** Enter an optional contact name.
- Step 6** Click **Add Owner**.
- Step 7** To edit an owner, click its name in the Owners pane on the left.
-

CLI Commands

Use **owner tag create name [attribute=value]** to create an owner. For example:

```
nrcmd> owner owner-1 create "First Owner" contact="Contact at owner-1"
```

Managing Regions

You can create regions to associate with address blocks, subnets, prefixes, links, and zones. You can list and add regions on a single page. Creating a region involves creating a tag name, full name, and contact name.

Local Advanced and Regional Advanced Web UI

-
- Step 1** From the **Administration** menu, choose **Regions** under the **Settings** submenu to open the List/Add Regions page. The regional cluster includes pull and push functions also.
- Step 2** Click the **Add Regions** icon in the Regions pane on the left.
- Step 3** Enter a unique region tag.
- Step 4** Enter a region name.
- Step 5** Enter an optional contact name.
- Step 6** Click **Add Region**.
- Step 7** To edit a region, click its name in the Regions pane on the left.
-

CLI Commands

Use `region tag create name [attribute=value]`. For example:

```
nrcmd> region region-1 create "Boston Region" contact="Contact at region-1"
```

Centrally Managing Owners and Regions

As a regional or local CCM administrator, you can:

- Push owners and regions to local clusters.
- Pull local cluster owners and regions to the central cluster.

Each of these functions involves having at least one regional CCM administrator subrole defined (see [Roles, Subroles, and Constraints, on page 39](#)).

The following table describes the subroles required for these operations.

Table 7: Subroles Required for Central Administrator Management

Central Administrator Management Action	Required Regional Subroles
Create, modify, pull, push, or delete owners or regions	owner-region

Related Topics

[Pushing and Pulling Owners or Regions, on page 71](#)

Pushing and Pulling Owners or Regions

You can push owners or regions to, and pull them from, local clusters on the List/Add Owners page or List/Add Regions page, respectively, in the regional cluster web UI.

Related Topics

[Pushing Owners or Regions to Local Clusters, on page 71](#)

[Pulling Owners and Regions from the Replica Database, on page 71](#)

Pushing Owners or Regions to Local Clusters

Pushing owners or regions to local clusters involves choosing one or more clusters and a push mode.

Regional Web UI

- Step 1** From the **Administration** menu, choose **Owners** or **Regions** under the **Settings** submenu.
- Step 2** On the List/Add Owners or List/Add Regions page, click the **Push All** icon in the left pane, or click **Push** at the top of the Edit Owner page or Edit Region page, for a particular owner or region. This opens the Push Owner or Push Region page.
- Step 3** Choose a push mode using one of the Data Synchronization Mode radio buttons.
- If you are pushing all the owners or regions, you can choose Ensure, Replace, or Exact.
 - If you are pushing a single owner or region, 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 owner or region data at the local cluster. Choose Exact only if you want to create an exact copy of the owner or region data at the local cluster, thereby deleting all owners or regions 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**.
- Step 6** On the View Push Owner Data Report or View Push Region Data Report page, view the push details, then click **OK** to return to the List/Add Owners or List/Add Regions page.
-

Pulling Owners and Regions from the Replica Database

When you pull an owner or region, you are actually pulling it from the regional cluster replica database. Creating the local cluster initially replicates the data, and periodic polling automatically updates the replication. However, to ensure that the replica data is current with the local cluster, you can force an update before pulling the data.

Regional Web UI

- Step 1** From the **Administration** menu in the regional cluster web UI, choose **Owners** or **Regions** under the **Settings** submenu.

- Step 2** On the List/Add Owners or List/Add Regions page, click the **Pull Data** icon in the left pane. This opens the Select Replica Owner Data to Pull or Select Replica Region Data to Pull page.
- Step 3** Click the **Replicate** icon in the Update Replica Data column for the cluster. (For the automatic replication interval, see [Replicating Local Cluster Data, on page 81.](#))
- Step 4** Choose a replication mode using one of the Mode radio buttons.
- Leave the default Replace mode enabled, unless you want to preserve any existing owner or region properties at the local cluster by choosing Ensure.
- Note** We do not recommend that you create an exact copy of the owner or region data at the local cluster by choosing Exact.
- Step 5** Click **Pull All Owners** or **Pull All Regions** next to the cluster, or expand the cluster name and click **Pull Owner** or **Pull Region** to pull an individual owner or region in the cluster.
- Step 6** On the Report Pull Replica Owners or Report Pull Replica Regions page, click **Run**.
- Step 7** On the Run Pull Replica Owners or Run Pull Replica Region page, view the change set data, then click **OK**. You return to the List/Add Owners or List/Add Regions page with the pulled owners or regions added to the list.
-



CHAPTER 6

Managing the Central Configuration

This chapter explains how to manage the central configuration at the Cisco Prime Network Registrar regional cluster.

- [Central Configuration Tasks](#), on page 73
- [Default Ports for Cisco Prime Network Registrar Services](#), on page 74
- [Licensing](#), on page 75
- [Configuring Server Clusters](#), on page 78
- [Central Configuration Management Server](#), on page 84
- [Trivial File Transfer](#), on page 84
- [Simple Network Management](#), on page 86
- [Integrating Cisco Prime Network Registrar SNMP into System SNMP](#), on page 95
- [Bring Your Own Device Web Server](#), on page 96
- [Polling Process](#), on page 98
- [Managing DHCP Scope Templates](#), on page 100
- [Managing DHCP Policies](#), on page 102
- [Managing DHCP Client-Classes](#), on page 103
- [Managing Virtual Private Networks](#), on page 105
- [Managing DHCP Failover Pairs](#), on page 106
- [Managing Lease Reservations](#), on page 106
- [Monitoring Resource Limit Alarms](#), on page 108
- [Local Cluster Management Tutorial](#), on page 111
- [Regional Cluster Management Tutorial](#), on page 117

Central Configuration Tasks

Central configuration management at the regional cluster can involve:

- Setting up server clusters, replicating their data, and polling subnet utilization and lease history data from them.
- Setting up routers (see [Managing Routers and Router Interfaces](#), on page 125).
- Managing network objects such as DHCP scope templates, policies, client-classes, options, networks, and virtual private networks (VPNs).
- Managing DHCP failover server pairs.

These functions are available only to administrators assigned the central-cfg-admin role. (The full list of functions for the central-cfg-admin are listed in [Table 5: Regional Cluster Administrator Predefined and Base Roles](#), on page 41.) Note that central configuration management does not involve setting up administrators and checking the status of the regional servers. These functions are performed by the regional administrator, as described in [Licensing](#), on page 75 and [Managing Servers](#), on page 129.

Default Ports for Cisco Prime Network Registrar Services

The following table lists the default ports used for the Cisco Prime Network Registrar services.

Table 8: Default Ports for Cisco Prime Network Registrar Services

Port Number	Protocol	Service
53	TCP/UDP	DNS
53	TCP/UDP	Caching DNS
67	UDP	DHCP client to server
68	UDP	DHCP server to client
69	UDP	TFTP (optional) client to server
80	HTTP	BYOD web server client to server web UI
162	TCP	SNMP traps server to server
389	TCP	DHCP server to LDAP server
443	HTTPS	BYOD web server secure client to server web UI
546	UDP	DHCPv6 server to client
547	UDP	DHCPv6 client to server
647	TCP	DHCP failover server to server
653	TCP	High-Availability (HA) DNS server to server
1234	TCP	Local cluster CCM server to server
1244	TCP	Regional cluster CCM server to server
4444	TCP	SNMP client to server
8080	HTTP	Local cluster client to server web UI

Port Number	Protocol	Service
8090	HTTP	Regional cluster client to server web UI
8443	HTTPS	Local cluster secure client to server web UI
8453	HTTPS	Regional cluster secure client to server web UI

Firewall Considerations

When DNS (caching or authoritative) servers are deployed behind a stateful firewall (whether physical hardware or software, such as contrack), it is recommended that:

- For at least UDP DNS traffic, stateful support be disabled if possible.
- If it is not possible to disable the stateful support, the number of allowed state table entries may need to be significantly increased.

DNS queries typically arrive from many different clients and requests from the same client may use different source ports. With thousands of queries per second, the number of these different sources can be large and if a firewall is using stateful tracking, it has to keep this state and does so for a period of time. Hence, you need to assure that the firewall can hold sufficient state - given the query traffic rates and the state time interval.

Licensing

Cisco Prime Network Registrar requires separate license for CCM, Authoritative DNS, Caching DNS, DHCP, and IPAM services or for combinations of these services. For more details on the Licensing, see the “License Files” section in the Overview chapter of the *Cisco Prime Network Registrar Installation Guide*.

See [Logging In to the Web UIs, on page 11](#) for entering license data the first time you try to log in. You can add the additional service based licenses in the regional server after you log in.

Whenever you log into a regional or local cluster, the overall licensing status of the system is checked. If there are any violations, you will be notified of the violation and the details. This notification is done only once for each user session. In addition, you will be able to see a message on each page indicating the violation.

Regional Web UI

Choose **Licenses** from **Administration > User Access** to open the List/Add Product Licenses page. Click **Browse** to locate the license file, click the file, then click **Open**. If the license ID in the file is valid, the license key appears in the list of licenses with the message “Successfully added license file *filename*.” If the ID is not valid, the License field shows the contents of the file and the message “Object is invalid” appears.

The License Utilization section at the top of the page lists the type of license, the number of nodes allowed for the license, and the actual number of nodes used. Expand the section by clicking the plus (+) sign. The license utilization for each licensed service is listed separately in this section.

The Right To Use and the In Use counts are displayed for each licensed service. The Right To Use value will be the aggregation of the counts across all added licenses for that service. The ‘total in use’ value will be the

aggregation of the latest utilization numbers obtained from all the local clusters. Only the services having a positive Right to use or In Use count will be listed in this section.

Licenses and usage count of earlier versions of Cisco Network Registrar will be listed under a separate section “ip-node”.

The **Expert** mode attribute lets you specify how often license utilization is collected from all the local clusters. Changes to this setting require a server restart to take effect. You can set this attribute at the Edit CCM Server page. The default value is 4 hours.

Adding License

Cisco will e-mail you one or more license files after you register the Cisco Prime Network Registrar Product Authorization Key (PAK) on the web according to the Software License Claim Certificate shipped with the product. Cisco administers licenses through a FLEXlm system.



Note If a license file fails to load, check that the file is properly formatted text file without any extraneous characters in it. Extracting the file from email and moving it between systems can sometimes result in these problems.

Once you have the file or files:

Regional Web UI

- Step 1** Locate the license file or files in a directory (or on the desktop) that is easy to find.
- Step 2** On the List/Add Product Licenses page, browse for each file by clicking the **Choose File** button.

Note The List/Add Product Licenses option is only available at the Regional.

- Step 3** In the Choose file window, find the location of the initial license file, then click **Open**.
- Step 4** If the license key is acceptable, the Add Superuser Administrator page appears immediately.
- Step 5** To add further licenses, from **Administration** menu, choose **Licenses** under the **User Access** submenu to open the List/Add Product Licenses page. Click **Choose File** to locate the additional license file, then click **Open**. If the key in the file is acceptable, the key, type, count, and expiration date appear, along with whether it is an evaluation key. If the key is not acceptable, the page shows the license text along with an error message. For the list of license types, see [Licensing, on page 75](#).

Above the table of licenses is a License Utilization area that, when expanded, shows the license types along with the total nodes that you can use and those actually used.

If Cisco Prime Network Registrar is installed as a distributed system, the license management is done from the regional cluster. You will not have the option of adding licenses in local cluster.

CLI Commands

Use **license file create** to register licenses that are stored in file. The file referenced should include its absolute path or path relative to where you execute the commands. For example:

```
nrcmd-R> license "C:\licenses\product.licenses" create
```

Use **license list** to list the properties of all the created licenses (identified by key), and **license listnames** to list just the keys. Use **license key show** to show the properties of a specific license key.

Registering a Local Cluster that is Behind a NAT

License management is done from the regional cluster when Cisco Prime Network Registrar is installed. You must install the regional cluster first, and load all licenses in the regional cluster. A local cluster can register with a regional either by registering with the regional cluster during the installation process. However, if the local cluster is behind a NAT instance, then the registration may fail because the initial request does not reach the regional cluster.

In Cisco Prime Network Registrar 8.3 and later, you can register a local cluster that is behind a NAT instance by initiating the registration from the local cluster. To register a local cluster that is spanned by a NAT instance, you must ensure that Cisco Prime Network Registrar 8.3 or later is installed on both the regional and local clusters. You can also verify the license utilization for the local cluster.



Note To register a local cluster when the regional cluster is behind a NAT instance, you need to register the local cluster from the regional server by registering the local cluster from the regional server, selecting the services and resynchronizing the data.

To register a local cluster that is behind a NAT instance, do the following:

Local Web UI

Step 1 From **Administration** menu, choose **Licenses** under the **User Access** submenu to open the List Licenses page.

On the List Licenses page, add the details of the regional cluster.

- a) Enter the IP address (IPv4 and/or IPv6) of the regional cluster.
- b) Enter the SCP port of the regional cluster (1244 is the preset value).
- c) Select the IP address (IPv4 and/or IPv6) of the local cluster that you want to register.
- d) Select the component services that you want to register for the local cluster.

Step 2 Click **Register**.

Note The regional CCM server maintains the license utilization history for all the local clusters in the Cisco Prime Network Registrar system for all counted services (DHCP, DNS, and CDNS).

To view the license utilization for the local cluster, click **Check Poll Status**.

CLI Commands

Use the following commands to register or re-register a local cluster:

```
nrcmd> license register [cdns|dns|dhcp[,...]] [<regional-ip>|<regional-ipv6>]
[<regional-port>]
```

```
nrcmd> license register cdns|dns|dhcp[,...] <regional-ip> <regional-ipv6> [<regional-port>]
```

License History

The License History page allows you to view the licenses utilized in the specified time frame.

Regional Web UI

- Step 1** Log into the regional cluster as superuser.
 - Step 2** From the **Administration** menu, choose **Administrators** to open the List/Add Administrators page for the local cluster version of this page, which is essentially identical.
 - Step 3** Click the **Add Administrators** icon in the Administrators pane, enter **example-regional-admin** in the Name field, then **examplereg** in the Password field in the Add Administrator dialog box, then click **Add Administrator**.
 - Step 4** Multiselect **central-cfg-admin-group** (for cluster administration) and **regional-admin-group** (for user administration) in the Groups drop-down list.
 - Step 5** Click **Save**.
-

CLI Command

Use **license showUtilHistory –full** view the number of utilized IP nodes against the RTUs (Right-to-Use) (see the **license** command in the CLIGuide.html file in the /docs directory for syntax and attribute descriptions).

Configuring Server Clusters

Server clusters are groupings of CCM, DNS, CDNS, DHCP, and TFTP servers at local cluster locations. For example, an organization might have Boston and Chicago clusters of DNS and DHCP servers. A central administrator might want to affect how addresses are allocated at these clusters, or poll subnet utilization or lease history data from them. The central administrator might even want to connect to those local clusters, if the required permissions exist, to view changes there or restart the servers.

View the created clusters on the View Tree of Cluster Servers page. To get there, click **Clusters**. Once the page is populated with clusters, it shows some rich information and provides some useful functions. The Go Local icon allows single sign-on to a local cluster web UI, if an equivalent administrator account exists at the local cluster.

The View Tree of Clusters page might have been populated by manually adding clusters on the List/Add Remote Clusters page, or automatically when adding and synchronizing with routers, which also creates server clusters. The cluster names are links that you can click to edit the cluster information. The resynchronization, replication, and polling functions are described further on in this chapter.

The DHCP server may have the Related Servers icon next to the DHCP server for the cluster. Click this icon to open the List Related Servers for DHCP Server page. These servers can be DNS, TFTP, or DHCP failover servers.

Related Topics

- [Adding Local Clusters, on page 79](#)
- [Editing Local Clusters, on page 80](#)
- [Connecting to Local Clusters, on page 80](#)
- [Synchronizing with Local Clusters, on page 81](#)

[Replicating Local Cluster Data, on page 81](#)

[Viewing Replica Data, on page 82](#)

[Purging Replica Data, on page 82](#)

[Deactivating, Reactivating, and Recovering Data for Clusters, on page 82](#)

[Polling Subnet Utilization and Lease History Data, on page 98](#)

[Enabling Subnet Utilization Collection, on page 99](#)

[Enabling Lease History Collection, on page 100](#)

Adding Local Clusters

Adding local clusters to the regional cluster is the core functionality of the central-cfg-admin role.

To enable subnet utilization and lease history data collection, see [Polling Subnet Utilization and Lease History Data, on page 98](#).

The minimum required values to add a cluster are its name, IP address (IPv4 and/or IPv6) of the machine, administrator username, and password. The cluster name must be unique and its IP address must match that of the host where the CNRDB database is located. Obtain the SCP and HTTP ports, username, and password from the local cluster administrator. The preset value at Cisco Prime Network Registrar installation for the SCP port is 1234 and the HTTP port is 8080.

You can also set whether you want outbound connections to local servers to be secure by setting the *use-ssl* attribute to optional or required. It is set to optional by default, and it requires the Cisco Prime Network Registrar Communications Security Option installed to be effective.

Regional Web UI

From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu. This opens the Manage Servers page. View the local clusters on this page. You can also add server clusters on the List/Add Remote Clusters page. The List/Add Remote Clusters page provides the following functions:

- Connect to a local cluster web UI for local administration.
- Resynchronize with a local cluster to reconcile updates there.
- Pull data over to a regional cluster replica database.
- Purge replica to clear the bad replica data without deleting/re-adding the cluster. Whenever you perform purge replica, you must perform manual replication to the get the replica data again.



Note This option appears only in Expert mode.

- Query subnet utilization data from a local cluster. This function appears only if you are assigned the regional-addr-admin role with at least the subnet-utilization subrole.
- Query lease history data from a local cluster. This function appears only if you are assigned the regional-addr-admin role with at least the lease-history subrole.

To add a cluster, click the **Add Cluster** icon in the **Manage Clusters** pane. This opens the Add Cluster dialog box. For an example of adding a local cluster, see [Create the Local Clusters, on page 119](#). Click **Add Cluster** to return to the List/Add Remote Clusters page.

Local Web UI

You can also manage clusters in the local web UI. See [Configuring Clusters in the Local Web UI, on page 19](#) for details.

CLI Commands

To add a cluster, use **cluster name create** <address | ipv6-address> [attribute=value ...] to give the cluster a name and address and set the important attributes. For example:

```
nrcmd> cluster example-cluster create 192.168.100.101 admin=admin password=changeme
```

Note that the administrator must be a superuser to fully synchronize at the local cluster.

Editing Local Clusters

Editing local clusters at the regional cluster is the core functionality of the central-cfg-admin role.

Regional Web UI

To edit a local cluster, click its name on the Manage Clusters pane to open the Edit Remote Cluster page. This page is essentially the same as the List/Add Remote Clusters page, except for an additional attribute unset function. You can choose the service (dhcp, dns, cdns, or none) that you want to run in the local by checking/unchecking the check boxes provided in the **Local Services** area. Make your changes, then click **Save**.

Local Web UI

You can also edit clusters in the local web UI. See [Configuring Clusters in the Local Web UI, on page 19](#) for details.

CLI Commands

To edit a local cluster, use **cluster name set attribute=value** [attribute=value ...] to set or reset the attributes. For example:

```
nrcmd> cluster Example-cluster set poll-replica-interval=8h
```

Connecting to Local Clusters

In the web UI, if you have an equivalent administrator account at the local cluster, you can single sign-on to the local cluster Manage Servers page by clicking the **Connect** icon on the List/Add Remote Clusters page. To return to the regional cluster web UI, click the **Return** icon at the top right corner of the local cluster page. If you do not have an equivalent account at the local cluster, the Connect icon opens the local cluster login page.

Synchronizing with Local Clusters

Synchronization is configuring regional and local clusters so that they can work together in a unified fashion. When you synchronize:

1. The list of local servers are copied to the regional cluster.
2. A shared secret is established between the regional and local clusters for single sign-on.

Synchronization occurs once when you create a local cluster at the regional cluster. However, changes might occur at the local cluster periodically, requiring you to re synchronize with it. For example, you might change the username and password used to make local connections. Resynchronization does not happen automatically—you must click the **Resynchronize** icon next to the cluster name on the List/Add Remote Clusters page. The result is a positive confirmation for success or an error message for a failure.

When you upgrade the local cluster, you should also resynchronize the cluster. For synchronization to be effective, the user account specified for the local cluster must be a superuser. If you get a synchronization error message, check the local cluster to ensure that it is running properly.



Note When you resynchronize clusters at the regional cluster, an automatic reinitialization of replica data occurs. The result is that for larger server configurations, resynchronization might take several minutes. The benefit, however, is that you do not need a separate action to update the replica data.

Replicating Local Cluster Data

Replication is copying the configuration data from a local server to the regional cluster replica database. Replication needs to occur before you can pull DHCP object data into the regional server database. During replication:

1. The current data from the local database is copied to the regional cluster. This usually occurs once.
2. Any changes made in the master database since the last replication are copied over.

Replication happens at a given time interval. You can also force an immediate replication by clicking the **Replicate** icon on the List/Add Remote Clusters page.

You can set the automatic replication interval on the Add Server Cluster page, or adjust it on the Edit Server Cluster page, using the *poll-replica-interval* attribute. This interval is preset at four hours. You can also set the fixed time of day to poll replica data by using the *poll-replica-offset* attribute; its default value is zero hours (no offset). The *poll-replica-rrs* attribute controls the replication of RR data without disabling other data replication. This attribute is present in Manage Servers and Manage Clusters page and has the values - none, all, and protected. If *poll-replica-rrs* is set to none, no RR data will be replicated for this cluster. If unset, the CCM server setting will apply.



Caution If the replica database is corrupted in any way, the regional CCM server will not start. If you encounter this problem, stop the regional service, remove (or move) the replica database files located in the *install-path* /regional/data/replica directory (and the log files in the /logs subdirectory), then restart the regional server. Doing so recreates the replica database without any data loss.

Viewing Replica Data

In the web UI, you can view the replica data cached in the replica database at the regional cluster by choosing **View Replica Data** from **Servers** submenu under the **Operate** menu. This opens the View Replica Class List page.

Regional Web UI

Select the:

1. Cluster in the Select Cluster list.
2. Object class in the Select Class list.
3. Replicate the data for the cluster and class chosen. Click the **Replicate Data for Cluster** button.
4. View the replica data. Click **View Replica Class List**, which opens a List Replica Data for Cluster page for the cluster and specific class of object you choose. On this page, you can:
 - Click the name of an object to open a View page at the regional cluster. Return to the List Replica page by clicking **Return to object List**.



Note The List Replica Address Blocks and List Replica Subnets pages do not provide this function. To view the address blocks or subnets for the local cluster, use the **Go Local** icon.

- Click the **Connect** icon to go to the List page for the object at the local cluster. Return to the List Replica *object* page by clicking the **Return** icon.

Click **Return** on the List Replica Data for Cluster page to return to the View Replica Class List page.

Purging Replica Data

In the regional web UI (only in Expert mode), you can clear the bad replica data without deleting/re-adding the clusters by clicking the **Purge Replica** icon on the List/Add Remote Clusters page. Whenever you perform purge replica, you must perform manual replication to get the replica data again.

Deactivating, Reactivating, and Recovering Data for Clusters

Deactivating a cluster might be necessary if you suspect that a hard disk error occurred where configuration data could have been lost. You can deactivate the cluster, remedy the problem, recover cluster data from the replica database, then reactivate the cluster. This saves you from having to delete and then recreate the cluster with all of its data lost in the process.

Deactivating, reactivating, and recovering the data for a cluster is available only in the web UI, and you must be an administrator assigned the central-config-admin role.

Data that is not recovered (and that you need to manually restore) includes:

- Contents of the **cnr.conf** file (see [Modifying the cnr.conf File, on page 157](#))
- Web UI configuration files
- Unprotected DNS resource records
- Administrator accounts



Note If the local secret db is lost, the old references are no longer valid, even though they are restored. To recover your passwords, you have to use central management for your admins, and then push them to your local clusters. Routers, since they have their own secrets, also need to be centrally managed and then should be re-pushed. For the local cluster partner objects, running the sync from regional will create valid objects, but the old cluster objects may need to be deleted first.

- Lease history
- Extension scripts



Note Restoring the data to a different IP address requires some manual reconfiguration of such things as DHCP failover server pair and High-Availability (HA) DNS server pair addresses.

Regional Web UI

Deactivate a cluster by clicking the **Deactivate** button for the cluster. This immediately changes the button to **Reactivate** to show the status of the cluster. Deactivating a cluster disables deleting, synchronizing, replicating data, and polling subnet utilization and lease history. These operations are not available while the cluster is deactivated.

Deactivating the cluster also displays the Recover icon in the Recover Data column of the cluster. Click this icon to recover the replica data. This opens a separate “in process” status window that prevents any operations on the web UI pages while the recovery is in process. As soon as the recovery is successful, the disabled functions are again enabled and available.

To reactivate the cluster, click the **Reactivate** button to change back to the **Deactivate** button and show the status as active.

CLI Commands

The following cluster commands are only available when connected to a regional cluster:

Table 9: Cluster Commands

Action	Command
Activate	cluster <i>name</i> activate
Deactivate	cluster <i>name</i> deactivate
Resynchronize	cluster <i>name</i> resynchronize
Synchronize	cluster <i>name</i> sync
Update Replica Data	cluster <i>name</i> updateReplicaData
Remove Replica Data	cluster <i>name</i> removeReplicaData
Recover Data	cluster <i>name</i> recoverData

Action	Command
Poll Lease History	cluster name pollLeaseHistory
Get Lease History State	cluster name getLeaseHistoryState
Poll Subnet Utilization	cluster name pollSubnetUtilization

Central Configuration Management Server

The CCM servers at the local and regional clusters provide the infrastructure for Cisco Prime Network Registrar operation and user interfaces. The CCM Server reads, writes, and modifies the Cisco Prime Network Registrar database (CCM DB). The main purpose of the CCM Server is to store and propagate data from the user to the protocol servers, and from the servers back to the user.

The change set is the fundamental unit of change to a data store. It sends incremental changes to a replicating server and provides an audit log for changes to the data store. Change sets consist of lists of change entries that are groups of one or more changes to a single network object. The web UI provides a view of the change sets for each data store.

Managing CCM Server

You can view logs and startup logs; edit the server attributes.

To view logs and startup logs, in the local cluster web UI, from the **Operate** menu, choose **Manage Servers** to open the Manage Servers page.

Editing CCM Server Properties

You can edit the CCM server properties using the Edit CCM Server page.

Local Basic or Advanced Web UI

-
- Step 1** To access the CCM server properties, choose **Manage Servers** under **Operate** menu to open the Manage Servers page.
 - Step 2** Click **Local CCM Server** in the Manage Servers pane on the left. The Edit Local CCM Server page appears. This page displays all the CCM server attributes.
 - Step 3** Modify the settings as per your requirement.
 - Step 4** Click **Save** to save the CCM server attribute modifications.
-

Trivial File Transfer

The Trivial File Transfer Protocol (TFTP) is a way of transferring files across the network using the User Datagram Protocol (UDP), a connectionless transport layer protocol. Cisco Prime Network Registrar maintains a TFTP server so that systems can provide device provisioning files to cable modems that comply with the Data Over Cable Service Interface Specification (DOCSIS) standard. The TFTP server buffers the DOCSIS

file in its local memory as it sends the file to the modem. After a TFTP transfer, the server flushes the file from local memory. TFTP also supports non-DOCSIS configuration files.

Here are some of the features of the Cisco Prime Network Registrar TFTP server:

- Complies with RFCs 1123, 1350, 1782, and 1783
- Includes a high performance multithreaded architecture
- Supports IPv6
- Caches data for performance enhancements
- Is configurable and controllable
- Includes flexible path and file access controls
- Includes audit logging of TFTP connections and file transfers
- Has a default root directory in the Cisco Prime Network Registrar install-path/data/tftp

Related Topics

[Viewing and Editing the TFTP Server, on page 85](#)

[Managing the TFTP Server Network Interfaces, on page 85](#)

Viewing and Editing the TFTP Server

At the local cluster, you can edit the TFTP server to modify its attributes. You must be assigned the server-management subrole of the ccm-admin role.

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 (see [Managing Servers, on page 129](#)).
- Step 2** Click the Local TFTP Server link in the left pane to open the Edit Local TFTP Server page.
- You can click the name of any attribute to open a description window for the attribute.
- Step 3** To unset any attribute value, check the check box in the Unset? column.
- Step 4** Click **Save** to save the changes or **Revert** to cancel the changes.
-

CLI Commands

Use **tftp** to show the attribute values. Use **tftp set attribute=value [attribute=value ...]** or **tftp enable attribute** to set or enable attributes. You can also use **tftp serverLogs show**, and **tftp serverLogs nlogs=number logsize=size**.

Managing the TFTP Server Network Interfaces

You can manage the network interfaces for the TFTP server.

Local Advanced Web UI

Manage the network interfaces associated with the TFTP server by clicking the **Network Interfaces** tab for the selected Local TFTP Server in the Manage Servers page. You can view the default configured network

interfaces, and create and edit additional ones. To create and edit them, you must be assigned the server-management subrole of the `ccm-admin` role.

The columns in the Network Interfaces page are:

- **Name**—Name of the network interface, such as the LAN adapter, loopback, and Fast Ethernet interfaces. If the name is under the Configured Interfaces column, you can edit and delete the interface. Clicking the name opens the Edit TFTP Server Network Interface page so that you can edit the interface name and addresses. Make the changes and then click **Save** on this page.
- **IP Address**—IP address of the network interface.
- **IPv6 Address**—IPv6 address, if applicable, of the network interface.
- **Flags**—Flags for whether the interface should be zero-broadcast, virtual, v4, v6, no-multicast, or receive-only.
- **Configure**—To configure a new network interface, click the **Configure** icon next to the interface name. This creates another interface based on the one selected, but with a more general IP address, and adds this interface to the Configured Interfaces for this TFTP Server.
- **List of available interfaces for this TFTP server**—User-configured network interfaces, showing each name and associated address. Click the interface name to edit it or click the **Delete** icon to delete it.

To return to managing the server, click **Revert**.

CLI Commands

Use the `tftp-interface` commands.

Simple Network Management

The Cisco Prime Network Registrar Simple Network Management Protocol (SNMP) notification support allows you to query the DHCP and DNS counters, be warned of error conditions and possible problems with the DNS and DHCP servers, and monitor threshold conditions that can indicate failure or impending failure conditions.

Cisco Prime Network Registrar implements SNMP Trap Protocol Data Units (PDUs) according to the SNMPv2c standard. Each trap PDU contains:

- Generic-notification code, if enterprise-specific.
- A specific-notification field that contains a code indicating the event or threshold crossing that occurred.
- A variable-bindings field that contains additional information about certain events.

Refer to the Management Information Base (MIB) for the details. The SNMP server supports only reads of the MIB attributes. Writes to the attributes are not supported.

The following MIB files are required:

- **Traps**—`CISCO-NETWORK-REGISTRAR-MIB.my` and `CISCO-EPM-NOTIFICATION-MIB.my`
- **DNS server**—`CISCO-DNS-SERVER-MIB.my`



Note The Caching DNS server requires only a subset of the DNS MIB when it is operating. Caching DNS server only supports the *server-start* and *server-stop* notification events.

- **DHCPv4 server**—CISCO-IETF-DHCP-SERVER-MIB.my
- **DHCPv4 server capability**—CISCO-IETF-DHCP-SERVER-CAPABILITY.my
- **DHCPv4 server extensions**—CISCO-IETF-DHCP-SERVER-EXT-MIB.my
- **DHCPv4 server extensions capability**—CISCO-IETF-DHCP-SERVER-EXT-CAPABILITY.my
- **DHCPv6 server**—CISCO-NETREG-DHCPV6-MIB.my (experimental)



Note The MIB, CISCO-NETREG-DHCPV6-MIB is defined to support query of new DHCP v6 related statistics and new DHCP v6 traps.

These MIB files are available in the /misc directory of the Cisco Prime Network Registrar installation path.

The following URL includes all files except the experimental CISCO-NETREG-DHCPV6-MIB.my file:

<ftp://ftp.cisco.com/pub/mibs/supportlists/cnr/cnr-supportlist.html>

The following dependency files are also required:

- **Dependency for DHCPv4 and DHCPv6**—CISCO-SMI.my
- **Additional dependencies for DHCPv6**—INET-ADDRESS-MIB.my

These dependency files are available along with all the MIB files at the following URL:

<ftp://ftp.cisco.com/pub/mibs/v2/>

To get the object identifiers (OIDs) for the MIB attributes, go to the equivalently named .oid file at:

<ftp://ftp.cisco.com/pub/mibs/oid/>

Related Topics

[Setting Up the SNMP Server, on page 87](#)

[How Notification Works, on page 88](#)

[Handling SNMP Notification Events, on page 91](#)

[Handling SNMP Queries, on page 94](#)

Setting Up the SNMP Server

To perform queries to the SNMP server, you need to set up the server properties.

Local Basic or Advanced and Regional Web UI

- Step 1** From the **Operate** menu, choose **Manage Servers** under the **Servers** submenu to open the Manage Servers page (see [Managing Servers, on page 129](#)).
- Step 2** Click the **Local SNMP Server** link to open the Edit Local SNMP Server page.
- Step 3** The *Community string* attribute is the password to access the server. (The community string is a read community string only.) The preset value is **public**.
- Step 4** You can specify the Log Settings, Miscellaneous Options and Settings, and Advanced Options and Settings:
- **trap-source-addr**—Optional sender address to use for outgoing traps.

- **trap-source-ip6address**— Optional sender IPv6 address to use for outgoing traps.
- **server-active**—Determines whether the SNMP server is active for queries. The default value is true. If set to false, the server will run, but is not accessible for queries and does not send out traps.
- **cache-ttl**—Determines how long the SNMP caches responds to queries, default to 60 seconds.

Step 5 To manage the SNMP server interfaces in the Advanced mode, click the **Network Interfaces** tab. You can view the default configured network interfaces, and create and edit additional ones. To create and edit them, you must be assigned the server-management subrole of the ccm-admin role. The interface properties are similar to those for the TFTP server (see [Managing the TFTP Server Network Interfaces, on page 85](#)).

Step 6 To manage trap recipients for the server:

- Click the **Trap Recipients** tab.
- Enter the name of the trap recipient.
- Enter the IPv4 and/or IPv6 address of a trap recipient.
- Click **Add Trap Recipient**.
- Repeat for each additional trap recipient.
- To set the port, community string, and agent address for a trap recipient, click its name on the Trap Recipients tab to open the Edit Trap Recipient page, then set the values.

Step 7 Complete the SNMP server setup by clicking **Save**.

CLI Commands

To set the community string in the CLI so that you can access the SNMP server, use **snmp set community=*name***. Use **snmp set trap-source-addr** to set the trap source IPv4 address. Use **snmp set trap-source-ip6address** to set the trap source IPv6 address. Use **snmp disable server-active** to deactivate the SNMP server and **snmp set cache-ttl=*time*** to set the cache time-to-live.

To set trap recipients, use **trap-recipient**, in the following syntax to include the IP address:

```
nrcmd> trap-recipient name create ip-addr=
nrcmd> trap-recipient name create ip6address=
```

You can also add the *agent-address*, *community*, and *port-number* values for the trap recipient.

Other SNMP-related commands include **snmp disable server-active** to prevent the server from running when started and the **snmp-interface** commands to configure the interfaces. The **addr-trap** command is described in [Managing the TFTP Server Network Interfaces, on page 85](#).

How Notification Works

Cisco Prime Network Registrar SNMP notification support allows a standard SNMP management station to receive notification messages from the DHCP and DNS servers. These messages contain the details of the event that triggered the SNMP trap.

Cisco Prime Network Registrar generates notifications in response to predetermined events that the application code detects and signals. Each event can also carry with it a particular set of parameters or current values. For example, the *free-address-low-threshold* event can occur in the scope with a value of 10% free. Other scopes and values are also possible for such an event, and each type of event can have different associated parameters.

The following table describes the events that can generate notifications.

Table 10: SNMP Notification Events

Event	Notification
Address conflict with another DHCP server detected (<i>address-conflict</i>)	An address conflicts with another DHCP server.
DNS queue becomes full (<i>dns-queue-size</i>)	The DHCP server DNS queue fills and the DHCP server stops processing requests. (This is usually a rare internal condition.)
Duplicate IP address detected (<i>duplicate-address</i> and <i>duplicate-address6</i>)	A duplicate IPv4 or IPv6 address occurs.
Duplicate IPv6 prefix detected (<i>duplicate-prefix6</i>)	A duplicate IPv6 prefix occurs.
Failover configuration mismatch (<i>failover-config-error</i>)	A DHCP failover configuration does not match between partners.
Free-address thresholds (<i>free-address-low</i> and <i>free-address-high</i> ; or <i>free-address6-low</i> and <i>free-address6-high</i>)	The high trap when the number of free IPv4 or IPv6 addresses exceeds the high threshold; or a low trap when the number of free addresses falls below the low threshold after previously triggering the high trap.
High-availability (HA) DNS configuration mismatch (<i>ha-dns-config-error</i>)	An HA DNS configuration does not match between partners.
HA DNS partner not responding (<i>ha-dns-partner-down</i>)	An HA DNS partner stops responding to the DNS server.
HA DNS partner responding (<i>ha-dns-partner-up</i>)	An HA DNS partner responds after having been unresponsive.
DNS masters not responding (<i>masters-not-responding</i>)	Master DNS servers stop responding to the DNS server.
DNS masters responding (<i>masters-responding</i>)	Master DNS servers respond after having been unresponsive.
Other server not responding (<i>other-server-down</i>)	A DHCP failover partner, or a DNS or LDAP server, stops responding to the DHCP server.
Other server responding (<i>other-server-up</i>)	DHCP failover partner, or a DNS or LDAP server, responds after having been unresponsive.
DNS secondary zones expire (<i>secondary-zone-expired</i>)	A DNS secondary server can no longer claim authority for zone data when responding to queries during a zone transfer.
Server start (<i>server-start</i>)	The DHCP or DNS server is started or reinitialized.
Server stop (<i>server-stop</i>)	The DHCP or DNS server is stopped.

Resource Monitoring SNMP Notifications

If SNMP traps are enabled for the resource limit alarms, Cisco Prime Network Registrar generates SNMP traps when the monitored resources exceed the critical or warning levels. SNMP traps are generated for resource limits:

- Whenever the resource's value exceeds the warning or critical limits (these are sent periodically while the value continues to exceed either threshold).
- Whenever the resource's value returns to a level below the warning limit.

The SNMP server generates a trap using the CISCO-EPM-NOTIFICATION-MIB. The mapping is as follows:

Trap Attribute Name	Object ID	Type	Value for Resource Events
cenAlarmVersion	1.3.6.1.4.1.99.311.1.1.2.1.2	SnmpAdminString (SIZE(1..16))	"1.2"
cenAlarmTimestamp	1.3.6.1.4.1.99.311.1.1.2.1.3	Timestamp	Time of last resource event state change
cenAlarmUpdatedTimeStamp	1.3.6.1.4.1.99.311.1.1.2.1.4	Timestamp	"current" time
cenAlarmInstanceID	1.3.6.1.4.1.99.311.1.1.2.1.5	SnmpAdminString (SIZE(1..20))	A unique id for the event - just hexadecimal digits
cenAlarmStatus	1.3.6.1.4.1.99.311.1.1.2.1.6	Integer32 (1..250)	1 (for Not acknowledged)
cenAlarmStatusDefinition	1.3.6.1.4.1.99.311.1.1.2.1.7	SnmpAdminString (SIZE(1..255))	"1,Not acknowledged"
cenAlarmType	1.3.6.1.4.1.99.311.1.1.2.1.8	Integer	Not Used
cenAlarmCategory	1.3.6.1.4.1.99.311.1.1.2.1.9	Integer32 (1..250)	100 (for Raw alarm)
cenAlarmCategoryDefinition	1.3.6.1.4.1.99.311.1.1.2.1.10	SnmpAdminString (SIZE(1..255))	"100,Raw alarm"
cenAlarmServerAddressType	1.3.6.1.4.1.99.311.1.1.2.1.11	InetAddressType	Cluster server address type - IPv4(1) or IPv6(2)
cenAlarmServerAddress	1.3.6.1.4.1.99.311.1.1.2.1.12	InetAddress	Cluster address (based on local cluster's object)
cenAlarmManagedObjectClass	1.3.6.1.4.1.99.311.1.1.2.1.13	SnmpAdminString (SIZE(1..255))	"Application"
cenAlarmManagedObjectAddressType	1.3.6.1.4.1.99.311.1.1.2.1.14	InetAddressType	Not used
cenAlarmManagedObjectAddress	1.3.6.1.4.1.99.311.1.1.2.1.15	InetAddress	Not used
cenAlarmDescription	1.3.6.1.4.1.99.311.1.1.2.1.16	OctetString (SIZE(1..1024))	Description formatted as " , "

Trap Attribute Name	Object ID	Type	Value for Resource Events
cenAlarmSeverity	1.3.6.1.4.1.99.311.1.1.2.1.17	Integer32	0 for Clear, 2 for Warning, and 5 for Critical
cenAlarmSeverityDefinition	1.3.6.1.4.1.99.311.1.1.2.1.18	SnmpAdminString (SIZE(1..255))	String alarm severity, one of "0,Clear", "2,Warning", or "5,Critical"
cenAlarmTriageValue	1.3.6.1.4.1.99.311.1.1.2.1.19	Integer32 (0..100)	Not used
cenEventIDList	1.3.6.1.4.1.99.311.1.1.2.1.20	OctetString (SIZE(1..1024))	Not used
cenUserMessage1	1.3.6.1.4.1.99.311.1.1.2.1.21	SnmpAdminString (SIZE(1..255))	Name of monitored resource
cenUserMessage2	1.3.6.1.4.1.99.311.1.1.2.1.22	SnmpAdminString (SIZE(1..255))	Server name (dhcp, dns, cdns, ...)
cenUserMessage3	1.3.6.1.4.1.99.311.1.1.2.1.23	SnmpAdminString (SIZE(1..255))	"Network Registrar"
cenAlarmMode	1.3.6.1.4.1.99.311.1.1.2.1.24	Integer	3 (event)
cenPartitionNumber	1.3.6.1.4.1.99.311.1.1.2.1.25	Gauge (0..100)	Not used
cenPartitionName	1.3.6.1.4.1.99.311.1.1.2.1.26	SnmpAdminString (SIZE(1..255))	Not used
cenCustomerIdentification	1.3.6.1.4.1.99.311.1.1.2.1.27	SnmpAdminString (SIZE(1..255))	Not used
cenCustomerRevision	1.3.6.1.4.1.99.311.1.1.2.1.28	SnmpAdminString (SIZE(1..255))	Not used
cenAlertID	1.3.6.1.4.1.99.311.1.1.2.1.29	SnmpAdminString (SIZE(1..20))	Same as cenAlarmInstanceID

For more information on resource limit alarms, see [Monitoring Resource Limit Alarms](#), on page 108.

Handling SNMP Notification Events

When Cisco Prime Network Registrar generates a notification, it transmits a single copy of the notification as an SNMP Trap PDU to each recipient. All events (and scopes or prefixes) share the list of recipients and other notification configuration data, and the server reads them when you initialize the notification.

You can set SNMP attributes in three ways:

- For the DHCP server, which includes the traps to enable and the default free-address trap configuration if you are not specifically configuring traps for scopes or prefixes (or their templates).

- On the scope or prefix (or its template) level by setting the *free-address-config* attribute.
- For the DNS server, which includes a *traps-enabled* setting.

To use SNMP notifications, you must specify trap recipients that indicate where trap notifications should go. By default, all notifications are enabled, but you must explicitly define the recipients, otherwise no notifications can go out. The IP address you use is often **localhost**.

The DHCP server provides special trap configurations so that it can send notifications, especially about free addresses for DHCPv4 and DHCPv6. You can set the trap configuration name, mode, and percentages for the low threshold and high threshold. The mode determines how scopes aggregate their free-address levels.

DHCP v4 Notification

The DHCP v4 modes and thresholds are (see also [Handling Deactivated Scopes or Prefixes, on page 93](#)):

- **scope mode**—Causes each scope to track its own free-address level independently (the default).
- **network mode**—Causes all scopes set with this trap configuration (through the scope or scope template *free-address-config* attribute) to aggregate their free-address levels if the scopes share the same *primary-subnet*.
- **selection-tags mode**—Causes scopes to aggregate their free-address levels if they share a primary subnet and have a matching list of selection tag values.
- **low-threshold**—Free-address percentage at which the DHCP server generates a low-threshold trap and re-enables the high threshold. The free-address level for scopes is the following calculation:

$$100 * \frac{\text{available-nonreserved-leases}}{\text{total-configured-leases}}$$

- **high-threshold**—Free-address percentage at which the DHCP server generates a high-threshold trap and re-enables the low threshold.

DHCP v6 Notification

The DHCP v6 modes and thresholds are (see also [Handling Deactivated Scopes or Prefixes, on page 93](#)):

- **prefix mode**—Causes each prefix to track its own free-address level independently.
- **link mode**—Causes all prefixes configured for the link to aggregate their own free-address levels if all prefixes share the same link.
- **v6-selection-tags mode**—Causes prefixes to aggregate their free-address levels if they share a link and have a matching list of selection tag values.
- **low-threshold**—Free-address percentage at which the DHCP server generates a low-threshold trap and re-enables the high threshold. The free-address level for prefixes is the following calculation:

$$100 * \frac{\text{max-leases} - \text{dynamic-leases}}{\text{max-leases}}$$

- **high-threshold**—Free-address percentage at which the DHCP server generates a high-threshold trap and re-enables the low threshold.

Handling Deactivated Scopes or Prefixes

A deactivated scope or prefix never aggregates its counters with other scopes or prefixes. For example, if you configure a prefix with **link** or **v6-selection-tags** trap mode, and then deactivate the prefix, its counters disappear from the total count on the aggregation. Any changes to the leases on the deactivated prefix do not apply to the aggregate totals.

Therefore, to detect clients for deactivated scopes or prefixes, you must set the event mode to **scope** or **prefix**, and not to any of the aggregate modes (**network**, **selection-tags**, **link**, or **v6-selection-tags**).

The use case for setting traps on deactivated prefixes, for example, is network renumbering. In this case, you might want to monitor both the new prefixes (as an aggregate, ensuring that you have enough space for all the clients) and old prefixes to ensure that their leases are freed up. You would probably also want to set the high threshold on an old prefix to 90% or 95%, so that you get a trap fired when most of its addresses are free.

Local Basic or Advanced Web UI

Access the SNMP attributes for the DHCP server by choosing **Manage Servers** from the **Operate** menu, then click **Local DHCP Server** in the left pane. You can view the SNMP attributes under SNMP (in Basic mode) or SNMP Settings (in Advanced mode) in the Edit DHCP Server page.

The four *lease-enabled* values (free-address6-low, free-address6-high, duplicate-address6, duplicate-prefix6) pertain to DHCPv6 only. Along with the traps to enable, you can specify the default free-address trap configuration by name, which affects all scopes and prefixes or links not explicitly configured.

To add a trap configuration, do the following:

-
- Step 1** In Advanced mode, from the **Deploy** menu, choose **Traps** under the **DHCP** submenu to access the DHCP trap configurations. The List/Add Trap Configurations page appears.
 - Step 2** Click the **Add Traps** icon in the left pane to open the Add AddrTrapConfig page.
 - Step 3** Enter the name, mode, and threshold percentages, then click **Add AddrTrapConfig**.
-

Editing Trap Configuration

To edit a trap configuration, do the following:

-
- Step 1** Click the desired trap name in the Traps pane to open the Edit Trap Configuration page
 - Step 2** Modify the name, mode, or threshold percentages.
 - Step 3** Click the **on** option for the *enabled* attribute to enable the trap configuration.
 - Step 4** Click **Save** for the changes to take effect.
-

Deleting Trap Configuration

To delete a trap configuration, select the trap in the Traps pane and click the **Delete** icon, then confirm or cancel the deletion.

Regional Basic or Advanced Web UI

In the regional web UI, you can add and edit trap configurations as in the local web UI. You can also pull replica trap configurations and push trap configurations to the local cluster on the List/Add Trap Configurations page.

Server Up/Down Traps

Every down trap must be followed by a corresponding up trap. However, this rule is not strictly applicable in the following scenarios:

1. If a failover partner or LDAP server or DNS server or HA DNS partner is down for a long time, down traps will be issued periodically. An up trap will be generated only when that server or partner returns to service.
2. If the DHCP or DNS server is reloaded or restarted, the prior state of the partner or related servers is not retained and duplicate down or up traps can result.



Note Other failover partner or LDAP server or DNS server or HA DNS partner up or down traps occur only to communicate with that partner or server, and therefore may not occur when the other partner or server goes down or returns to service.

CLI Commands

To set the trap values for the DHCP server at the local cluster, use **dhcp set traps-enabled=value**. You can also set the *default-free-address-config* attribute to the trap configuration. For example:

```
nrcmd> dhcp set traps-enabled=server-start,server-stop,free-address-low,free-address-high
```

```
nrcmd> dhcp set default-free-address-config=v4-trap-config
```



Note If you do not define a *default-free-address-config* (or *v6-default-free-address-config* for IPv6), Cisco Prime Network Registrar creates an internal, unlisted trap configuration named **default-aggregation-addr-trap-config**. Because of this, avoid using that name for a trap configuration you create.

To define trap configurations for DHCPv4 and DHCPv6, use **addr-trap name create** followed by the *attribute=value* pairs for the settings. For example:

```
nrcmd> addr-trap v4-trap-conf create mode=scope low-threshold=25% high-threshold=30%
```

```
nrcmd> addr-trap v6-trap-conf create mode=prefix low-threshold=20% high-threshold=25%
```

Handling SNMP Queries

You can use SNMP client applications to query the following MIBs:

- CISCO-DNS-SERVER-MIB.my
- CISCO-IETF-DHCP-SERVER-MIB.my
- CISCO-IETF-DHCP-SERVER-EXT-MIB.my
- CISCO-NETREG-DHCPV6-MIB.my (experimental)

When the SNMP server receives a query for an attribute defined in one of these MIBs, it returns a response PDU containing that attribute value. For example, using the NET-SNMP client application (available over the Internet), you can use one of these commands to obtain a count of the DHCPDISCOVER packets for a certain address:

```
C:\net-snmplib\bin>snmpget -m ALL -v 2c -c public
192.168.241.39:4444.iso.org.dod.internet.private.enterprises.cisco.ciscoExperiment.
ciscoIetfDhcpSrvMIB.ciscoIetfDhcpv4SrvMIBObjects.cDhcpv4Counters.cDhcpv4CountDiscovers
```

```
CISCO-IETF-DHCP-SERVER-MIB::cDhcpv4CountDiscovers.0 = Counter32: 0
C:\net-snmplib\bin>snmpget -m ALL -v 2c -c public
192.168.241.39:4444
1.3.6.1.4.1.9.10.102.1.3.1
```

```
CISCO-IETF-DHCP-SERVER-MIB::cDhcpv4CountDiscovers.0 = Counter32: 0
```

Both commands return the same results. The first one queries the full MIB attribute name, while the second one queries its OID equivalent (which can be less error prone). As previously described, the OID equivalents of the MIB attributes are located in the relevant files at the following URL:

<ftp://ftp.cisco.com/pub/mibs/oid/>

For example, the CISCO-IETF-DHCP-SERVER-MIB.oid file includes the following OID definition that corresponds to the previous query example:

```
"cDhcpv4CountDiscovers" "1.3.6.1.4.1.9.10.102.1.3.1"
```

Here are some possible SNMP query error conditions:

- The community string sent in the request PDU does not match what you configured.
- The version in the request PDU is not the same as the supported version (SNMPv2).
- If the object being queried does not have an instance in the server, the corresponding variable binding type field is set to SNMP_NOSUCHINSTANCE. With a GetNext, if there is no next attribute, the corresponding variable binding type field is set to SNMP_ENDOFMIBVIEW.
- If no match occurs for the OID, the corresponding variable binding type field is set to SNMP_NOSUCHOBJECT. With a GetNext, it is set to SNMP_ENDOFMIBVIEW.
- If there is a bad value returned by querying the attribute, the error status in the response PDU is set to SNMP_ERR_BAD_VALUE.

Integrating Cisco Prime Network Registrar SNMP into System SNMP

You can integrate the Cisco Prime Network Registrar SNMP server into the SNMP server for the system it runs on. The integration can be done in a way where the system will respond to queries for Cisco Prime Network Registrar MIB entries. On systems using NET-SNMP (and compatible servers) this is done by adding the following entries to the /etc/snmp/snmpd.conf configuration file:

- For IPv4:

```
view systemview included .1.3.6.1.4.1.9.9
view systemview included .1.3.6.1.4.1.9.10

proxy -v 2c -c public 127.0.0.1:4444 .1.3.6.1.4.1.9.9
```

```
proxy -v 2c -c public 127.0.0.1:4444 .1.3.6.1.4.1.9.10
```

- For IPv6:

```
view systemview included .1.3.6.1.4.1.9.9
view systemview included .1.3.6.1.4.1.9.10

proxy -v 2c -c public [::1]:4444 .1.3.6.1.4.1.9.9
proxy -v 2c -c public [::1]:4444 .1.3.6.1.4.1.9.10
```

The community string **public** and the port number **4444** may have to be replaced if the Cisco Prime Network Registrar SNMP server has been configured with different values for those settings.

NET-SNMP is commonly available on Linux and other Unix-like systems. On other systems, similar mechanisms may also be available.

Bring Your Own Device Web Server

The BYOD web server at the regional cluster provides the infrastructure for Cisco Prime Network Registrar BYOD operation. The main purpose of the BYOD Web Server is to authenticate the user against AD and collect the device metadata by registering the user's own device in Cisco Prime Network Registrar.

Managing BYOD Web Server

You can view logs and startup logs; edit the server attributes.

To view logs and startup logs, in the regional cluster web UI, from the **Operate** menu, choose **Manage Servers** under the **Server** submenu to open the Manage Servers page.

Editing BYOD Web Server Properties

You can edit the BYOD web server properties using the Edit Local BYOD Web Server page.

Regional Basic or Advanced or Expert Web UI

-
- Step 1** To access the BYOD web server properties, choose **Manage Servers** under **Operate** menu to open the Manage Servers page.
- Step 2** Click **Local BYOD Web Server** in the Manage Servers pane on the left. The Edit Local BYOD Web Server page appears. This page displays the BYOD web server attributes.
- **KeyStore Settings:** Redirects the "http call" of the BYOD web server to secure "https" with a combination of key store file and key store password.
 - **LDAP Settings:** Specifies the remote LDAP server used for client registration.
 - **Additional Attributes (Auto- start):** Indicates if the BYOD server should be started automatically after every server agent restart.
- Step 3** Modify the settings as per your requirement.
- Step 4** Click **Save** to save the BYOD web server attribute modifications.
- Step 5** Click **Start Server** or **Restart Server** to apply the modifications to the BYOD web server.
-

Setting Up BYOD Theme and Content

You can create the content and multiple BYOD themes at the regional cluster which can be applied in BYOD web server interface.

Adding and Previewing BYOD Themes

You can create your own themes on the regional cluster using the BYOD Theme page and apply the created theme to the BYOD web server so that the logo, background, font, and other properties of the BYOD interface are displayed as per your customization. The created theme can be previewed prior to publishing it to the BYOD web server.

To add and preview a theme:

Regional Advanced or Expert Web UI

-
- Step 1** From the **Deploy** menu, choose **Theme** under the **BYOD** submenu to open the List/Add Custom Theme page.
- Step 2** Click the **Add Theme** icon in the Theme pane.
The **Add Custom Theme** window appears.
- Step 3** Enter the Theme Name in the Add Custom Theme window.
- Step 4** Click **Add Custom Theme** to create a new BYOD Theme.
- Step 5** Update the Edit Custom Theme page with required theme attributes.
- Step 6** Click the **Review Theme** icon in the top right corner of the List/Add Custom Theme page.
The Theme Preview window appears displaying the BYOD page with the newly added theme.
- Note** You can navigate between the BYOD pages with **Register** and **Reboot** to view how the theme is applied to the BYOD pages. By default, the **Theme preview** window loads the BYOD Device Registration page.
- Step 7** Click **Reboot** to preview your theme in the Device Activation page.
- Note** You must close the Theme Preview window after preview to return to the List/Add Custom Theme page in the regional server.
- Step 8** Click **Save** in the List/Add Custom Theme page in the regional server to apply the theme to the BYOD web server or click **Revert** to change the attribute values prior to saving the Custom Theme.
- Note** You can modify and preview the theme any number of times. Only the recently saved theme is applied to the BYOD web server.
-

Adding and Previewing BYOD Content

You can create the BYOD web server contents such as login page message, about, terms of services, contact details, and help message on the BYOD content page of the regional cluster, and preview it prior to publishing it to the BYOD web server. These contents can be published in the BYOD web server interface for the device registration and login pages.

To add and review content:

Regional Advanced or Expert Web UI

-
- Step 1** From the **Deploy** menu, choose **Content** under the **BYOD** submenu to open the Edit BYOD content page.
- Step 2** Upload the file or enter relevant text in the Edit BYOD content page.
- Note** You must upload only .html, .htm, or .txt files.
- Step 3** Click **Review** to preview the content in the Edit BYOD content page before saving. A **Content Review** window containing the contents appears.
- Step 4** Click on **About/Terms of Service/Contact/Help** in the content review page to preview the content added in the EDIT BYOD content page of the regional server.
- Step 5** Click **Save** to publish the added BYOD content to the BYOD web server.
-

Polling Process

When the regional cluster polls the local cluster for subnet utilization or lease history, it first requests all available data up to the current time. This time is recorded in the history databases, and subsequent polls request only new data from this time forward. All times are stored relative to each local cluster time, adjusted for that cluster time zone.

If the times on each server are not synchronized, you might observe odd query results. For example, if the regional cluster time lags behind that of a local cluster, the collected history might be in the future relative to the time range queries at the regional cluster. If so, the result of the query would be an empty list. Data merged from the several clusters could also appear out of sequence, because of the different time skews between local clusters. This type of inconsistency would make it difficult to interpret trends. To avoid these issues, using a network time service for all clusters is strongly recommended.

Polling Subnet Utilization and Lease History Data

Subnet utilization and lease history data are automatically collected at any regional cluster where these features are enabled for the DHCP server or failover pair. The default polling interval to update the regional databases is 4 hours. You can poll the servers by clicking the **Lease History** icon on the List/Add Remote Clusters page. For this manual polling, if the server is in a failover relationship, data is only retrieved for the subnets where the server is the main.

If you have address space privileges (you are assigned the regional-addr-admin role with at least the subnet-utilization and lease-history subroles), you can query the subnet utilization or lease history data by choosing Current Utilization or Lease History from **Operate** menu (see the *"Generating Subnet Utilization History Reports"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*, or the *"Running IP Lease Histories"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*).

Related Topics

[Polling Process, on page 98](#)

[Adjusting the Polling Intervals, on page 99](#)

Adjusting the Polling Intervals

You can adjust the automatic polling interval for subnet utilization and lease history, along with other attributes. These attributes are set in three places at the regional cluster, with the following priority:

1. **Cluster**—These values override the server-wide settings, unless they are unset, in which case the server values are used. The cluster values are set when adding or editing the cluster. In the CLI, set the attributes listed in the table below, using the **cluster** command.
2. **Regional CCM server** (the preset polling interval is 4 hours)—This is set on the Edit CCM Server page, accessible by clicking **Servers**, then the Local CCM Server link. In the CLI, set the attributes listed in the table below using the **ccm** command.



Note If subnet utilization or lease history collection is not explicitly turned on at the local cluster DHCP server (see [Enabling Subnet Utilization Collection, on page 99](#) and [Enabling Lease History Collection, on page 100](#)), no data is collected, even though polling is on by default. Subnet utilization collection at the DHCP server is distinct from polling at the regional cluster, and polling does not automatically trigger collection. Subnet utilization collection must occur before new polling picks up any new data. Because this collection is preset to every 15 minutes, the polling interval should be set higher than this interval (the automatic polling interval is preset to every 4 hours). This also means that subsequent explicit polling performed before the next *collect-addr-util-interval* will not return any new subnet utilization data.

Table 11: Subnet Utilization and Lease History Polling Regional Attributes

Attribute Type	Subnet Utilization	Lease History
Polling interval—How often to poll data	<i>poll-subnet-util-interval</i> 0 (no polling) to 1 year, preset to 4 hours for the CCM server	<i>poll-lease-hist-interval</i> 0 (no polling) to 1 year, preset to 4 hours for the CCM server
Retry interval—How often to retry after an unsuccessful polling	<i>poll-subnet-util-retry</i> 0 to 4 retries	<i>poll-lease-hist-retry</i> 0 to 4 retries
Offset—Hour of the day to guarantee polling	<i>poll-subnet-util-offset</i> 0 to 24h (0h= midnight)	<i>poll-lease-hist-offset</i> 0 to 24h (0h=midnight)

The polling offset attribute ensures that polling occurs at a specific hour of the day, set as 24-hour time, in relation to the polling interval. For example, if you set the interval to 4h and the offset to 6h (6 A.M.), the polling occurs at 2 A.M., 6 A.M., 10 A.M., 2 P.M., 6 P.M., and 10 P.M. each day.

Enabling Subnet Utilization Collection

Step 1 Configure the local cluster DHCP server with scopes and address ranges so that clients have requested leases.

Step 2 Explicitly enable subnet utilization collection. The DHCP server attributes to set are:

- ***collect-addr-util-duration***—Maximum period the DHCP server maintains data. You must change this from the preset value of 0 (no collection) to some reasonable value (see the context sensitive help for this attribute for the impact on memory).

- ***collect-addr-util-interval***—Frequency the server collects snapshots of the data (preset to 15 minutes). How you juggle this value with that of the *collect-addr-util-duration* attribute determines how much memory you use (see the context sensitive help for this attribute).

In the CLI, set the attributes using the **dhcp set** command.

- Step 3** If you are in staged dhcp edit mode, reload the local cluster DHCP server.
- Step 4** At the regional cluster, create the cluster that includes this DHCP server.
- Step 5** Go to the Subnet Utilization Settings section of the List/Add Remote Clusters. Set the attributes described in [Table 11: Subnet Utilization and Lease History Polling Regional Attributes, on page 99](#).
- Step 6** Click **Save**.
- Step 7** In the regional web UI, click the **Subnet History** icon for the cluster to obtain the initial set of subnet utilization data. This data is refreshed automatically at each polling interval. Note that if you subsequently click the **Subnet History** icon, new subnet utilization data does not appear until after the next collection interval (*collect-addr-util-interval*) on the DHCP server (preset to 15 minutes).

Enabling Lease History Collection

- Step 1** Configure the local cluster DHCP server with scopes and address ranges so that clients have requested leases.
- Step 2** Explicitly enable lease history data collection. The DHCP server attributes to set are:
- ***ip-history***—Enable or disable the lease history database for v4-only (DHCPv4), v6-only (DHCPv6), or both.
 - ***ip-history-max-age***—Limit on the age of the history records (preset to 4 weeks).
- In the CLI, set the attributes using the **dhcp set ip-history=<value> (v4-only, v6-only, both, or disable)** command.
- Step 3** If in staged dhcp edit mode, reload the local cluster DHCP server.
- Step 4** At the regional cluster, create the cluster that includes this DHCP server.
- Step 5** In the regional web UI, go to the Lease History Settings section of the List/Add Remote Clusters page.
- Step 6** Set the attributes in [Table 11: Subnet Utilization and Lease History Polling Regional Attributes, on page 99](#).
- Step 7** Click **Save**.
- Step 8** On the List/Add Remote Clusters page, click the **Replica** icon next to the cluster name.
- Step 9** Click the **Lease History** icon for the cluster involved to obtain the initial set of lease history data. This data is refreshed automatically at each polling interval.

Managing DHCP Scope Templates

Scope templates apply certain common attributes to multiple scopes. These common attributes include a scope name based on an expression, policies, address ranges, and an embedded policy options based on an expression. The scope templates you add or pull from the local clusters are visible on the List/Add DHCP Scope Templates page (choose **Scope Templates** from the **Design > DHCPv4** menu).

For details on creating and editing scope templates, and applying them to scopes, see the *"Creating and Applying Scope Templates"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*. The regional cluster web UI has the added feature of pushing scope templates to local clusters and pulling them from local clusters.

Related Topics

[Pushing Scope Templates to Local Clusters, on page 101](#)

[Pulling Scope Templates from Replica Data, on page 101](#)

Pushing Scope Templates to Local Clusters

You can push the scope templates you create from the regional cluster to any of the local clusters. In the web UI, go to the List/Add DHCP Scope Templates page, and do any of the following:

- if you want to push a specific template to a cluster, select the scope template from the Scope Templates pane on the left, and click **Push** (at the top of the page). This opens the Push DHCP Scope Template page.
- If you want to push all of the available scope templates, click the **Push All** icon at the top of the Scope Templates pane. This opens the Push Data to Local Clusters page.

Regional Web UI

The Push DHCP Scope Template page and Push Data to Local Clusters page identify the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure** (preset value)—Ensures that the local cluster has new data without affecting any existing data.
- **Replace**—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for “push all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field.



Tip The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

After making these choices, click **Push Data to Clusters**. This opens the View Push Scope Template Data Report page.

Pulling Scope Templates from Replica Data

You may choose to pull scope templates from the replica data of the local clusters instead of explicitly creating them. (You may first want to update the policy replica data by clicking the **Replicate** icon next to the cluster name.) To pull the scope templates in the regional web UI, click the **Pull Data** icon at the top of the Scope Templates pane.

Regional Web UI

The Select Replica DHCP Scope Template Data to Pull page shows a tree view of the regional server replica data for the local clusters' scope templates. The tree has two levels, one for the local clusters and one for the scope templates in each cluster. You can pull individual scope templates from the clusters, or you can pull all of their scope templates. To pull individual scope templates, expand the tree for the cluster, then click **Pull Scope Template** next to its name. To pull all the scope templates from a cluster, click **Pull All Scope Templates**.

To pull the scope templates, you must also choose a synchronization mode:

- **Ensure**—Ensures that the regional cluster has new data without affecting any existing data.
- **Replace** (preset value)—Replaces data without affecting other objects unique to the regional cluster.
- **Exact**—Available for “pull all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the regional cluster.

Managing DHCP Policies

Every DHCP server must have one or more policies defined for it. Policies define lease duration, gateway routers, and other configuration parameters, in what are called DHCP options. Policies are especially useful if you have multiple scopes, because you need only define a policy once and apply it to the multiple scopes.

For details on creating and editing DHCP policies, and applying them to scopes, see the *"Configuring DHCP Policies" section in Cisco Prime Network Registrar 9.0 DHCP User Guide*. The regional cluster web UI has the added feature of pushing policies to, and pulling them from, the local clusters.

Related Topics

[Pushing Policies to Local Clusters, on page 102](#)

[Pulling Policies from Replica Data, on page 103](#)

Pushing Policies to Local Clusters

You can also push the policies you create from the regional cluster to any of the local clusters. In the regional web UI, go to List/Add DHCP Policies page, and do any of the following:

- If you want to push a specific policy to a cluster, select the policy from the Policies pane on the left, and click **Push** (at the top of the page).
- If you want to push all the policies, click the **Push All** icon at the top of the Policies pane.

Regional Web UI

The Push DHCP Policy Data to Local Clusters page identifies the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure** (preset value)—Ensures that the local cluster has new data without affecting any existing data.
- **Replace**—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for push-all operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field. Then click **Push Data to Clusters** to open the View Push Policy Data Report page.



Tip The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

Pulling Policies from Replica Data

You may choose to pull policies from the replica data of the local clusters instead of explicitly creating them. (In the regional web UI, you may first want to update the policy replica data by clicking the **Replicate** icon next to the cluster name). To pull the policies, click the **Pull Data** icon at the top of the Policies pane.

Regional Web UI

The Select Replica DHCP Policy Data to Pull page shows a tree view of the regional server replica data for the local clusters' policies. The tree has two levels, one for the local clusters and one for the policies in each cluster. You can pull individual policies from the clusters, or you can pull all of their policies. To pull individual policies, expand the tree for the cluster, then click **Pull Policy** next to its name. To pull all the policies from a cluster, click **Pull All Policies**.

To pull all the policies, you must also choose a synchronization mode:

- **Ensure**—Ensures that the regional cluster has new data without affecting any existing data.
- **Replace** (preset value)—Replaces data without affecting other objects unique to the regional cluster.
- **Exact**—Available for “pull all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the regional cluster.

Managing DHCP Client-Classes

Client-classes provide differentiated services to users that are connected to a common network. You can group your user community based on administrative criteria, and then ensure that each user receives the appropriate class of service. Although you can use the Cisco Prime Network Registrar client-class facility to control any configuration parameter, the most common uses are for:

- **Address leases**—How long a set of clients should keep its addresses.
- **IP address ranges**—From which lease pool to assign clients addresses.
- **DNS server addresses**—Where clients should direct their DNS queries.
- **DNS hostnames**—What name to assign clients.
- **Denial of service**—Whether unauthorized clients should be offered leases.

For details on creating and editing client-classes, see the *"Managing Client-Classes and Clients" chapter in Cisco Prime Network Registrar 9.0 DHCP User Guide*. The regional cluster web UI has the added feature of pushing client-classes to, and pulling them from, the local clusters.

Related Topics

[Pushing Client-Classes to Local Clusters, on page 104](#)

[Pushing Client-Classes to Local Clusters, on page 104](#)

Pushing Client-Classes to Local Clusters

You can also push the client-classes you create from the regional cluster to any of the local clusters. In the Regional web UI, go to the List/Add DHCP Client Classes page, and do any of the following:

- If you want to push a specific client-class to a cluster in the web UI, select the client-class from the Client Classes pane on the left, and click **Push** (at the top of the page). This opens the Push DHCP Client Class page.
- If you want to push all the client-classes, click the **Push All** icon at the top of the Client Classes pane. This opens the Push Data to Local Clusters page.

Regional Web UI

The Push DHCP Client Class page and Push Data to Local Clusters page identifies the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure** (preset value)—Ensures that the local cluster has new data without affecting any existing data.
- **Replace**—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for “push all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field. Then click **Push Data to Clusters** to open the View Push Client-Class Data Report page.



Tip The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

Pulling Client-Classes from Replica Data

You may choose to pull client-classes from the replica data of the local clusters instead of explicitly creating them. (In the web UI, you might first want to update the client-class replica data by clicking the **Replicate** icon next to the cluster name.) To pull the client-classes, click the **Pull Data** icon at the top of the Client Classes pane.

Regional Web UI

The Select Replica DHCP Client-Class Data to Pull page shows a tree view of the regional server replica data for the local clusters’ client-classes. The tree has two levels, one for the local clusters and one for the client-classes in each cluster. You can pull individual client-classes from the clusters, or you can pull all of their client-classes. To pull individual client-classes, expand the tree for the cluster, then click **Pull Client-Class** next to its name. To pull all the client-classes from a cluster, click **Pull All Client-Classes**.

To pull the client-classes, you must also choose a synchronization mode:

- **Ensure**—Ensures that the regional cluster has new data without affecting any existing data.
- **Replace** (preset value)—Replaces data without affecting other objects unique to the regional cluster.
- **Exact**—Available for “pull all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the regional cluster.

Managing Virtual Private Networks

A virtual private network (VPN) is a specialized address space identified by a key. A VPN allows address overlap in a network, because the addresses are distinguished by separate keys. Most IP addresses exist in the global address space outside of a VPN. You can create regional VPNs only if you are an administrator assigned the dhcp-management subrole of the central-cfg-admin role.

For details on creating and editing VPNs, and applying them to various network objects, see the "*Configuring Virtual Private Networks Using DHCP*" section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*. The regional web UI has the added feature of pushing VPNs to local clusters and pulling them from local clusters.

Related Topics

[Pushing VPNs to Local Clusters, on page 105](#)

[Pulling VPNs from Replica Data, on page 105](#)

Pushing VPNs to Local Clusters

You can push the VPNs you create from the regional cluster to any of the local clusters. In the Regional web UI, go to the List/Add VPNs page, and do any of the following:

- If you want to push a specific VPN to a cluster in the web UI, select the VPN from the VPNs pane on the left, and click **Push** (at the top of the page). This opens the Push VPN page.
- If you want to push all the VPNs, click the **Push All** icon at the top of the VPNs pane. This opens the Push Data to Local Clusters page.

Regional Web UI

The Push VPN page and Push Data to Local Clusters page identify the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure** (preset value)—Ensures that the local cluster has new data without affecting any existing data.
- **Replace**—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for “push all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field. Then click **Push Data to Clusters** to open the View Push VPN Data Report page.



Tip The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

Pulling VPNs from Replica Data

Instead of explicitly creating VPNs, you can pull them from the local clusters. (In the regional web UI, you may first want to update the VPN replica data by clicking the **Replica** icon next to the cluster name.) To pull

the replica data, click the **Pull Data** icon at the top of the VPNs pane on the left, to open the Select Replica VPN Data to Pull page.

This page shows a tree view of the regional server replica data for the local clusters' VPNs. The tree has two levels, one for the local clusters and one for the VPNs in each cluster. You can pull individual VPNs or you can pull all of them. To pull individual VPNs, expand the tree for the cluster, then click **Pull VPN** next to its name. To pull all the VPNs, click **Pull All VPNs**.

To pull the VPNs, you must choose a synchronization mode:

- **Ensure**—Ensures that the regional cluster has new data without affecting any existing data.
- **Replace** (preset value)—Replaces data without affecting other objects unique to the regional cluster.
- **Exact**—Available for “pull all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the regional cluster.

Managing DHCP Failover Pairs

With DHCP failover, a backup DHCP server can take over for a main server if the latter comes off the network for any reason. You can use failover to configure two servers to operate as a redundant pair. If one server is down, the other server seamlessly takes over so that new DHCP clients can get, and existing clients can renew, their addresses. Clients requesting new leases need not know or care about which server responds to their lease request. These clients can obtain leases even if the main server is down.

In the regional web UI, you can view any created failover pairs on the List/Add DHCP Failover Pairs page. To access this page, click **DHCP**, then **Failover**. This functionality is available only to administrators who are assigned the dhcp-management subrole of the central-cfg-admin role.

For details on creating and editing failover pairs, see the *"Setting Up Failover Server Pairs" section in Cisco Prime Network Registrar 9.0 DHCP User Guide*. The regional cluster web UI has the added feature of pulling addresses from local clusters to create the failover pairs.

To pull the address space for a failover pair, you must have regional-addr-admin privileges.

Regional Web UI

-
- Step 1** On the List/Add DHCP Failover Pairs page or View Unified Address Space page, click the **Pull v4 Data** or **Pull v6 Data** icon in the **Failover Pairs** pane.
 - Step 2** Choose the data synchronization mode (**Update**, **Complete**, or **Exact**) on the Select Pull Replica Address Space page. The results of choosing these modes are described in the table on the page.
 - Step 3** Click the **Report** button in the Synchronize Failover Pair tab and click **Return**.
 - Step 4** Click **Run** on the Report Pull Replica Address Space page.
 - Step 5** Click **OK** on the Run Pull Replica Address Space page.
-

Managing Lease Reservations

You can push lease reservations you create from the regional cluster to any of the local clusters. In the regional cluster web UI, go to the List/Add DHCPv4 Reservations page or List/Add DHCPv6 Reservations page, and click the **Push All** icon in the Reservations pane on the left. Note that you cannot push individual reservations.

If the cluster pushed to is part of a DHCP failover configuration, pushing a reservation also pushes it to the partner server.

Related Topics

[DHCPv4 Reservations, on page 107](#)

[DHCPv6 Reservations, on page 107](#)

DHCPv4 Reservations

To create DHCPv4 reservations, the parent subnet object must exist on the regional server. If there are pending reservation edits at regional, these can be pushed to the subnet local cluster or failover pair. If the subnet has never been pushed, the parent scope is added to the local cluster or pair.

Once a subnet is pushed to a local cluster or pair, reservations are pushed to that cluster or pair. To move the scopes and subnet to another local cluster or failover pair, the subnet must first be reclaimed.

DHCPv6 Reservations

To create DHCPv6 reservations, the parent prefix must exist on the regional server. When there are pending reservation or prefix changes, you can push the updates to the local cluster.

Once a prefix is pushed to a local cluster, it can only update that local cluster. To move the prefix to another local cluster, it must first be reclaimed.

Regional Web UI

The ensuing page identifies the data to push, how to synchronize it with the local cluster, and the cluster or clusters to which to push it. The data synchronization modes are:

- **Ensure**—Ensures that the local cluster has new data without affecting any existing data.
- **Replace** (preset value)—Replaces data without affecting other objects unique to the local cluster.
- **Exact**—Available for “push all” operations only. Use this with caution, because it overwrites the data and deletes any other objects unique to the local cluster.

Choose the destination cluster or clusters in the Available field and move it or them to the Selected field.



Tip The synchronization mode and cluster choice settings are persistent for the duration of the current login session, so that they are in effect each time you access this page, unless you change them.

After making these choices, click **Push Data to Clusters**. This opens the View Push Reservations Data Report page. Click **OK** on this page.

You can also pull the replica address space on the List/Add DHCP v6 Reservations page, and opt whether to omit reservations when doing so. You should use this option only to reduce processing time when you are sure that there are no pending changes to reservations to merge. To omit reservations for the pull, check the *Omit Reservations?* check box, then click **Pull Data**.

See the *"Managing DHCPv6 Addresses"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*.

Monitoring Resource Limit Alarms

Resource limit alarms enable you to monitor Cisco Prime Network Registrar system resources and provide an indication when one or more product resources has entered potentially dangerous level and requires attention. Resource limit alarms are designed to convey the resource limit information in an organized and consolidated way.



Note The log messages related to resource limits are logged to the `ccm_monitor_log` files. For more information on log files, see [Log Files, on page 132](#).

You can reset the predefined threshold levels for both critical and warning levels for each monitored resource.

Cisco Prime Network Registrar reports the current status, the current value, and the peak value of the monitored resources in the web UI and CLI. The peak value is compared to the configured warning or critical limit for the resource limit alarm and the status of the resource limit alarm is displayed as OK, Warning, or Critical. Cisco Prime Network Registrar displays the alarms on the web UI and CLI until the resulting condition no longer occurs and the peak value is reset.

The resource limit alarms are updated at regular intervals based on the polling interval you configure. For more information on setting up the polling interval, see [Setting Resource Limit Alarms Polling Interval, on page 109](#).

If SNMP traps are enabled for the resource limit alarms, Cisco Prime Network Registrar generates SNMP traps when the monitored resources exceed the critical or warning levels. SNMP traps are generated whenever the current value exceeds the configured warning or critical level.

The resource limit alarms can be configured both at the regional and in the local cluster. The resource limit alarms data is consolidated at the individual local cluster level. The resource limits alarms available on the regional cluster level pertain to only the regional cluster. The table below lists the types of resource limit alarms that are available on the regional or the local cluster.

Table 12: Resource Limit Alarms

	Regional Cluster	Local Cluster
Data Free Space in../Data Partition	✓	✓
Shadow Backup Time	✓	✓
CCM Memory	✓	✓
CNR Server Agent Memory	✓	✓
Tomcat Memory	✓	✓
DHCP Memory	x	✓
CDNS Memory	x	✓
DNS Memory	x	✓
SNMP Memory	✓	✓

TFTP Memory	x	✓
Lease Count	x	✓
Zone Count	x	✓
Resource Records Count	x	✓

Configuring Resource Limit Alarm Thresholds

You can configure the warning and critical limits for the resource limit alarms using the **Edit CCM Server** page.

Local and Regional Web UI

-
- Step 1** To access the CCM server properties, choose **Manage Servers** under the **Operate** menu to open the Manage Servers page.
- Step 2** Click **Local CCM Server** in the Manage Servers pane on the left. The Edit Local CCM Server page appears. This page displays all the CCM server attributes.
- Step 3** Click the **Configure Resource Limits** tab.
- Step 4** Modify the settings as per your requirement.
- Note** To enable the SNMP traps for the resource limit alarms, select the Enable Traps option in the Trap Configuration group.
- Step 5** Click **Save** to save the CCM server attribute modifications.
-

CLI Commands

To set the resource limit alarms on the local or regional cluster, use **resource set attribute=value**. Use **resource show** to review the current setting and use **resource report [all | full | levels]** command to report on the resources.

To view the defined warning and critical levels, use **resource report levels** command.

A 109 status message is reported (if at least one resource is in the critical or warning state) under the following scenarios.

- Execute **resource report** command.
- Connect to a cluster via CLI.
- Exit from CLI.

Setting Resource Limit Alarms Polling Interval

You can set how often Cisco Prime Network Registrar polls for alarm data from the server and updates the web UI data. The *stats-history-sample-interval* controls the CCM server system polling rate.

-
- Step 1** To edit the alarm poll interval, you need to edit the user preferences by going to **User Preferences** under the Settings drop-down list (at the top of the main page).
- Step 2** After making the user preference settings, click **Modify User Preferences**.
-

Viewing Resource Limit Alarms

Resource limit alarms are displayed on the Alarms toolbar. To see a summary of the alarms, in the Cisco Prime Network Registrar web UI, click the **Alarms** toolbar on the bottom of the web UI. This opens the Alarms toolbar overlay which displays the status, resource values (current, configured warning, and critical value), and the peak value for each resource limit alarm. Based on the peak value for each resource limit, the status of resource limit is displayed as OK, Warning, or Critical on the web UI and CLI. The alarms are updated at regular intervals based on the polling interval you configure. For more information on setting up the polling interval, see [Setting Resource Limit Alarms Polling Interval, on page 109](#).



Note When a resource is in a warning or critical state, the resource limit alarm is also displayed on the Configuration Summary page.

Resetting Resource Limit Alarms Peak Value

Cisco Prime Network Registrar maintains the peak values for each resource limit. The peak value is updated only when the current value exceeds the peak value. The peak value is compared to the configured warning or critical limit for the resource limit alarm and the status of the resource limit alarm is displayed as OK, Warning, or Critical.

When the peak value exceeds the configured warning or critical limit the status of the resource limit alarm is shown as Warning or Critical (on the web UI and CLI) respectively until the peak value is explicitly reset. To reset the peak value, perform the following steps:

-
- Step 1** On the **Alarms** toolbar, select the Alarm for which you want to reset the peak value.
- Step 2** Click **Reset Alarm** to clear the peak value.
-

CLI Commands

To reset the peak value on the local or regional cluster, use **resource reset** [*name* [,*name* [...]]].



Note If no resource name is provided, all are reset.

Export Resource Limit Alarms Data

You can export the resource limit alarms data to a CSV file. To export the resource limit alarms:

-
- Step 1** Select the alarm in the Alarms toolbar at the top of the web UI.
- Step 2** Click **Export to CSV**.
- Step 3** The File Download pop-up window displays. Click **Save**.
- Step 4** In the Save As pop-up window, choose the location you want to save the file to and click **Save**.
-

Local Cluster Management Tutorial

This tutorial describes a basic scenario on a local cluster of the Example Company. Administrators at the cluster are responsible for users, zone data, DHCP data, address space data, and the servers in general. The task is to set up two zones (example.com and boston.example.com), hosts in the zones, and a subnet. The local cluster must also create a special administrator account so that the regional cluster in San Jose can perform the central configuration and replicate the local cluster administrators and address space at another cluster, as described in [Regional Cluster Management Tutorial, on page 117](#).

Related Topics

- [Administrator Responsibilities and Tasks, on page 111](#)
- [Create the Administrators, on page 112](#)
- [Create the Address Infrastructure, on page 112](#)
- [Create the Zone Infrastructure, on page 113](#)
- [Create a Host Administrator Role with Constraints, on page 115](#)
- [Create a Group to Assign to the Host Administrator, on page 116](#)
- [Test the Host Address Range, on page 117](#)

Administrator Responsibilities and Tasks

The local cluster administrators have the following responsibilities and tasks:

- **example-cluster-admin**—Created by the superuser:
 - At the Boston cluster, creates the other local administrators (example-zone-admin and example-host-admin).
 - Creates the basic network infrastructure for the local clusters.
 - Constrains the example-host-role to an address range in the boston.example.com zone.
 - Creates the example-host-group (defined with the example-host-role) that the example-zone-admin will assign to the example-host-admin.
- **example-zone-admin**:
 - Creates the example.com and boston.example.com zones, and maintains the latter zone.
 - Assigns the example-host-group to the example-host-admin.
- **example-host-admin**—Maintains local host lists and IP address assignments.

Create the Administrators

For this example, the superuser in Boston creates the local cluster, zone, and host administrators, as described in the [Administrator Responsibilities and Tasks, on page 111](#).

Local Basic Web UI

- Step 1** At the Boston local cluster, log in as superuser (usually **admin**).
- Step 2** In Basic mode, from the **Administration** menu, choose **Administrators**.
- Step 3** Add the local cluster administrator (with superuser access)—On the List/Add Administrators page:
- Click the **Add Administrators** icon in the Administrators pane, enter **example-cluster-admin** in the Name field.
 - Enter **exampleadmin** in the Password and Confirm Password fields, then click **Add Admin**.
 - Check the Superuser check box.
 - Do not choose a group from the Groups list.
 - Click **Save**.
- Step 4** Add the local zone administrator on the same page:
- Click the **Add Administrators** icon in the Administrators pane, enter **example-zone-admin** in the Name field, and **examplezone** in the Password field, then click **Add Admin**.
 - Multiselect **ccm-admin-group**, **dns-admin-group**, and **host-admin-group** in the Groups drop-down list. The **dns-admin-group** is already predefined with the **dns-admin** role to administer DNS zones and servers. The **ccm-admin-group** guarantees that the **example-zone-admin** can set up the **example-host-admin** with a constrained role later on. The **host-admin-group** is mainly to test host creation in the zone.
 - Click **Save**.
- Step 5** Add the local host administrator on the same page:
- Click the **Add Administrators** icon in the Administrators pane, enter **example-host-admin** in the Name field, and **examplehost** in the Password field, then click **Add Admin**.
 - Do not choose a group at this point. (The **example-zone-admin** will later assign **example-host-admin** to a group with a constrained role.)
 - Click **Save**.

Note For a description on how to apply constraints to the administrator, see the [Create a Host Administrator Role with Constraints, on page 115](#).

Create the Address Infrastructure

A prerequisite to managing the zones and hosts at the clusters is to create the underlying network infrastructure. The network configuration often already exists and was imported. However, this tutorial assumes that you are starting with a clean slate.

The local **example-cluster-admin** next creates the allowable address ranges for the hosts in the **boston.example.com** zone that will be assigned static IP addresses. These addresses are in the **192.168.50.0/24** subnet with a range of hosts from 100 through 200.

Local Advanced Web UI

- Step 1** At the local cluster, log out as superuser, then log in as the **example-cluster-admin** user with password **exampleadmin**. Because the administrator is a superuser, all features are available.
- Step 2** Click **Advanced** to go to Advanced mode.
- Step 3** Click **Design**, then **Subnets** under DHCPv4 submenu.
- Step 4** On the List/Add Subnets page, enter the boston.example.com subnet address:
- Click the **Add Subnets** icon in the Subnets pane, enter **192.168.50** in the Address field.
 - Choose **24** in the mask drop-down list—This subnet will be a normal Class C network.
 - Leave the Owner, Region, and Address Type fields as is. Add description if desired.
 - Click **Add Subnet**.
- Step 5** Click the 192.168.50.0/24 address to open the Edit Subnet page.
- Step 6** In the IP Ranges fields, enter the static address range:
- Enter **100** in the Start field. Tab to the next field.
 - Enter **200** in the End field.
 - Click **Add IP Range**. The address range appears under the fields.
- Step 7** Click **Save**.
- Step 8** Click **Address Space** to open the View Unified Address Space page. The 192.168.50.0/24 subnet should appear in the list. If not, click the **Refresh** icon.
-

Create the Zone Infrastructure

For this scenario, example-cluster-admin must create the Example Company zones locally, including the example.com zone and its subzones. The example-cluster-admin also adds some initial host records to the boston.example.com zone.

Related Topics

- [Create the Forward Zones, on page 113](#)
- [Create the Reverse Zones, on page 114](#)
- [Create the Initial Hosts, on page 114](#)

Create the Forward Zones

First, create the example.com and boston.example.com forward zones.

Local Basic Web UI

- Step 1** At the local cluster, log in as the **example-zone-admin** user with password **examplezone**.
- Step 2** From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu. This opens the List/Add Forward Zones page.
- Step 3** Create the example.com zone (tab from field to field):
- Click the **Add Forward Zone** icon in the Forward Zones pane, enter **example.com** in the Name field.

- b) In the Nameserver FQDN field, enter **ns1**.
- c) In the Contact E-Mail field, enter **hostmaster**.
- d) In the Serial Number field, enter the serial number.
- e) Click **Add Zone**.

- Step 4** Create the **boston.example.com** zone in the same way, using the same values as in the previous steps:
- a) Creating a zone with a prefix added to an existing zone opens the Create Subzone in Parent Zone page, because the zone can be a potential subzone. Because you do want to create this zone as a subzone to example.com, click **Create as Subzone** on the Create Subzone in Parent Zone page.
 - b) Because nameservers are different in each zone, you must create a glue Address (A) record to tie the zones together. Enter 192.168.50.1 in the A record field, then click **Specify Glue Records**. Then click **Report, Run, and Return**.
 - c) The List/Add Zones page should now list example.com and boston.example.com.
- Step 5** Click **Advanced**, then **Show Forward Zone Tree** to show the hierarchy of the zones. Return to list mode by clicking **Show Forward Zone List**.

Create the Reverse Zones

Next, create the reverse zones for example.com and boston.example.com. This way you can add reverse address pointer (PTR) records for each added host. The reverse zone for example.com is based on the 192.168.50.0 subnet; the reverse zone for boston.example.com is based on the 192.168.60.0 subnet.

Local Basic Web UI

- Step 1** At the local cluster, you should be logged in as the example-zone-admin user, as in the previous section.
- Step 2** From the **Design** menu, choose **Reverse Zones** under the **Auth DNS** submenu.
- Step 3** On the List/Add Reverse Zones page, click the **Add Reverse Zone** icon in the Reverse Zones pane, enter **50.168.192.in-addr.arpa** in the Name field. (There is already a reverse zone for the loopback address, 127.in-addr.arpa.)
- Step 4** Enter the required fields to create the reverse zone, using the forward zone values:
- a) **Nameserver**—Enter **ns1.example.com**. (be sure to include the trailing dot).
 - b) **Contact E-Mail**—Enter **hostmaster.example.com**. (be sure to include the trailing dot).
 - c) **Serial Number**—Enter the serial number.
- Step 5** Click **Add Reverse Zone** to add the zone and return to the List/Add Reverse Zones page.
- Step 6** Do the same for the boston.example.com zone, using **60.168.192.in-addr.arpa** as the zone name and the same nameserver and contact e-mail values as in **Step 4**. (You can cut and paste the values from the table.)

Create the Initial Hosts

As a confirmation that hosts can be created at the Boston cluster, the example-zone-admin tries to create two hosts in the example.com zone.

Local Advanced Web UI

- Step 1** As the example-zone-admin user, click **Advanced** to enter Advanced mode.

- Step 2** From the **Design** menu, choose **Hosts** under the **Auth DNS** submenu. This opens the List/Add Hosts for Zone page. You should see `boston.example.com` and `example.com` in the Select Zones box on the left side of the window.
- Step 3** Click `example.com` in the list of zones.
- Step 4** Add the first static host with address `192.168.50.101`:
- Enter **userhost101** in the Name field.
 - Enter the complete address **192.168.50.101** in the IP Address(es) field. Leave the IPv6 Address(es) and Alias(es) field blank.
 - Ensure that the Create PTR Records? check box is checked.
 - Click **Add Host**.
- Step 5** Add the second host, **userhost102**, with address **192.168.50.102**, in the same way. The two hosts should now appear along with the nameserver host on the List/Add Hosts for Zone page.

Create a Host Administrator Role with Constraints

In this part of the tutorial, the Boston `example-cluster-admin` creates the `example-host-role` with address constraints in the `boston.example.com` zone.

Local Advanced Web UI

- Step 1** Log out as the `example-zone-admin` user and log in as the **example-cluster-admin** user (with password **exampleadmin**).
- Step 2** Click **Advanced** to enter Advanced mode.
- Step 3** From the **Administration** menu, choose **Roles** under **User Access** submenu to open the List/Add Administrator Roles page.
- Step 4** Add the `example-host-role`:
- Click the **Add Role** icon in the Roles pan to open the Add Roles dialog box.
 - Enter **example-host-role** in the Name field.
 - Click **Add Role**. The `example-host-role` should now appear in the list of roles on the List/Add Administrator Roles page.
- Step 5** Add the constraint for the role:
- Click **Add Constraint**.
 - On the Add Role Constraint for Role page, scroll down to Host Restrictions.
 - For the `all-forward-zones` attribute, click the **false** radio button.
 - For the `zones` attribute, enter **boston.example.com**.
 - For the `ipranges` attribute, enter the range **192.168.50.101–192.168.50.200**.
 - The `zone-regexp` and `host-regexp` attribute fields are for entering regular expressions to match zones and hosts, respectively, in regex syntax. (See the following table for the commonly used regex values.)

Table 13: Common Regex Values

Value	Matches
<code>.</code> (dot)	Any character (a wildcard). Note that to match a literal dot character (such as in a domain name), you must escape it by using a backslash (<code>\</code>), such that <code>\.com</code> matches <code>.com</code> .

Value	Matches
<code>\char</code>	Literal character (<i>char</i>) that follows, or the <i>char</i> has special meaning. Used especially to escape metacharacters such as the dot (.) or another backslash. Special meanings include <code>\d</code> to match decimal digits, <code>\D</code> for nondigits, <code>\w</code> for alphanumerics, and <code>\s</code> for whitespace.
<code>char?</code>	Preceding <i>char</i> once or not at all, as if the character were optional. For example, <code>example\?.?com</code> matches <code>example.com</code> or <code>examplecom</code> .
<code>char*</code>	Preceding <i>char</i> zero or more times. For example, <code>ca*t</code> matches <code>ct</code> , <code>cat</code> , and <code>caaat</code> . This repetition metacharacter does iterative processing with character sets (see <code>[charset]</code>).
<code>char+</code>	Preceding <i>char</i> one or more times. For example, <code>ca+t</code> matches <code>cat</code> and <code>caaat</code> (but not <code>ct</code>).
<code>[charset]</code>	Any of the characters enclosed in the brackets (a character set). You can include character ranges such as <code>[a-z]</code> (which matches any lowercase character). With the <code>*</code> repetition metacharacter applied, the search engine iterates through the set as many times as necessary to effect a match. For example, <code>a[bcd]*b</code> will find <code>abcdb</code> (by iterating through the set a second time). Note that many of the metacharacters (such as the dot) are inactive and considered literal inside a character set.
<code>[^charset]</code>	Anything but the <i>charset</i> , such that <code>[^a-zA-Z0-9]</code> matches any nonalphanumeric character (which is equivalent to using <code>\W</code>). Note that the caret outside a character set has a different meaning.
<code>^</code>	Beginning of a line.
<code>\$</code>	End of a line.

g) Click **Add Constraint**. The constraint should have an index number of 1.

Step 6 Click **Save**.

Create a Group to Assign to the Host Administrator

The Boston `example-cluster-admin` next creates an `example-host-group` that includes the `example-host-role` so that the `example-zone-admin` can assign this group to the `example-host-admin`.

Local Advanced Web UI

- Step 1** As example-cluster-admin, still in Advanced mode, from the **Administration** menu, choose **Groups** submenu to open the List/Add Administrator Groups page.
- Step 2** Create the example-host-group and assign the example-host-role to it:
- Click the **Add Groups** icon in the Groups pane, enter **example-host-group** in the Name field.
 - From the Base Role drop-down list, choose **example-host-role**.
 - Click **Add Group**.
 - Add a description such as **Group for the example-host-role**, then click **Save**.
- Step 3** Log out as example-cluster-admin, then log in as the **example-zone-admin** user (with password **examplezone**).
- Step 4** As example-zone-admin, assign the example-host-group to the example-host-admin:
- In Basic mode, from the **Administration** menu, choose **Administrators**.
 - On the List/Add Administrators page, click example-host-admin to edit the administrator.
 - On the Edit Administrator page, choose **example-host-group** in the Available list, then click << to move it to the Selected list.
 - Click **Save**. The example-host-admin should now show the example-host-group in the Groups column on the List/Add Administrators page.
-

Test the Host Address Range

The example-host-admin next tests an out-of-range address and then adds an acceptable one.

Local Advanced Web UI

- Step 1** At the local cluster, log out as example-zone-admin, then log in as **example-host-admin** (with password **examplehost**).
- Step 2** Click **Advanced** to enter Advanced mode.
- Step 3** From the **Design** menu, choose **Hosts** from the **Auth DNS** submenu.
- Step 4** On the List/Add Hosts for Zone page, try to enter an out-of-range address (note the range of valid addresses in the Valid IP Ranges field):
- Enter **userhost3** in the Name field.
 - Deliberately enter an out-of-range address (**192.168.50.3**) in the IP Address(es) field.
 - Click **Add Host**. You should get an error message.
- Step 5** Enter a valid address:
- Enter **userhost103**.
 - Enter **192.168.50.103** in the IP Address(es) field.
 - Click **Add Host**. The host should now appear with that address in the list.
-

Regional Cluster Management Tutorial

This tutorial is an extension of the scenario described in the [Local Cluster Management Tutorial, on page 111](#). In the regional cluster tutorial, San Jose has two administrators—a regional cluster administrator and a central

configuration administrator. Their goal is to coordinate activities with the local clusters in Boston and Chicago so as to create DNS zone distributions, router configurations, and DHCP failover configurations using the servers at these clusters. The configuration consists of:

- One regional cluster machine in San Jose.
- Two local cluster machines, one in Boston and one in Chicago.
- One Cisco uBR7200 router in Chicago.

Related Topics

[Administrator Responsibilities and Tasks, on page 118](#)

[Create the Regional Cluster Administrator, on page 118](#)

[Create the Central Configuration Administrator, on page 119](#)

[Create the Local Clusters, on page 119](#)

[Add a Router and Modify an Interface, on page 120](#)

[Add Zone Management to the Configuration Administrator, on page 121](#)

[Create a Zone for the Local Cluster, on page 121](#)

[Pull Zone Data and Create a Zone Distribution, on page 122](#)

[Create a Subnet and Pull Address Space, on page 122](#)

[Push a DHCP Policy, on page 123](#)

[Create a Scope Template, on page 123](#)

[Create and Synchronize the Failover Pair, on page 124](#)

Administrator Responsibilities and Tasks

The regional administrators have the following responsibilities and tasks:

- **example-regional-admin**—Created by the superuser at the San Jose regional cluster, who creates the example-cfg-admin.
- **example-cfg-admin**:
 - Defines the Boston and Chicago clusters and checks connectivity with them.
 - Adds a router and router interfaces.
 - Pulls zone data from the local clusters to create a zone distribution.
 - Creates a subnet and policy, and pulls address space, to configure DHCP failover pairs in Boston and Chicago.

Create the Regional Cluster Administrator

The regional superuser first creates the example-regional-administrator, defined with groups, to perform cluster and user administration.

Regional Web UI

- Step 1** Log into the regional cluster as superuser.
 - Step 2** From the **Administration** menu, choose **Administrators** to open the List/Add Administrators page for the local cluster version of this page, which is essentially identical.
 - Step 3** Click the **Add Administrators** icon in the Administrators pane, enter **example-regional-admin** in the Name field, then **examplereg** in the Password field in the Add Administrator dialog box, then click **Add Administrator**.
 - Step 4** Multiselect **central-cfg-admin-group** (for cluster administration) and **regional-admin-group** (for user administration) in the Groups drop-down list.
 - Step 5** Click **Save**.
-

Create the Central Configuration Administrator

As part of this tutorial, the example-regional-admin next logs in to create the example-cfg-admin, who must have regional configuration and address management capabilities.

Regional Web UI

- Step 1** Log out as superuser, then log in as **example-regional-admin** with password **examplereg**. Note that the administrator has all but host and address space administration privileges.
 - Step 2** From the **Administration** menu, choose **Administrators** to open the List/Add Administrators page.
 - Step 3** Click the **Add Administrators** icon in the Administrators pane, enter **example-cfg-admin** in the Name field, then **cfgadmin** in the Password field in the Add Administrator dialog box, then click **Add Administrator**.
 - Step 4** Multiselect **central-cfg-admin-group** and **regional-addr-admin-group** in the Groups drop-down list.
 - Step 5** Click **Save**. The example-cfg-admin now appears with the two groups assigned.

You can also add constraints for the administrator. Click **Add Constraint** and, on the Add Role Constraint for Role page, choose the read-only, owner, or region constraints, then click **Add Constraint**.
-

Create the Local Clusters

The example-cfg-admin next creates the two local clusters for Boston and Chicago.

Regional Web UI

- Step 1** Log out as example-regional-admin, then log in as **example-cfg-admin** with password **cfg admin**.
- Step 2** From the **Operate** menu, choose **Manage Clusters** from the **Servers** submenu to open the List/Add Remote Clusters page.
- Step 3** Click the **Add Manage Clusters** icon in the **Manage Clusters** pane.
- Step 4** On the Add Cluster dialog box, create the Boston cluster based on data provided by its administrator:
 - a) Enter **Boston-cluster** in the name field.
 - b) Enter the IPv4 address of the Boston server in the IPv4 Address field.

- c) Enter the IPv6 address of the Boston server in the IPv6 Address field.
- d) Enter **example-cluster-admin** in the admin field, then **exampleadmin** in the password field.
- e) Enter in the SCPO-port field the SCP port to access the cluster as set at installation (**1234** is the preset value).
- f) Click **Add Cluster**.

- Step 5** Create the Chicago cluster in the same way, except use **Chicago-cluster** in the name field, enter the remaining values based on data provided by the Chicago administrator, then click **Add Cluster**. The two clusters should now appear on the List/Add Remote Clusters page.
- Step 6** Connect to the Boston cluster. Click the **Go Local** icon next to Boston-cluster. If this opens the local cluster Manage Servers page, this confirms the administrator connectivity to the cluster. To return to the regional cluster web UI, click the **Go Regional** icon.
- Step 7** Connect to the Chicago cluster to confirm the connectivity in the same way.
- Step 8** Confirm that you can replicate data for the two forward zones from the Boston cluster synchronization:
- a) From the **Operate** menu, choose **Replica Data** from the **Servers** submenu.
 - b) On the View Replica Class List page, click Boston-cluster in the Select Cluster list.
 - c) In the Select Class list, click **Forward Zones**.
 - d) Click the **Replicate** icon in the Replicate Data column.
 - e) Click **View Replica Class List**. On the List Replica Forward Zones for Cluster page, you should see the boston.example.com and example.com zones.

Add a Router and Modify an Interface

The example-cfg-admin next takes over at the regional cluster to add a router and modify one of its interfaces to configure the DHCP relay agent. Add the subnets manually.

Regional Advanced Web UI

- Step 1** As example-cfg-admin, from the **Deploy** menu, choose **Router List** from the **Router Configuration** submenu.
- Step 2** On the List/Add Routers page, click the **Add Router** icon in the Router List pane.
- Step 3** On the Add Router dialog box, add the router based on data from its administrator:
- a) Give the router a distinguishing name in the name field. For this example, enter **router-1**.
 - b) Enter the router description in the description field.
 - c) Enter the management interface address for the router in the address field.
 - d) Enter the IPv6 management interface address for the router in the ip6address field.
 - e) Choose a owner and a region.
 - f) Click **Add Router**. The router should now appear on the List/Add Routers page.
- Step 4** Confirm that the router is created. Click **Router Tree** to view the hierarchy of router interfaces for router-1 on the View Tree of Routers page.
- Step 5** Configure a DHCP relay agent for the router:
- a) Create a new interface for the router.
 - b) Click the interface names on the View Tree of Routers page to open the Edit Router Interface page. (Alternatively, from the List/Add Routers page, click the **Interfaces** icon associated with the router, then click the interface name on the List Router Interfaces for Router page.)
 - c) On the Edit Router Interface page, enter the IP address of the DHCP server in the ip-helper field.

d) Click **Save** at the bottom of the page.

Step 6 Confirm with the router administrator that the DHCP relay agent was successfully added.

Add Zone Management to the Configuration Administrator

Because there are no zones set up at the Chicago cluster, the example-cfg-admin can create a zone at the regional cluster to make it part of the zone distribution. However, the example-regional-admin must first modify the example-cfg-admin to be able to create zones.

Regional Web UI

Step 1 Log out as example-cfg-admin, then log in as **example-regional-admin**.

Step 2 From the **Administration** menu, choose **Administrators** under the **User Access** submenu.

Step 3 On the List/Add Administrators page, click example-cfg-admin from the Administrators pane.

Step 4 On the Edit Administrator page, click central-dns-admin-group in the Groups Available list, then move it (using <<) to the Selected list. The Selected list should now have central-cfg-admin-group, regional-addr-admin-group, and central-dns-admin-group.

Step 5 Click **Save**. The change should be reflected on the List/Add Administrators page.

Create a Zone for the Local Cluster

The example-cfg-admin next creates the chicago.example.com zone for the zone distribution with the Boston and Chicago zones.

Regional Web UI

Step 1 Log out as example-regional-admin, then log in as **example-cfg-admin**.

Step 2 From the **Design** menu, choose **Forward Zones** under the **Auth DNS** submenu.

Step 3 Click the **Add Forward Zone** icon in the **Forward Zones** pane.

Step 4 On the Add DNS Zone dialog box, enter:

- a) **Name**—chicago.example.com.
- b) **Nameserver FQDN**—ns1.
- c) **Contact E-mail**—hostmaster.
- d) **Nameservers**—ns1 (click **Add Nameserver**).
- e) Click **Add DNS Zone**.

Step 5 Click the **Reverse Zones** submenu.

Step 6 On the List/Add Reverse Zones page, create the **60.168.192.in-addr.arpa** reverse zone for the Chicago zone, with the proper attributes set.

Pull Zone Data and Create a Zone Distribution

The example-cfg-admin next pulls zone data from Boston and Chicago and creates a zone distribution.

Regional Web UI

-
- Step 1** As example-cfg-admin, from the **Design** menu, choose **Views** under the **Auth DNS** submenu to view the List/Add Zone Views page.
- Step 2** On the List/Add Zone Views page, pull the zone from the replica database:
- Click the **Pull Data** icon in the **Views** pane.
 - On the Select Replica DNS View Data to Pull dialog box, leave the Data Synchronization Mode defaulted as Update, then click **Report** to open the Report Pull Replica Zone Data page.
 - Notice the change sets of data to pull, then click **Run**.
 - On the Run Pull Replica Zone Data page, click **OK**.
- Step 3** On the List/Add Zone Views page, notice that the Boston cluster zone distribution is assigned an index number (**1**) in the Name column. Click the number.
- Step 4** On the Edit Zone Views page, in the Primary Server field, click Boston-cluster. (The IP address of the Boston-cluster becomes the first master server in the Master Servers list.)
- Step 5** Because we want to make the Chicago-cluster DNS server a secondary server for the Boston-cluster:
- Click **Add Server** in the Secondary Servers area.
 - On the Add Zone Distribution Secondary Server page, choose **Chicago-cluster** in the Secondary Server drop-down list.
 - Click **Add Secondary Server**.
- Step 6** On the Edit Zone Distribution page, in the Forward Zones area, move **chicago.example.com** to the Selected list.
- Step 7** In the Reverse Zones area, move **60.168.192.in-addr.arpa** to the Selected list.
- Step 8** Click **Modify Zone Distribution**.
-

Create a Subnet and Pull Address Space

The example-cfg-admin next creates a subnet at the regional cluster. This subnet will be combined with the other two pulled subnets from the local clusters to create a DHCP failover server configuration.

Regional Web UI

-
- Step 1** As example-cfg-admin, from the **Design** menu, choose **Subnets** under the **DHCPv4** submenu to open the List/Add Subnets page. You should see the subnets created by adding the router (in the [Add a Router and Modify an Interface, on page 120](#)).
- Step 2** Create an additional subnet, 192.168.70.0/24 by clicking the **Add Subnets** icon in the Subnets pane:
- Enter **192.168.70** (the abbreviated form) as the subnet network address in the Address/Mask field.
 - Leave the **24** (255.255.255.0) selected as the network mask.
 - Click **Add Subnet**.
- Step 3** Click **Address Space** to confirm the subnet you created.
- Step 4** On the View Unified Address Space page, click **Pull Replica Address Space**.

- Step 5** On the Select Pull Replica Address Space page, leave everything defaulted, then click **Report**.
- Step 6** The Report Pull Replica Address Space page should show the change sets for the two subnets from the clusters. Click **Run**.
- Step 7** Click **OK**. The two pulled subnets appear on the List/Add Subnets page.
-

Push a DHCP Policy

The example-cfg-admin next creates a DHCP policy, then pushes it to the local clusters.

Regional Web UI

- Step 1** As example-cfg-admin, from the **Design** menu, choose **Policies** under the **DHCP Settings** submenu.
- Step 2** On the List/Add DHCP Policies page, click the **Add Policies** icon in the **Policies** pane.
- Step 3** On the Add DHCP Policy dialog box, create a central policy for all the local clusters:
- Enter **central-policy-1** in the Name field. Leave the Offer Timeout and Grace Period values as is.
 - Enter a lease period. In the DHCP > DHCPv4 > Options drop-down list, choose **dhcp-lease-time [51] (unsigned time)**, then enter **2w** (two weeks) for the lease period in the Value field.
 - Click **Add Option**.
 - Click **Add Policy**. The central-policy-1 should appear on the List/Add DHCP Policies page.
- Step 4** Push the policy to the local clusters:
- Select the policy, central-policy-1 and click the **Push** button.
 - On the Push DHCP Policy to Local Clusters page, leave the Data Synchronization Mode as **Ensure**. This ensures that the policy is replicated at the local cluster, but does not replace its attributes if a policy by that name already exists.
 - Click **Select All** in the Destination Clusters section of the page.
 - Click << to move both clusters to the Selected field.
 - Click **Push Data to Clusters**.
 - View the push operation results on the View Push DHCP Policy Data Report page, then click **OK**.
-

Create a Scope Template

The example-cfg-admin next creates a DHCP scope template to handle failover server pair creation.

Regional Web UI

- Step 1** As the example-cfg-admin user, from the **Design** menu, choose **Scope Templates** under the **DHCPv4** submenu.
- Step 2** On the List/Add DHCP Scope Templates page, click the **Add Scope Templates** icon in the **Scope Templates** pane. Enter **scope-template-1** in the Name field, then click **Add DHCP Scope Template**.
- Step 3** The template should appear on the List/Add DHCP Scope Templates page. Set the basic properties for the scope template—Enter or choose the following values in the fields:

- a) **Scope Name Expression**—To autogenerate names for the derivative scopes, concatenate the example-scope string with the subnet defined for the scope. To do this, enter (**concat “example-scope-” subnet**) in the field (including the parentheses).
- b) **Policy**—Choose **central-policy-1** in the drop-down list.
- c) **Range Expression**—Create an address range based on the remainder of the subnet (the second through last address) by entering (**create-range 2 100**).
- d) **Embedded Policy Option Expression**—Define the router for the scope in its embedded policy and assign it the first address in the subnet by entering (**create-option “routers” (create-ipaddr subnet 1)**).

Step 4 Click **Save**.

Create and Synchronize the Failover Pair

The example-cfg-admin next creates the failover server pair relationship and synchronizes the failover pair. The DHCP server at Boston becomes the main, and the server at Chicago becomes the backup.

Regional Web UI

Step 1 As the example-cfg-admin user, from the **Deploy** menu, choose **Failover Pairs** under the **DHCP** submenu.

Step 2 On the List/Add DHCP Failover Pairs page, click the **Add Failover Pair** icon in the **Failover Pairs** pane.

Step 3 On the Add DHCP Failover Pair dialog box, enter or choose the following values:

- a) **Failover Pair Name**—Enter **central-fo-pair**.
- b) **Main Server**—Click **Boston-cluster**.
- c) **Backup Server**—Click **Chicago-cluster**.
- d) **Scope Template**—Click **scopetemplate-1**
- e) Click **Add Failover Pair**.

Step 4 Synchronize the failover pair with the local clusters:

- a) On the List/Add DHCP Failover Pairs page, click the **Report** icon in the Synchronize column.
- b) On the Report Synchronize Failover Pair page, accept **Local Server** as the source of network data.
- c) Accept **Main to Backup** as the direction of synchronization.
- d) Accept the operation **Update**.
- e) Click **Report** at the bottom of the page.
- f) On the View Failover Pair Sync Report page, click **Run Update**.
- g) Click **Return**.

Step 5 Confirm the failover configuration and reload the server at the Boston cluster:

- a) On the List/Add DHCP Failover Pairs page, click the **Go Local** icon next to Boston-cluster.
- b) On the Manage DHCP Server page, click the **Reload** icon.
- c) Click the **Go Regional** icon at the top of the page to return to the regional cluster.

Step 6 Confirm the failover configuration and reload the server at the Chicago cluster in the same way.



CHAPTER 7

Managing Routers and Router Interfaces

This chapter explains how to add and edit routers and router interfaces in Cisco Prime Network Registrar.

- [Adding Routers, on page 125](#)
- [Editing Routers, on page 126](#)
- [Viewing and Editing the Router Interfaces, on page 126](#)
- [Pushing and Reclaiming Subnets for Routers, on page 127](#)

Adding Routers

Local Advanced and Regional Web UI

- Step 1** From the **Deploy** menu, choose **Router List** (in regional web UI) or **Routers** (in local web UI) under the **Router Configuration** submenu. This opens the List/Add Routers page.
- Step 2** Click the **Add Routers** icon. This opens the Add Router page.
- Step 3** On the Add Router dialog box, add the router based on data from its administrator:
- a) Give the router a distinguishing name in the name field.
 - b) Enter the router description in the description field.
 - c) Enter the router IP address in the address field.
 - d) Enter the management interface address for the router in the address field.
 - e) Enter the IPv6 management interface address for the router in the ip6address field.
 - f) Choose a owner and region.
- Step 4** Click **Add Router**.
-

CLI Commands

Add a router using `router name create address [attribute=value]`. The *address* can be either IPv4 or IPv6. For example:

```
nrcmd> router router-1 create 192.168.121.121
```

Editing Routers

Editing routers involves modifying some of the router attributes.

Local Advanced and Regional Web UI

Click the router name in the Router Tree pane or Router List pane on the left. In the Edit Router page, you can enter values for the different attributes. Additionally, you can use the **Unset** checkbox also to disable the attributes. Make your changes, then click **Save**.

CLI Commands

Edit a router attribute using **router name set attribute=value [attribute=value ...]**. For example:

```
nrcmd> router router-1 set owner=owner-1
```

Viewing and Editing the Router Interfaces

Editing a router interface involves modifying some of its attributes.

Local Advanced and Regional Web UI

If you click the **Interfaces** tab associated with the router on the List/Add Routers page, the list of related cable or Ethernet interfaces appears. Both from this page and the Router Tree pane on the left, you can click the interface name to edit it. The Interfaces tab also contains the option to delete the interface (click the **Delete** icon corresponding to the interface). Editing the interface also includes an additional attribute **Unset** function. You can add, edit, or delete interfaces for virtual routers without restrictions. A vpn-id that qualifies the addresses, subnets, and prefixes for the router interface can also be selected in the Edit Router Interface page.



Note Modifying a router interface is done as a delete and then an add of the router interface.

CLI Commands

Edit a router interface attribute using **router-interface name set attribute=value**. For example:

```
nrcmd> router-interface Ethernet1/0 set ip-helper=192.168.121.122
```

Related Topics

[Changeable Router Interface Attributes, on page 126](#)

[Bundling Interfaces, on page 127](#)

Changeable Router Interface Attributes

If you are editing the attributes of the router interface, you can change the following attributes:

- Name
- MAC address
- Description
- Address of the primary subnet address on the interface
- Addresses of the secondary subnets on the interface
- Address of any IP helper (DHCP relay agent) for the interface
- Address of any cable helper of the DHCP server to accept unicast packets for the interface
- Link associated with the router interface
- IPv6 address of the router interface
- IPv6 DHCP relay destination addresses configured for the interface

Bundling Interfaces

An interface bundle provides load balancing among the router interfaces. When you define a bundle, all the participating interfaces in the bundle must have the same bundle identifier (ID), which is the name of the interface specified as the master.

If you want to use bundling, the following attributes are in the Interface Bundling Settings section of the Edit Router Interface page, or set them using the **router-interface** command in the CLI:

- **bundle-id**—Interface bundle identifier, the name of the master interface. All participating interfaces in the bundle must have the same bundle ID.
- **is-master**—This interface is the master interface in the bundle.

Pushing and Reclaiming Subnets for Routers

You can push subnets to, and reclaim subnets from, a router interface (see the *"Reclaiming Subnets" section in Cisco Prime Network Registrar 9.0 DHCP User Guide*). When you push or reclaim a subnet with a virtual router, all primary and secondary relationships that are set for the router interface are also set for the related subnets and scopes.



CHAPTER 8

Maintaining Servers and Databases

This chapter explains how to administer and control your local and regional server operations.

- [Managing Servers, on page 129](#)
- [Scheduling Recurring Tasks, on page 131](#)
- [Logs, on page 132](#)
- [Running Data Consistency Rules, on page 137](#)
- [Monitoring and Reporting Server Status, on page 140](#)
- [Troubleshooting DHCP and DNS Servers, on page 156](#)
- [Using the TAC Tool, on page 160](#)
- [Troubleshooting and Optimizing the TFTP Server, on page 161](#)

Managing Servers

If you are assigned the server-management subrole of the ccm-admin role, you can manage the Cisco Prime Network Registrar servers as follows:

- **Start**—Load the database and start the server.
- **Stop**—Stop the server.
- **Reload**—Stop and restart the server. (Note that you do not need to reload the server for all RR updates, even protected RR updates. For details, see the *"Managing DNS Update"* chapter in *Cisco Prime Network Registrar 9.0 DHCP User Guide*.)
- **Check statistics**—See the [Displaying Statistics, on page 142](#).
- **View logs**—See the [Searching the Logs, on page 135](#).
- **Manage interfaces**—See the specific protocol pages for how to manage server interfaces.

Starting and stopping a server is self-explanatory. When you reload the server, Cisco Prime Network Registrar performs three steps—stops the server, loads configuration data, and restarts the server. Only after you reload the server does it use your changes to the configuration.



Note The CDNS, DNS, DHCP, and SNMP servers are enabled by default to start on reboot. The TFTP server is not enabled by default to start on reboot. You can change this using `[server] type enable` or `disable start-on-reboot` in the CLI.



Note If **exit-on-stop** attribute of DHCP, DNS, or TFTP server is enabled, then the statistics and scope utilization data only from the last start (reload) is reported while if the attribute is disabled, information across reloads is displayed.

Local Basic or Advanced and Regional Web UI

You can manage the protocol servers in the following ways depending on if you are a:

- **Local or regional cluster administrator**—Choose **Manage Servers** from the **Operate** menu to open the Manage Servers page.

The local and regional cluster web UI access to server administration is identical, even though the available functions are different. As a regional administrator, you can check the state and health of the regional CCM server and server agent. However, you cannot stop, start, reload, or view statistics, logs, or interfaces for them.

At the local cluster, to manage the DHCP, DNS, CDNS, TFTP, or SNMP servers, select the server in the Manage Servers pane and do any of the following:

- Click the **Statistics** tab to view statistics for the server. (See the [Displaying Statistics, on page 142.](#))
- Click the **Logs** tab in the View Log column to view the log messages for the server. (See the [Searching the Logs, on page 135.](#))
- Click the **Start Server** button to start the server.
- Click the **Stop Server** button stop the server.
- Click the **Restart Server** button to reload the server.

- **Local cluster DNS administrator**—Choose **DNS Server** from the **Deploy** menu to open the Manage DNS Authoritative Server page.

Along with the Statistics, Startup Logs, Logs, HA DNS Server Status, Start Server, Stop Server, and Restart Server functions, you can also perform other functions when you click the **Commands** button to open the DNS Commands dialog box.

The server command functions are:

- **Forcing all zone transfers** (see the *"Enabling Zone Transfers" section in Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*)—Click the **Run** icon. This is the equivalent of **dns forceXfer secondary** in the CLI.
- **Scavenging all zones** (see the *"Scavenging Dynamic Records" section in Cisco Prime Network Registrar 9.0 DHCP User Guide*)—Click the **Run** icon. This is the equivalent of **dns scavenge** in the CLI.

- **Local cluster Caching DNS server**—Choose **CDNS Server** from the **Deploy** menu to open the Manage DNS Caching Server page.

Along with the Statistics, Startup Logs, Logs, Start Server, Stop Server, and Restart Server functions, you can also perform other functions when you click the **Commands** button to open the CDNS Commands dialog box.

In Advanced and Expert modes, you can flush Caching CDNS cache and flush the resource records. Click the Commands button to execute the commands.

- **Local cluster DHCP administrator**—Click **DHCP Server** from the **Deploy** menu to open the Manage DHCP Server page.

Along with the Statistics, Startup Logs, Logs, Start Server, Stop Server, and Restart Server functions, you can also perform other functions when you click the **Commands** button to open the DHCP Server Commands dialog box.

This page provides the Get Leases with Limitation ID feature, to find clients that are associated through a common limitation identifier (see the *"Administering Option 82 Limitation" section in Cisco Prime Network Registrar 9.0 DHCP User Guide*). Enter at least the IP address of the currently active lease in the IP Address field, then click the **Run** icon. You can also enter the limitation ID itself in the form *nn:nn:nn* or as a string ("*nnnn*"), in which case the IP address becomes the network in which to search. This function is the equivalent of **dhcp limitationList ipaddress [limitation-id] show** in the CLI.

CLI Commands

In the CLI, the regional cluster allows CCM server management only:

- To start the server, use **server type start** (or simply **type start**; for example, **dhcp start**).
- To stop the server, use **server type stop** (or simply **type stop**; for example, **dhcp stop**). If stopping the server, it is advisable to save it first using the **save** command.
- To reload the server, use **server type reload** (or simply **type reload**; for example, **dhcp reload**). Cisco Prime Network Registrar stops the server you chose, loads the configuration data, and then restarts the server.
- To set or show attributes for the server, use **[server] type set attribute=value** or **[server] type show**. For example:

```
nrcmd> ccm set ipaddr=192.168.50.10
```

Scheduling Recurring Tasks

In Basic and Advanced user mode in the local cluster web UI, you can schedule a number of recurring tasks. These tasks are:

- Reloading the DHCP server.
- Reloading the DNS server.
- Synchronizing DHCP failover server pairs:
 - If in staged dhcp edit mode, reload the main DHCP server.
 - Synchronize the failover configuration to the backup DHCP server.
 - If in staged dhcp edit mode, reload the backup DHCP server.
- Synchronizing High-Availability (HA) DNS server pairs:
 - If in staged dhcp edit mode, reload the main DNS server.
 - Synchronize the HA DNS configuration to the backup DNS server.
 - If in staged dhcp edit mode, reload the backup DNS server.
- Synchronizing zone distribution maps:
 - If in staged dhcp edit mode, reload the main DNS server.
 - If in staged dhcp edit mode, reload the backup HA DNS server.

- Synchronize the zone distribution maps.
- If in staged dhcp edit mode, reload the secondary DNS server or servers.

Local Basic or Advanced Web UI

To set up one or more of these recurring server tasks:

-
- Step 1** From the **Operate** menu, choose **Schedule Tasks** to open the List/Add Scheduled Tasks page.
- Step 2** Click the **Add Scheduled Task** icon in the Scheduled Tasks pane on the left to open the Add Scheduled Task page.
- Step 3** Enter values in the appropriate fields:
- Name of the scheduled task. This can be any identifying text string.
 - Pull down from the available list of task types, which are:
 - **dhcp-reload**—Reloads the DHCP server
 - **dns-reload**—Reloads the DNS server
 - **cdns-reload**—Reloads the Caching DNS server
 - **sync-dhcp-pair**—Synchronizes the DHCP failover server pair
 - **sync-dns-pair**—Synchronizes the HA DNS failover server pair
 - **sync-zd-map**—Synchronizes zone distribution maps
 - **sync-dns-update-map**—Synchronizes DNS update maps
 - Indicate the time interval for the scheduled task, such as 60m or 4w2d.
- Step 4** Click **Add Scheduled Task**.
- Step 5** If you click the name of the task on the List/Add Scheduled Tasks page, on the Edit Scheduled Task page you can view (in the Task Status section) the last status or the list of last errors (if any) that occurred during the task execution. Click **Run Now** to run the task immediately.
- Note** The DNS server startup and background loading slows down when HA is enabled before the HA DNS server communicates to its partner. You need to allow the HA DNS server to communicate with its partner before reloading or restarting the DNS server.
-

Logs

Log Files

The following table describes the Cisco Prime Network Registrar log files in the *install-path/logs* directory.

Table 14: Log Files in .../logs Directory

Component	File in /logs Directory	Local/Regional	Logs
Installation	install_cnr_log	Both	Installation process
Upgrade	ccm_upgrade_status_log	Both	Upgrade process
	dns_upgrade_status_log	Local	Upgrade process

Component	File in /logs Directory	Local/Regional	Logs
Server agent	agent_server_1_log	Both	Server agent starts and stops
Port check	checkports_log	Both	Network ports
DNS server	name_dns_1_log	Local	DNS activity
	dns_startup_log	Local	DNS startup activity
CDNS server	cdns_log	Local	CDNS activity
	cdns_startup_log	Local	CDNS startup activity
DHCP server	name_dhcp_1_log	Local	DHCP activity
	dhcp_startup_log	Local	DHCP startup activity
TFTP server	file_tftp_1_log file_tftp_1_trace	Local	TFTP activity
SNMP server	cnrsnmp_log	Both	SNMP activity
CCM database	config_ccm_1_log	Both	CCM configuration, starts, stops
	ccm_startup_log		CCM startup activity
Web UI	cnrwebui_log	Both	Web UI state
Tomcat/web UI (in cnrwebui subdirectory)	catalina.date .log.txt jsui_log.date .txt cnrwebui_access_log.date .txt	Both	CCM database for Tomcat server and web UI (Because new files are created daily, periodically archive old log files.)
Resource Limits	ccm_monitor_log	Both	Resource limit activity.

DNS, DHCP, CDNS, CCM, and TFTP servers can generate a number of log files, each with a preconfigured maximum size of 10 MB. This preconfigured value applies to new installs only.



Note Upgrades from pre-9.0 versions will use the old preconfigured (or explicitly configured) value of 1,000,000 bytes for log files.

The first log file name has the `_log` suffix. When this file reaches its maximum size, it gets the `.01` version extension appended to its name and a new log file is created without the version extension. Each version extension is incremented by one for each new file created. When the files reach their configured maximum number, the oldest file is deleted and the next oldest assumes its name. The usual maximum number is 10 for the DNS, DHCP, CDNS, CCM, and TFTP servers.

Cisco Prime Network Registrar also has `server_startup_log` files. This applies to the CCM, DHCP, DNS, and TFTP servers. These files log the start up and shut down phases of the server (the information is similar to the normal log file information). Server startup log files are useful in diagnosing problems that have been reported when the server was last started.

The number of these start-up logs is fixed at four for a server, and the size is fixed at 10 MB per server.



Note Some user commands can create *User authentication* entries in the Server Agent log because of separate connections to the cluster. Do not interpret these as a system security violation by another user.

Logging can also be directed to syslog. See [Modifying the cnr.conf File, on page 157](#).

CLI Commands

You can check the configured maximums for the DNS, DHCP, and TFTP servers using `[server] type serverLogs show` in the CLI, which shows the maximum number (*nlogs*) and size (*logsize*) of these protocol server log files. You can adjust these parameters using `[server] type serverLogs set nlogs=value` and `[server] type serverLogs set logsize=value`. You cannot adjust these maximums for any of the other log files.



Note A change to the server logs will not take effect until you restart Cisco Prime Network Registrar.

Logging Server Events

When you start Cisco Prime Network Registrar, it automatically starts logging Cisco Prime Network Registrar system activity. Cisco Prime Network Registrar maintains all the logs by default on:

- **Windows**—*install-path*\logs
- **Linux**—*install-path*/logs (to view these logs, use the **tail -f** command)



Tip To avoid filling up the Windows Event Viewer and preventing Cisco Prime Network Registrar from running, in the Event Log Settings, check the **Overwrite Events as Needed** box. If the events do fill up, save them to a file, then clear them from the Event Log.

Local Basic or Advanced and Regional Web UI

Server logging is available in the web UI when you open the Manage Servers page for a server (see the [Managing Servers, on page 129](#)), then click the **Log** icon in the View Log column for the server. This opens the Log for Server page. The log is in chronological order with the page with the latest entries shown first. If you need to see earlier entries, click the left arrow at the top or bottom of the page.

Related Topics

[Searching the Logs, on page 135](#)

[Logging Format and Settings, on page 134](#)

Logging Format and Settings

The server log entries include the following categories:

- **Activity**—Logs the activity of your servers.
- **Info**—Logs standard operations of the servers, such as starting up and shutting down.

- **Warning**—Logs warnings, such as invalid packets, user miscommunication, or an error in a script while processing a request.
- **Error**—Logs events that prevent the server from operating properly, such as out of memory, unable to acquire resources, or errors in configuration.



Note Warnings and errors go to the Event Viewer on Windows (see the Tip in [Logging Server Events, on page 134](#)). For a description of the log messages for each server module, see the *install-path/docs/msgid/MessageIdIndex.html* file.

Local Basic or Advanced and Regional Web UI

You can affect which events to log. For example, to set the logging for the local cluster DNS and DHCP server:

- **DNS**—From the **Deploy** menu, choose **DNS Server** under the **DNS** submenu to open the Manage DNS Server page. Click the name of the server to open the Edit DNS Server page. Expand the Log Settings section to view the log settings. Make changes to the attributes as desired, click **Save**, and then reload the server. (See *Table 4 in the "Troubleshooting DNS Servers" section in Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide* for the log settings to maximize DNS server performance.)
- **DHCP**—From the **Deploy** menu, choose **DHCP Server** under the **DHCP** submenu to open the Manage DHCP Server page. Click the name of the server to open the Edit DHCP Server page. Expand the Log Settings section to view the log settings. Make changes to the attributes as desired, click **Save**, and then reload the server. (See *Table 6 in the "Tuning the DHCP Server" section in Cisco Prime Network Registrar 9.0 DHCP User Guide* for the log settings to maximize DHCP server performance.)

CLI Commands

Use `dns set log-settings=value`, `dhcp set log-settings=value`, and `tftp set log-settings=value` for the respective servers.

Searching the Logs

The web UI provides a convenient way to search for entries in the activity and startup log files. You can locate specific message text, log message IDs, and message timestamps using a regular expression string entry. When you click the **Log** icon in the View Log or View Startup Log column on the Manage Servers page (or one of the specific server pages), this opens a Log for Server page. In the text field next to the Search icon at the top or bottom of the page, enter the search string in the regular expression syntax. (For example, entering **name?** searches for occurrences of the string *name* in the log file.) Click the **Search** icon to view the results of log search.

Click the name of the log message, which opens the Log for Server page with the full message text. To view the full message text, click the name of the log message. Change between Table and Text view by clicking the **Log** icon. Click **Close** on the Log Search Result page to close the browser window.

View Change Log

In the web UI, you can view the change logs and tasks associated with configurations you make.

Local Basic and Advanced Web UI

From the **Operate** menu, choose **Change Log**. To view the change log, you must be assigned the database subrole of the ccm-admin or regional-admin role:

- The View Change Log page shows all the change logs, sorted by DBSN name. To get to the bottom of the list, click the right arrow at the bottom left of the page. Click the DBSN number of the change log entry to open a View Change Set page for it.

On the View Change Log page, you can filter the list, manually trim it, and save it to a file. You can filter the list by:

- Start and end dates
- Administrator who initiated the changes
- Configuration object class
- Specific object
- Object identifier (ID), in the format OID-00:00:00:00:00:00:00:00
- Server
- Database

Click **Filter List** or **Clear Filter** (to clear the filter that persists through the session). You can initiate a trim of the change log by setting how many days old you want the record to get before trimming it, by setting a number of days value in the “older than” field and clicking the **Delete** icon.

To save the change log entries to a comma-separated values (CSV) file, click the **Save** icon.

If a task is associated with a change log, it appears on the View Change Set page. You can click the task name to open the View CCM Task page for it.

Dynamic Update on Server Log Settings

The DHCP and the DNS servers register the changes on the server logs only during the server configuration, which happens during a reload. Reloading the servers is time consuming. Cisco Prime Network Registrar allows the DHCP and DNS servers to register the changes to log settings, without a reload.

Local Basic or Advanced Web UI

To dynamically update DHCP server log settings, do the following:

-
- Step 1** From the **Deploy** menu, choose **DHCP Server** under the **DHCP** submenu. The Manage DHCP Server page appears.
 - Step 2** Click the name of the DHCP server in the left pane to open the Edit DHCP Server page.
 - Step 3** Modify the log settings as desired.
 - Step 4** Click **Save** at the bottom of the page. The new log settings are applied to the DHCP server. The Manage DHCP Server page is displayed with an updated page refresh time.
-

Local Basic or Advanced Web UI

To dynamically update DNS server log settings, do the following:

-
- Step 1** From the **Deploy** menu, choose **DNS Server** under the **DNS** submenu. This opens the Manage DNS Server page.

- Step 2** Click the name of the DNS server in the left pane to open the Edit DNS Server page.
- Step 3** Modify the log settings as desired.
- Step 4** Click **Save** at the bottom of the page. The new log settings are applied to the DNS server. The Manage DNS Server page is displayed with an updated page refresh time.
- Note** If the `dhcp-edit-mode` or `dns-edit-mode` is set to synchronous, and if the server running, the change in server log settings is communicated to the server.

CLI Commands

To dynamically update the DHCP or DNS server log settings using the CLI, you must have the appropriate `edit-mode` set to synchronous. After changing the server log settings, use the `save` command to save the settings.

For example:

```
nrcmd> session set dhcp-edit-mode=synchronous
nrcmd> dhcp set log-settings=new-settings
nrcmd> save
```

Running Data Consistency Rules

Using consistency rules, you can check data inconsistencies such as overlapping address ranges and subnets. You can set data consistency rules at the regional and local clusters.

The table on the List Consistency Rules page explains these rules. Check the check box next to the rule that you want to run.



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`.

The List Consistency Rules page includes functions to select all rules and clear selections. You can show the details for each of the rule violations as well as view the output. The rule selections you make are persistent during your user session.

Local Basic or Advanced and Regional Web UI

To run consistency rules, do the following:

-
- Step 1** From the **Operate** menu, choose **Consistency Reports** under the **Reports** submenu. The List Consistency Rules page appears.
- Step 2** Check the check boxes for each of the listed consistency rules that you want to apply.
- To select all the rules, click the **Select All Rules** link.
 - To clear all selections, click the **Clear Selection** link.
- Step 3** Click **Run Rules**.

The Consistency Rules Violations page appears. The rules are categorized by violation type.

- To show details for the violations, click the **Show Details** link.
- To show the output, click the page icon.

Step 4 Click **Return to Consistency Rules** to return to the List Consistency Rules page.

CLI Tool

Use the **cnr_rules** consistency rules tool from the command line to check for database inconsistencies. You can also use this tool to capture the results of the rule in a text or XML file.

The **cnr_rules** tool is located at:

- **Windows**—...\bin\cnr_rules.bat
- **Linux**—.../usrbin/cnr_rules

To run the **cnr_rules** tool, enter:

```
> cnr_rules -N username -P password [options]
```

- **-N username** —Authenticates using the specified username.
- **-P password** —Authenticates using the specified password.
- **options** —Describes the qualifying options for the tool, as described in the following table. If you do not enter any options, the command usage appears.

Table 15: cnr_rules Options

Option	Description
Example	
-list	Lists the available consistency rules. Note The list of available commands is tailored to the permissions of the administrator specified in the value of the -N option.
	<pre>> cnr_rules -N admin -P changeme -list</pre>

Option	Description
<p><code>-run [rule-match]</code></p>	<p>Run the available rules. Optionally, you can run a subset of the available rules by applying a case-insensitive rule-match string.</p> <ul style="list-style-type: none"> • Runs all rules: <pre data-bbox="963 464 1243 510">> cnr_rules -N admin -P changeme -run</pre> <ul style="list-style-type: none"> • Runs only the rules whose names contain the string “dhcp”: <pre data-bbox="963 611 1500 636">> cnr_rules -N admin -P changeme -run dhcp</pre> <p>Tip To match a string containing spaces, enclose the string using double-quotation marks ("). For example: <code>> cnr_rules -N admin -P changeme -run "router interface"</code></p>
<p><code>-details</code></p>	<p>Includes details of the database objects that violate consistency rules in the results.</p> <p>Runs the DNS rules, and includes details of the database object in the results:</p> <pre data-bbox="963 1020 1487 1066">> cnr_rules -N admin -P changeme -run DNS -details</pre>
<p><code>-xml</code></p>	<p>Generates rule results in an XML file.</p> <p>Note When using the <code>-xml</code> option, the <code>-details</code> option is ignored because the XML file includes all the detailed information.</p> <pre data-bbox="963 1266 1500 1291">> cnr_rules -N admin -P changeme -run -xml</pre>
<p><code>-path classpath</code></p>	<p>Changes the Java classpath that is searched to locate the available consistency rules (optional).</p> <p>In order to run a new, custom consistency rule, you can use this option. You must get the support of a support engineer to do this.</p>

Option	Description
-interactive	<p>Runs the tool in an interactive session.</p> <pre>> cnr_rules -N admin -P changeme -run -interactive RuleEngine [type ? for help] > ? Commands: load <class> // load the specified rule class run <rule-match> // run rules matching a string, or '*' for all list // list rules by name xml // toggle xml mode detail // toggle detail mode (non-xml only) quit // quit RuleEngine</pre>
-both	Displays domain names in both Unicode and ASCII.

You can redirect the output of any of these preceding commands to another file. Use the following syntax to capture the rule results in a:

- Text file:

```
> cnr_rules -N username -P password -run -details > filename.txt
```

- XML file:

```
> cnr_rules -N username -P password -run -xml > filename.xml
```

Monitoring and Reporting Server Status

Monitoring the status of a server involves checking its:

- State
- Health
- Statistics
- Log messages
- Address usage
- Related servers (DNS and DHCP)
- Leases (DHCP)

Related Topics

[Server States, on page 141](#)

[Displaying Health, on page 141](#)

[Displaying Statistics, on page 142](#)

[Displaying IP Address Usage, on page 153](#)

[Displaying Related Servers, on page 153](#)

[Displaying Leases, on page 156](#)

Server States

All Cisco Prime Network Registrar protocol servers (DNS, DHCP, SNMP, and TFTP) pass through a state machine consisting of the following states:

- **Loaded**—First step after the server agent starts the server (transitional).
- **Initialized**—Server was stopped or fails to configure.
- **Unconfigured**—Server is not operational because of a configuration failure (transitional).
- **Stopped**—Server was administratively stopped and is not running (transitional).
- **Running**—Server is running successfully.

The two essential states are initialized and running, because the server transitions through the states so quickly that the other states are essentially invisible. Normally, when the server agent starts the server, it tells the server to be up. The server process starts, sets its state to loaded, then moves up to running. If you stop the server, it walks down the states to initialized, and if you restart, it moves up to running again. If it fails to configure for some reason, it drops back to initialized, as if you had stopped it.

There is also an exiting state that the server is in very briefly when the process is exiting. The user interface can also consider the server to be disabled, but this rarely occurs and only when there is no server process at all (the server agent was told not to start one).

Displaying Health

You can display aspects of the health of a server, or how well it is running. The following items can decrement the server health, so you should monitor their status periodically. For the:

- Server agent (local and regional clusters)
- CCM server (local and regional clusters)
- DNS server (local cluster):
 - Configuration errors
 - Memory
 - Disk space usage
 - Inability to contact its root servers
- Caching DNS server (local cluster)
- DHCP server (local cluster):
 - Configuration errors
 - Memory
 - Disk space usage
 - Packet caching low
 - Options not fitting in the stated packet limit
 - No more leases available
- TFTP server (local cluster):
 - Memory
 - Socket read or write error
 - Exceeding the overload threshold and dropping request packets

Server Health Status

The server health status varies from the value 0 to 10. The value 0 means the server is not running and 10 means the server is running. Some of the servers report only 0 or 10, and not anything in between. When a server reports a value from 1 to 9, it means that it detected conditions that indicate possible problems. It has nothing to do with the actual performance of the server. So, if the health of the server is a value from 1 to 9, the server log files need to be reviewed to see what errors were logged.



Note Depending on the level of activity and the size and number of log files, the condition that reduced the server health might not be visible in the log files. It is important to review the log files, but the servers do not log all the conditions that reduce the server health.

The following conditions can reduce the DHCP server health:

- Configuration errors (occurs when the server is getting started or restarting)
- When the server detects out-of-memory conditions
- When packet receive failures occur
- When packets are dropped because the server is out of request or response buffers
- When the server is unable to construct a response packet

Similar conditions exist for the TFTP server.



Tip Health values range from 0 (the server is not running) to 10 (the highest level of health). It is recommended that the health status can be ignored, with the understanding that zero means server is not running and greater than zero means server is running. On Linux, you can run the `cnr_status` command, in the `install-path /usrbin/` directory, to see if your local cluster server is running. For more information on how to check whether the local cluster server is running, see the *Cisco Prime Network Registrar Installation Guide*.

Local Basic or Advanced and Regional Web UI

From the **Operate** menu, select **Manage Servers**. Check the Manage Servers page for the state and health of each server.

CLI Commands

Use `[server] type getHealth`. The number 10 indicates the highest level of health, 0 that the server is not running.

Displaying Statistics

To display server statistics, the server must be running.

Local Basic or Advanced and Regional Web UI

Go to the Manage Servers page, click the name of the server in the left pane, then click the **Statistics** tab, if available. On the Server Statistics page, click the name of the attribute to get popup help.

The DHCP, DNS, and CDNS statistics are each divided into two groups of statistics. The first group is for total statistics and the second group is for sample statistics. The total statistics are accumulated over time. The

sample statistics occur during a configurable sample interval. The names of the two categories vary per server and per user interface, and are identified in the following table.

Table 16: Server Statistics Categories

Server	User Interface	Total Statistics (Command)	Sample Statistics (Command)
DHCP	Web UI	Total Statistics	Activity Summary
	CLI	Total Counters since the start of the last DHCP server process (dhcp getStats)	Sampled counters since the last sample interval (dhcp getStats sample)
DNS	Web UI	Total Statistics	Sample Statistics
	CLI	Total Counters since the start of the last server process (dns getStats)	Sampled counters since the last sample interval (dns getStats sample)
CDNS	Web UI	Total Statistics	Sample Statistics
	CLI	Total Counters since the start of the last server process (cdns getStats total)	Sampled counters since the last sample interval (cdns getStats sample)

To set up the sample counters, you must activate either the *collect-sample-counters* attribute for the server or a *log-settings* attribute value called *activity-summary*. You can also set a *log-settings* value for the sample interval for each server, which is preset to 5 minutes. The *collect-sample-counters* attribute is preset to true for the DNS server, but is preset to false for the DHCP server. For example, to enable the sample counters and set the interval for DHCP, set the following attributes for the DHCP server:

- Enable *collect-sample-counters* (**dhcp enable collect-sample-counters**)
- Set *log-settings* for *activity-summary* (**dhcp set log-settings=activity-summary**)
- Set *activity-summary-interval* to 5m (**dhcp set activity-summary-interval=5m**)

CLI Commands

In the CLI, if you use **[server] type getStats**, the statistics are encoded in curly braces followed by sets of digits, as described in [Table 17: DNS Statistics](#) for DNS, [Table 19: DHCP Statistics](#) for DHCP, and [Table 20: TFTP Statistics](#) for TFTP. The **server type getStats all** command is more verbose and identifies each statistic on a line by itself. Using the additional **sample** keyword shows the sample statistics only.

Reset the counters and total statistic by using **dhcp resetStats**, **dns resetStats**, or **cdns resetStats**.

DNS Statistics

The DNS server statistics in the web UI appear on the DNS Server Statistics page, click on the statistic's name to read its description. You can refresh the DNS Server Statistics.

The DNS server statistics that you can view are:

- **Attribute**—Displays server statistics such as server identifier, recursive service, process uptime, time since reset, and so on.

Total Statistics

- **Performance Statistics**—Displays the total statistics of the DNS Server performance.
- **Query Statistics**—Displays the total statistics of the queries.
- **HA Statistics**—Displays the total statistics of the HA DNS Server.
- **Push Notification Statistics**—Displays the total statistics of DNS Push Notifications.
- **Security Statistics**—Displays the total statistics of the security.
- **IPv6 Statistics**—Displays the total statistics of the IPv6 packets received and sent.
- **Error Statistics**—Displays the total statistics of the errors.
- **Max Counter Statistics**—Displays the total statistics of the maximum number of concurrent threads, RRs, DNS update latency, concurrent packets, and so on.

Sample Statistics

- **Performance Statistics**—Displays the sample statistics about the DNS Server performance.
- **Query Statistics**—Displays the sample statistics about the queries.
- **HA Statistics**—Displays the sample statistics about the HA DNS Server.
- **Push Notification Statistics**—Displays the sample statistics of DNS Push Notifications.
- **Security Statistics**—Displays the sample statistics about the security.
- **IPv6 Statistics**—Displays the sample statistics about the IPv6 packets received and sent.
- **Error Statistics**—Displays the sample statistics about the errors.



Note To get the most recent data, click the **Refresh Server Statistics** icon at the top left of the Statistics page.

The **dns getStats** command has the following options:

```
dns getStats [performance | query | errors | security | maxcounters | ha | ipv6 | dns-pn |
all] [total | sample]
```

The **dns getStats all** command is the most commonly used. The **dns getStats** command without **all** option returns the statistics in a single line of positional values in the following format (the table below shows how to read these values):

```
nrcmd> dns getStats

100 Ok
{1} 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17
```

Table 17: DNS Statistics

Digit	Statistic	Description
{1}	id	Implementation ID (release and build information).
2	config-recurs	Recursion services—(1) available, (2) restricted, (3) unavailable.

Digit	Statistic	Description
3	config-up-time	Time (in seconds) elapsed since the last server startup.
4	config-reset-time	Time (in seconds) elapsed since the last server reset (restart).
5	config-reset	Status or action to reinitializes any name server state—If using the (2) reset action, reinitializes any persistent name server state; the following are read-only statuses: (1) other—server in some unknown state, (3) initializing, or (4) running.
6	counter-auth-ans	Number of queries answered authoritatively.
7	counter-auth-no-names	Number of queries returning authoritative no such name responses.
8	counter-auth-no-data-resps	Number of queries returning authoritative no such data (empty answer) responses. (Deprecated statistics)
9	counter-non-auth-datas	Number of queries answered nonauthoritatively (cached). (Deprecated statistics)
10	counter-non-auth-no-datas	Number of queries answered nonauthoritatively with no data.
11	counter-referrals	Number of queries forwarded to other servers.
12	counter-errors	Number of responses answered with errors (RCODE values other than 0 or 3).
13	counter-rel-names	Number of requests received for names of only one label (relative names).
14	counter-req-refusals	Number of refused queries.
15	counter-req-unparses	Number of unparseable requests.
16	counter-other-errors	Number of aborted requests due to other errors.
17	total-zones	Total number of configured zones.

CDNS Statistics

The CDNS server statistics in the web UI appear on the DNS Caching Server Statistics page, click on the name of the statistics to read its description. You can refresh the CDNS Server Statistics.

Table 18: CDNS Statistics

Digit	Statistic	Description
{1}	name	Name identifying the DNS Caching Server.
2	time-current	The current time given by the CDNS Server.
3	time-up	The amount of time the server has been up and running.
4	time-elapsed	The elapsed since last statistics poll.
5	queries-total	Total number of queries received by the CDNS Server.
6	queries-over-tcp	Total number of queries received over TCP by the CDNS Server.
7	queries-over-ipv6	Total number of queries received over TCP by the CDNS Server.
8	queries-with-edns	Number of queries with EDNS OPT RR present.
9	queries-with-edns-do	Number of queries with EDNS OPT RR with DO (DNSSEC OK) bit set.
10	queries-type-A	Number of A queries received.
11	queries-type-AAAA	Number of AAAA queries received.
12	queries-type-CNAME	Number of CNAME queries received.
13	queries-type-PTR	Number of PTR queries received.
14	queries-type-NS	Number of NS queries received.
15	queries-type-SOA	Number of SOA queries received.
16	queries-type-MX	Number of MX queries received.
17	queries-type-DS	Number of DS queries received.

Digit	Statistic	Description
18	queries-type-DNSKEY	Number of DNSKEY queries received.
19	queries-type-RRSIG	Number of RRSIG queries received.
21	queries-type-NSEC	Number of NSEC queries received.
22	queries-type-NSEC3	Number of NSEC3 queries received.
23	queries-type-other	Number of queries received of type 256+.
24	queries-with-flag-QR	Number of incoming queries with QR (query response) flag set. These queries are dropped.
25	queries-with-flag-AA	Number of incoming queries with AA (auth answer) flag set. These queries are dropped.
26	queries-with-flag-TC	Number of incoming queries with TC (truncation) flag set. These queries are dropped.
27	queries-with-flag-RD	Number of incoming queries with RD (recursion desired) flag set.
28	queries-with-flag-RA	Number of incoming queries with RA (recursion available) flag set.
29	queries-with-flag-Z	Number of incoming queries with Z flag set.
30	queries-with-flag-AD	Number of incoming queries with AD flag set.
31	queries-with-flag-CD	Number of incoming queries with CD flag set.
32	queries-failing-acl	Number of queries being dropped or refused due to ACL failures.
33	cache-hits	The total number of queries that were answered from cache.
34	cache-misses	The total number of queries that were not found in the cache.
35	cache-prefetches	Number of prefetches performed.

Digit	Statistic	Description
36	requestlist-total	The total number of queued requests waiting for recursive replies.
37	requestlist-total-user	The total number of queued user requests waiting for recursive replies.
38	requestlist-total-system	The total number of queued system requests waiting for recursive replies.
39	requestlist-total-average	The average number of requests on the request list.
40	requestlist-total-max	The maximum number of requests on the request list.
41	requestlist-total-overwritten	The number of requests on the request list that were overwritten by newer entries.
42	requestlist-total-exceeded	The number of requests dropped because the request list was full.
43	recursive-replies-total	The total number of recursive queries replies.
44	recursive-time-average	The average time to complete a recursive query.
45	recursive-time-median	The median time to complete a recursive query.
46	mem-process	An estimate of the memory in bytes of the CDNS process.
47	mem-cache	Memory in bytes of RRSet cache.
48	mem-query-cache	Memory in bytes of incoming query message cache.
49	mem-iterator	Memory in bytes used by the CDNS iterator module.
50	mem-validator	Memory in bytes used by the CDNS validator module.
51	answers-with-NOERROR	Number of answers from cache or recursion that result in rcode of NOERROR being returned to client.

Digit	Statistic	Description
52	answers-with- NXDOMAIN	Number of answers from cache or recursion that result in rcode of NXDOMAIN being returned to client.
53	answers-with-NODATA	Number of answers that result in pseudo rcode of NODATA being returned to client.
54	answers-with-other-errors	Number of answers that result in pseudo rcode of NODATA being returned to client.
55	answers-secure	Number of answers that correctly validated.
56	answers-unsecure	Number of answers that did not correctly validate.
57	answers-rrset-unsecure	Number of RRsets marked as bogus by the validator.
58	answers-unwanted	Number of replies that were unwanted or unsolicited. High values could indicate spoofing threat.
59	reset-time	Reports the most recent time the stats were reset (i.e. cdns resetStats in nrcmd).
60	sample-time	Reports the time the server collected the last set of sample statistics.
61	sample-interval	Reports the sample interval used by the server when collecting sample statistics.

DHCP Statistics

The DHCP server statistics in the web UI appear on the DHCP Server Statistics page, click on the statistic's name to read its description.

The DHCP server statistics details are available for:

- **Attribute**—Displays the server statistics such as server start time, server reload time, server up time, and statistics reset time.
- **Total Statistics**—Displays the total statistics of the scopes, request buffers, response buffers, packets and so on.

- Lease Counts (IPv4)—Displays the sample statistics of the IPv4 lease counts such as active leases, configured leases, reserved leases, and reserved active leases.
- Packets Received (IPv4)—Displays the sample statistics of the IPv4 packets received.
- Packets Sent (IPv4)—Displays the sample statistics of the IPv4 packets sent.
- Packets Failed (IPv4)—Displays the statistics of the failed IPv4 packets.

The Additional Attributes are:

- Failover Statistics—Displays the statistics of the DHCP failover server.
- IPv6 Statistics—Displays the statistics of the IPv6 prefixes configured, timed-out IPv6 offer packets and so on.
- Lease Counts (IPv6)—Displays the statistics of the IPv6 lease counts of active leases, configured leases, reserved leases, and reserved active leases.
- Packets Received (IPv6)—Displays the statistics of the IPv6 packets received.
- Packets Sent (IPv6)—Displays the statistics of the IPv6 packets sent.
- Packets Failed (IPv6)—Displays the statistics of the failed IPv6 packets.

Additional Attributes also includes Top Utilized Aggregations.



Note To get the most recent data, click the **Refresh Server Statistics** icon at the top left of the Statistics page.

The **dhcp getStats** command has the following options:

```
dhcp getStats [[all | server [,] failover [,] dhcpv6] [total | sample]
```

The **dhcp getStats all** command is the most commonly used. The **dhcp getStats** command without **all** option returns the statistics in a single line of positional values in the following format (the table below shows how to read these values):

```
nrcmd> dhcp getStats
```

```
100 Ok
{1} 2 3 4 5 6 7 8
```

Table 19: DHCP Statistics

Digit	Statistic	Description
{1}	start-time-str	Date and time of last server reload, as a text string.
2	total-discovers	Number of DISCOVER packets received.
3	total-requests	Number of REQUEST packets received.
4	total-releases	Number of RELEASED packets received.
5	total-offers	Number of OFFER packets sent.
6	total-acks	Number of acknowledgement (ACK) packets sent.

Digit	Statistic	Description
7	total-naks	Number of negative acknowledgement (NAK) packets sent.
8	total-declines	Number of DECLINE packets received.

TFTP Statistics

The TFTP server statistics in the web UI appear on the TFTP Server Statistics page, click on the statistic's name to read its description. The following table shows the TFTP statistics encoded as output to the generic **tftp getStats** command.

When the TFTP server starts up, it allocates sessions (tftp-max-sessions) and packets (tftp-max-packets) for its use. The TFTP session represents the communication between the TFTP client and TFTP server.

When a read request reaches the TFTP server, the server assigns a packet for the request, increments the total-packets-in-use and total-read-requests values by one, and responds to the user with a data packet. The TFTP server backs up the latest communication packet to resend, if needed. The TFTP server picks another packet from the pool to use it as data packet. When the TFTP server receives an acknowledgment for the block of data sent to the client, it sends the next data block. The TFTP server queues up packets associated with a session, if the session is not able to work on the packets immediately.

The TFTP server statistics details are available for:

- Attribute—Displays the server statistics such as port number, default device, home directory, use home directory as root, and so on.
- Log Settings—Displays the statistics of the log level, log settings, and packet trace level.



Note To get the most recent data, click the **Refresh Server Statistics** icon at the top left of the page.

TFTP statistics is encoded as an output to the generic **tftp getStats** command in the following format:

```
nrcmd> tftp getStats
100 Ok
{1} 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
```

Table 20: TFTP Statistics

Digit	Attribute	Description
{1}	id	Implementation ID (release and build information).
2	server-state	State of the server (up or down).
3	server-time-since-start	Running time since last start.
4	server-time-since-reset	Running time since last reset.
5	total-packets-in-pool	Number of packets in the pool.

Digit	Attribute	Description
6	total-packets-in-use	Number of packets the server is using.
7	total-packets-received	Number of packets received since the last start or reload.
8	total-packets-sent	Number of packets sent since the last start or reload.
9	total-packets-drained	Number of packets read and discarded since the last start or reload.
10	total-packets-dropped	Number of packets dropped since the last start or reload.
11	total-packets-malformed	Number of packets received that were malformed since the last start or reload.
12	total-read-requests	Number of packets read since the last start or reload.
13	total-read-requests-completed	Number of read packets completed since the last start or reload.
14	total-read-requests-refused	Number of read packets refused since the last start or reload.
15	total-read-requests-ignored	Number of read packets ignored since the last start or reload.
16	total-read-requests-timed-out	Number of read packets that timed out since the last start or reload.
17	total-write-requests	Number of read packets that were write requests since the last start or reload.
18	total-write-requests-completed	Number of write requests completed since the last start or reload.
19	total-write-requests-refused	Number of write requests refused since the last start or reload.
20	total-write-requests-ignored	Number of write requests ignored since the last start or reload.
21	total-write-requests-timed-out	Number of write requests that timed out since the last start or reload.
22	total-docsis-requests	Number of DOCSIS requests received since the last start or reload.
23	total-docsis-requests-completed	Number of DOCSIS requests completed since the last start or reload.
24	total-docsis-requests-refused	Number of DOCSIS requests refused since the last start or reload.
25	total-docsis-requests-ignored	Number of DOCSIS requests ignored since the last start or reload.
26	total-docsis-requests-timed-out	Number of DOCSIS requests that timed out since the last start or reload.
27	read-requests-per-second	Number of read requests per second.
28	write-requests-per-second	Number of write requests per second.
29	docsis-requests-per-second	Number of DOCSIS requests per second.

Displaying IP Address Usage

Displaying IP address usage gives an overview of how clients are currently assigned addresses.

Local Advanced and Regional Web UI

You can look at the local or regional cluster address space, or generate a subnet utilization or lease history report at the regional cluster, to determine IP address usage. These functions are available in both web UIs in the **Design > DHCPv4** menu, if you have address space privileges at the local or regional cluster.

You can determine the current address space utilization by clicking the **View** icon in the Current Usage column for the unified address space, address block, and subnet (see the *"Viewing Address Utilization for Address Blocks, Subnets, and Scopes"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*). You can also get the most current IP address utilization by querying the lease history (see the *"Querying Leases"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*). In the latter case, the regional CCM server references the appropriate DHCP server directly. To ensure this subnet-to-server mapping, you must update the regional address space view so that it is consistent with the relevant local cluster. Do this by pulling the replica address space, or reclaiming the subnet to push to the DHCP server (see the *"Reclaiming Subnets"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*). Also ensure that the particular DHCP server is running.

CLI Commands

You can generate an IP address usage report using the **report** command. The command has the following syntax:

```
report [column-separator=string]
       [dhcp-only]
       [dhcpv4]
       [dhcpv6]
       [file=outputfile]
       [vpn=name]
```

The column-separator specifies the character string that separates the report columns (the preset value is the space character). If you want to include more than one space, precede them with the backslash (\) escape character (enclosed in quotation marks). You can specify DHCPv4 or DHCPv6 addresses (**dhcp-only** is the same as **dhcpv4**). Not specifying the VPN returns the addresses in the current VPN only.

Displaying Related Servers

Cisco Prime Network Registrar displays the relationship among servers in a DNS zone distribution or a DHCP failover configuration. In the web UI, you can view a related servers page when you click the **Related Servers** icon on various pages. You can use the display of related servers to diagnose and monitor misconfigured or unreachable servers.

Related Topics

[Monitoring Remote Servers Using Persistent Events, on page 153](#)

[DNS Zone Distribution Servers, on page 155](#)

[DHCP Failover Servers, on page 155](#)

Monitoring Remote Servers Using Persistent Events

To service clients that require updates to DNS and LDAP related servers, the DHCP server uses a persistent event algorithm to ensure updates to related servers when a related server is temporarily unavailable. In

In addition, the algorithm prevents a misconfigured or offline DNS server from using up all the available update resources.

At startup, the DHCP server calculates the number of related servers in the configuration that require persistent events. A preconfigured Maximum Pending Events attribute (an Expert mode attribute that specifies the number of in-memory events that is preset to 40,000) is divided by the number of servers to obtain a limit on the number of events permitted for each remote server. This calculation covers related DNS and LDAP servers (DHCP failover does not use persistent storage for events). The DHCP server uses this calculation to issue log messages and take the action described in the following table. The table shows a hypothetical case of a DHCP server with four related DNS servers each having a limit of 10K events.

Table 21: Persistent Event Algorithm

Event Reached	DHCP Server Action
50% of the calculated per-server limit (Maximum Pending Events value divided by the number of total related servers); for example, 5K events on a related server out of a total of 40K maximum pending events	Issues an INFO log message every 2 minutes, as long as the limits are exceeded: The queue of events for the <i>name</i> remote server at <i>address</i> has <i>x</i> events, and has reached the info limit of <i>y/2</i> events out of an upper limit of <i>y</i> events per remote server. The remote server may be misconfigured, inoperative, or unreachable.
100% of the calculated per-server limit and less than 50% of the Maximum Pending Events value; for example, 10K events on a related server, with fewer than 10K total maximum pending events	Issues a WARNING log message every 2 minutes, as long as the limits are exceeded: The queue of events for the <i>name</i> remote server at <i>address</i> has <i>x</i> events, has exceeded the limit of <i>y</i> events per remote server, but is below the limit of <i>z</i> total events in memory. The remote server may be misconfigured, inoperative, or unreachable.
100% of the calculated per-server limit and 50% or more of the Maximum Pending Events value; for example, 10K events on a related server, with 20K total maximum pending events	Issues an ERROR log message every 2 minutes, as long as the limits are exceeded: The queue of events for the <i>name</i> remote server at <i>address</i> has <i>x</i> events, and has grown so large that the server cannot continue to queue new events to the remote server. The limit of <i>y</i> events per remote server and <i>z/2</i> total events in memory has been reached. This and future updates to this server will be dropped. The current eventID <i>n</i> is being dropped. The server drops the current triggering event and all subsequent events with that server.
100% of the Maximum Pending Events value; for example, 40K events across all related servers	Issues an ERROR log message: The queue of pending events has grown so large that the server cannot continue to queue new events. The queue's size is <i>z</i> , and the limit is <i>z</i> . The server drops all subsequent events with all related servers.

SNMP traps and DHCP server log messages also provide notification that a related server is unreachable.

DNS Zone Distribution Servers

A DNS zone distribution simplifies creating multiple zones that share the same secondary server attributes. You can view and set the primary and secondary DNS servers in a zone distribution.

Local Basic or Advanced Web UI

From the **Deploy** menu, click **Zone Distribution** under the **DNS** submenu. This opens the List/Add Zone Distributions page. The local cluster allows only one zone distribution, the default. Click this zone distribution name to open the Edit Zone Distribution page, which shows the authoritative and secondary servers in the zone distribution.

Regional Web UI

From the **Deploy** menu, choose **Zone Distribution** under the **DNS** submenu. This opens the List/Add Zone Distributions page. The regional cluster allows creating more than one zone distribution. Click the zone distribution name to open the Edit Zone Distribution page, which shows the name of the zone distribution map, primary, authoritative, and secondary servers in the zone distribution.



Note Default zone distribution names are not editable. However, non-default zone distribution names are editable and can be saved.

CLI Commands

Create a zone distribution using **zone-dist name create primary-cluster [attribute=value]**, then view it using **zone-dist list**. For example:

```
nrcmd> zone-dist distr-1 create Boston-cluster  
  
nrcmd> zone-dist list
```

DHCP Failover Servers

Related servers in a DHCP failover pair relationship can show the following information:

- **Type**—Main or backup DHCP server.
- **Server name**—DNS name of the server.
- **IP address**—Server IP address in dotted octet format.
- **Requests**—Number of outstanding requests, or two dashes if not applicable.
- **Communication status**—OK or INTERRUPTED.
- **Cluster state**—Failover state of this DHCP server.
- **Partner state**—Failover state of its partner server.

For details on DHCP failover implementation, see the *"Managing DHCP Failover"* section in *Cisco Prime Network Registrar 9.0 DHCP User Guide*

Local Basic or Advanced Web UI

From the **Deploy** menu, choose **Failover Pairs** under the **DHCP** submenu. The List/Add DHCP Failover Pairs page shows the main and backup servers in the failover relationship.

CLI Commands

Use **dhcp getRelatedServers** to display the connection status between the main and partner DHCP servers. If there are no related servers, the output is simply 100 Ok.

Displaying Leases

After you create a scope, you can monitor lease activity and view lease attributes.

Local Basic or Advanced Web UI

From the **Design** menu, choose **Scopes** under the **DHCPv4** submenu; or from the **Design** menu, choose **Prefixes** under the **DHCPv6** submenu. On the List/Add DHCP Scopes or List/Add DHCPv6 Prefixes page, click the **View** icon in the Leases column to open the List DHCP Leases for Scope or List DHCP Leases for Prefix page.

Regional Advanced Web UI

From the **Operate** menu, choose **DHCPv4 Lease History** or **DHCPv6 Lease History** under the **Reports** submenu. Set the query parameters and then query the lease history. (See the *"Querying Leases"* section in *Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide*.)

Troubleshooting DHCP and DNS Servers

The following sections describe troubleshooting the configuration and the DNS, DHCP, and TFTP servers.

Related Topics

[Immediate Troubleshooting Actions, on page 156](#)

[Modifying the cnr.conf File, on page 157](#)

[Troubleshooting Server Failures, on page 159](#)

[Troubleshooting and Optimizing the TFTP Server, on page 161](#)

[Linux Troubleshooting Tools, on page 160](#)

[Using the TAC Tool, on page 160](#)

Immediate Troubleshooting Actions

When facing a problem, it is crucial not to cause further harm while isolating and fixing the initial problem. Here are things to do (or avoid doing) in particular:

- Have 512 MB or more of memory and 2.5 GB or more of a data partition.
- Do not reboot a cable modem termination system (CMTS).
- Enable DHCP failover.
- Do not reload, restart, or disrupt Cisco Prime Network Registrar with failover resynchronization in progress.

Modifying the cnr.conf File

Cisco Prime Network Registrar uses the **cnr.conf** file for basic configuration parameters. This file is normally located in the *install-path* /conf directory. Cisco Prime Network Registrar creates the file during installation and processes it line by line.

You can edit this file if configuration parameters change. Note that during normal operation, you would not want to change the values. However, certain conditions might require you to modify certain values, such as when you move the data files for disk space reasons.

The format of the **cnr.conf** file consists of parameter name-value pairs, one per line; for example, for a Windows local cluster installation:

```
cnr.rootdir=C:\\Program Files\\Network Registrar\\Local
cnr.ccm-port=1234
cnr.cisco-gss-appliance-integration=n
cnr.datadir=C:\\NetworkRegistrar\\Local\\data
cnr.java-home=C:\\Program Files\\Java\\jre1.5.0_12
cnr.logdir=C:\\NetworkRegistrar\\Local\\logs
cnr.https-port=8443
cnr.tempdir=C:\\NetworkRegistrar\\Local\\temp
cnr.http-port=8080
cnr.ccm-mode=local
cnr.ccm-type=cnr
cnr.http-enabled=y
cnr.https-enabled=n
cnr.keystore-file=C:
cnr.keystore-password=unset
cnr.backup-time=23:45
```

Directory paths must be in the native syntax for the operating system. The format allows the use of colons (:) in directory paths, but not as name-value pair separators; it does not allow line continuation or embedded unicode characters. Other modifications to the file might include the location of the log directory (see [Log Files, on page 132](#)) or the time cnr_shadow_backup backups should occur (see [Setting Automatic Backup Time, on page 165](#)).

In rare cases, you might want to modify the file; for example, to exclude certain data from daily backups due to capacity issues. To do this, you need to add the appropriate settings manually.



Caution

We recommend that you use the default settings in this file. If you must change these settings, do so only in consultation with the Cisco Technical Assistance Center (TAC) or the Cisco Prime Network Registrar development team.

The following settings are supported:

- **cnr.backup-dest**—Specify the destination to place backed up databases. Defaults to cnr.datadir if not specified.
- **cnr.backup-dbs**—Provide a comma-separated list of the databases you want to backup. For a local cluster the default is ccm,dhcp,dns,mcd. For a regional cluster it is ccm,lease6hist,leasehist,subnetutil,replica.
- **cnr.backup-files**—Provide a comma-separated list of files and the complete path to the files that you want copied as part of the backup. Files are copied to cnr.backup-dest.
- **cnr.dbrecover-backup**—Specify whether to run db recover and db verify on a backed up Oracle Berkeley database. The default is true. This setting is used for daily backups only. Manual backups ignore this setting. Disabling the automatic operation means that you must run the operation manually, preferably on a separate machine, or at a time when the Cisco Prime Network Registrar servers are relatively idle.

- `cnr.daily-backup`—Specify whether to run the daily back up. The default is true.

Modifying the cnr.conf File for Syslog Support

Cisco Prime Network Registrar supports logging to a Syslog server (on Linux). The Syslog support is not enabled by default. To configure which messages need to be logged, based on logging levels, the `cnr.conf` file must be updated.

In addition, on Windows, event logging for Warnings and Errors is enabled by default (for Windows Event log). In this release, you can log more (or less) to the event log by changing the log settings.

The following `cnr.conf` configuration parameters are supported:

- `cnr.syslog.enable`—Specifies whether logging to Syslog server or Windows Event log is enabled for Prime Network Registrar servers.
 - To disable all logging, the value can be 0, off, or disabled.
 - To enable all logging, the value can be 1, on, or enabled.
 - By default, this parameter is disabled for Linux and enabled for Windows.
- `cnr.syslog.levels`—Specifies the severity levels to be logged to Syslog or Windows Event log. If Syslog is enabled, this defaults to warning and error. The value can be a case-blind, comma separated, list of the following keywords: error, warning, info, activity, and debug. This parameter is ignored if Syslog is disabled.



Note While it is possible to enable all of the severity levels and thus all messages written to the server log files are also logged to Syslog, this is not recommended unless the server log settings are reviewed and minimized. The performance impact on Syslog and the servers may vary greatly depending on how logging is configured. Syslog may rate limit the messages, so useful messages may also be lost.

- `cnr.syslog.facility`—Specifies the facility under which Syslog logs (Linux OS). This parameter is ignored for Windows. The valid facility keywords are daemon (the default), local0, local1, local2, local3, local4, local5, local6, local7.



-
- Note**
- These parameters apply to all Cisco Prime Network Registrar servers (`cnrservagt`, `ccm`, `cdns`, `cnrsnmp`, `dns`, `dhcp`, and `tftp`).
 - To apply any change to the `cnr.conf` parameters, Cisco Prime Network Registrar must be restarted.
-

The following `cnr.conf` configuration parameters allow server-specific overrides of the above parameters. server is one of `cnrservagt`, `ccm`, `cdns`, `cnrsnmp`, `dns`, `dhcp`, and `tftp`.

- `cnr.syslog.server.enable`—Specifies whether Syslog or Windows Event logging is enabled for the specified server (`cnr.syslog.enable` is ignored for that server).
- `cnr.syslog.server.levels`—Specifies the severity levels for the specified server (`cnr.syslog.levels` is ignored for that server).

- `cnr.syslog.server.facility`—Specifies the Syslog facility for the specified server (`cnr.syslog.facility` is ignored for that server).

The server specific configuration value is used, if specified. Otherwise, all parameters of the server are used. For example, to enable Syslog only for DHCP, add the following to the `cnr.conf` file:

```
cnr.syslog.dhcp.enable=1
```

To change the severity levels to include all non-debug logging (this assumes logging has been enabled for some or all servers), use:

```
cnr.syslog.enable=1
cnr.syslog.levels=error,warning,info,activity
```

To enable Syslog only for the DNS server:

```
cnr.syslog.dns.enable=1
cnr.syslog.dns.levels=error,warning,info,activity
```



Tip Syntax or other errors in the `cnr.conf` parameters are not reported and are ignored (that is, if a `levels` keyword is mistyped, that keyword is ignored). Therefore, if a configuration change does not work, check if the parameter(s) have been specified correctly.

Troubleshooting Server Failures

The server agent processes (`nwreglocal` and `nwregregion`) normally detect server failures and restart the server. You can usually recover from the failure and the server is not likely to fail again immediately after restarting. On rare occasions, the source of the server failure prevents the server from successfully restarting, and the server fails again as soon as it restarts. In such instances, perform the following steps:

Step 1 If the server takes a significantly long time to restart, stop and restart the server agent. On:

- Windows:

```
net stop nwreglocal or nwregregion
net start nwreglocal or nwregregion
```

- Linux:

```
/etc/rc.d/init.d/nwreglocal stop or nwregregion stop
/etc/rc.d/init.d/nwreglocal start or nwregregion start
```

Step 2 Keep a copy of all the log files. Log files are located in the *install-path*/logs directory on Linux, and the *install-path*\logs folder on Windows. The log files often contain useful information that can help isolate the cause of a server failure.

Step 3 Use the TAC tool, as described in [Using the TAC Tool, on page 160](#), or save the core or user.dmp file, if one exists, depending on the operating system:

- **Windows**—The `user.dmp` file is located in the system directory, which varies depending on the Windows system. Search for this file and save a renamed copy.
- **Linux**—The core file is located in the *install-path*. Save a renamed copy of this file that Cisco Prime Network Registrar does not overwrite.

- Step 4** On Windows, use the native event logging application to save the System and Application event logs to files. You can do this from the Event Viewer. These event logs often contain data that helps debug Cisco Prime Network Registrar server problems. For a description of the log messages for each server module, see the *install-path* /docs/msgid/MessageIdIndex.html file.

Linux Troubleshooting Tools

You can also use the following commands on Linux systems to troubleshoot Cisco Prime Network Registrar. To:

- See all Cisco Prime Network Registrar processes:

```
ps -leaf | grep nwr
```

- Monitor system usage and performance:

```
top
vmstat
```

- View login or bootup errors:

```
grep /var/log/messages*
```

- View the configured interfaces and other network data:

```
ifconfig -a
```

Using the TAC Tool

There may be times when any amount of troubleshooting steps will not resolve your problem and you have to resort to contacting the Cisco Technical Assistance Center (TAC) for help. Cisco Prime Network Registrar provides a tool so that you can easily assemble the server or system error information, and package this data for TAC support engineers. This eliminates having to manually assemble this information with TAC assistance. The resulting package from this tool provides the engineers enough data so that they can more quickly and easily diagnose the problem and provide a solution.

The **cnr_tactool** utility is available in the bin directory of the Windows, and usrbin directory of the UNIX or Linux, installation directories. Execute the **cnr_tactool** utility:

```
> cnr_tactool -N username -P password [-d output-directory] [-n]
```

The output directory is optional and normally is the temp directory of the installation directories (in the /var path on Linux). You may specify the **-n** option to indicate that when the cnr_exim tool is run, it is run without exporting any resource records (this specifies the -a none option to cnr_exim). If you do not supply the username and password on the command line, you are prompted for them:

```
> cnr_tactool
```

```
user:
```

```
password:  
[processing messages....]
```

The tool generates a packaged tar file whose name includes the date and version. The tar file contains all the diagnostic files.

Troubleshooting and Optimizing the TFTP Server

You can set certain attributes to troubleshoot and optimize TFTP server performance.

Related Topics

[Tracing TFTP Server Activity, on page 161](#)

[Optimizing TFTP Message Logging, on page 161](#)

[Enabling TFTP File Caching, on page 162](#)

Tracing TFTP Server Activity

To trace TFTP server activity, set the *packet-trace-level* attribute to a value of 1 through 4, depending on the level of verbosity you want the TFTP server to use to write messages to the trace file. The trace files are located in the /logs subdirectory of the installation directory. Windows tracing goes to the file_tftp_1_log file; Linux tracing goes to the /var/nwreg2/{local | regional}/logs/file_tftp_1_log and file_tftp_1_trace files.

Here are the trace levels, with each higher level being cumulative:

- **0**—Disables all server tracing (the default).
- **1**—Displays all the log messages in the trace file.
- **2**—Displays the client IP address and port number for all packets.
- **3**—Displays the packet header information.
- **4**—Displays the first 32 bytes of the packet.



Note Setting and getting the trace level only works if the TFTP server is started. Turn on packet tracing only for debugging purposes, and then not for any extended time, for performance reasons.

Optimizing TFTP Message Logging

You can improve TFTP server performance by restricting logging and tracing. By default, the server logs error, warning, and informational messages to file_tftp_1_log files. You can set the log levels using a few TFTP server parameters:

- **Log level** (use the *log-level* attribute)—Master controller of server logging, which is preset to, and is best left at, level 3 (logs all error, warning, and informational messages). As with packet tracing, the higher logging levels are cumulative. If set to 0, no server logging occurs.
- **Log settings** (use the *log-settings* attribute)—This is the second level of logging control and takes only two values, *default* or *no-success-messages*. The *default* log setting does not alter the default value of log level 3 (error, warning, and informational messages). However, you may want to disable writing

success informational messages, and thereby improve server performance, by changing the log settings to *no-success-messages*.

- **Log file count and size** (use the *log-file-count* attribute)—Sets how many log files to maintain and how large to allow them to get in the /logs directory. The default value is to maintain a maximum of ten files of 10 MB each.



Note Reload the TFTP server after changing these values.

Enabling TFTP File Caching

You can improve TFTP server performance significantly by enabling file caching on the server. You must do this explicitly, because it is preset to disabled. You must also create and point to a file cache directory, and you can set the maximum size of this directory. Here are the steps:

-
- Step 1** Determine where you want the TFTP cache files to go. This becomes a subdirectory of the TFTP home directory, which is preset to *install-path* /data/tftp (on Linux, it is /var/nwreg2/{local | regional}/data/tftp). If you want a different location, set the *home-directory* attribute.
 - Step 2** Change to the TFTP home directory and create the cache directory, such as CacheDir, in the home directory, using the **mkdir CACHEDIR** command. Note that Cisco Prime Network Registrar ignores any files in any subdirectories of this cache directory.
 - Step 3** Use the *file-cache-directory* attribute to set up the TFTP server to point to the cache directory. You cannot use absolute path or relative path in the directory name. The file-cache-directory name is either appended to the path given in the home-directory or the default home directory path (if you do not specify one).
 - Step 4** Use the *file-cache-max-memory-size* attribute to set the maximum memory size, in bytes, of the cache. The preset value is 32 KB. Cisco Prime Network Registrar loads all files into cache that cumulatively fit this memory size. If you set the value to 0, Cisco Prime Network Registrar does not cache any data, even if you enable file caching.
 - Step 5** Copy all of the files you want cached into the cache directory, and not into any subdirectory. Because all files in this directory are loaded into cache, do not include large files.
 - Step 6** Enable the *file-cache* attribute to enable file caching, then reload the server. Cisco Prime Network Registrar logs the name of each cached file, and skips any it cannot load. It reads in all files as binary data and translates them as the TFTP client requests. For example, if a client requests a file as NetASCII, the client receives the cached data in that form.
 - Step 7** Writing to cache is not allowed. If you need to update a cache file, overwrite it in the cache directory, then reload the server.
-



CHAPTER 9

Backup and Recovery

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

- [Backing Up Databases, on page 163](#)
- [Syntax and Location, on page 164](#)
- [Backup Strategy, on page 164](#)
- [Backing Up CNRDB Data, on page 166](#)
- [Database Recovery Strategy, on page 167](#)
- [Virus Scanning While Running Cisco Prime Network Registrar, on page 170](#)
- [Troubleshooting Databases, on page 170](#)

Backing Up Databases

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

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

Related Topics

[Syntax and Location, on page 164](#)

[Backup Strategy, on page 164](#)

[Database Recovery Strategy, on page 167](#)

[Backing Up CNRDB Data, on page 166](#)

[Backing Up All CNRDBs Using tar or Similar Tools, on page 166](#)

[Recovering CNRDB Data from Backups, on page 168](#)

[Recovering All CNRDBs Using tar or Similar Tools, on page 169](#)

[Recovering Single CNRDB from tar or Similar Tools, on page 169](#)

[Virus Scanning While Running Cisco Prime Network Registrar, on page 170](#)

Syntax and Location

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

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

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

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



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

Backup Strategy

The backup strategy involves either:

- Making CCM perform a nightly shadow backup for you (See the [Setting Automatic Backup Time, on page 165](#)) and using the shadow backups for permanent backup and then doing an explicit backup - either using the `cnr_shadow_backup` utility and backing up the backup files (*.bak DBs)

or

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

Manual Backup (Using `cnr_shadow_backup` utility)

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

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



Note If you change the location of the data directory, you must edit the `cnr.conf` file, which is located in `.../conf` (see [Modifying the cnr.conf File, on page 157](#)). Change the `cnr.datadir` variable to the full path to the data directory. For example, the following is the default value on Windows:

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

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

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

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

**Caution**

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

Related Topics

[Setting Automatic Backup Time, on page 165](#)

[Performing Manual Backups, on page 165](#)

[Using Third-Party Backup Programs with `cnr_shadow_backup`, on page 165](#)

Setting Automatic Backup Time

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

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

**Note**

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

Performing Manual Backups

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

**Note**

To restore DHCP data from a failover partner that is more up to date than a backup, see [Restoring DHCP Data from a Failover Server, on page 177](#).

Using Third-Party Backup Programs with `cnr_shadow_backup`

You should avoid scheduling third-party backup programs while `cnr_shadow_backup` is operating. Third-party backup programs should be run either an hour earlier or later than the `cnr_shadow_backup` operation. As described in [Setting Automatic Backup Time, on page 165](#), the default shadow backup time is daily at 23:45.

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

The operational files are listed in [Backup Strategy, on page 164](#). On Linux, Cisco Prime Network Registrar also maintains lock files in the following directories:

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

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

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

Backing Up CNRDB Data

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

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

The actual file naming convention is:

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

Backing Up All CNRDBs Using tar or Similar Tools

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

Step 1 Shut down Cisco Prime Network Registrar.

Backups cannot be done using tar or similar tools if Cisco Prime Network Registrar is running.

Step 2 Back up the entire data directory and subdirectories:

```
> /var/nwreg2/local/data or /var/nwreg2/regional/data
> /opt/nwreg2/*/conf
```

Step 3 Restart Cisco Prime Network Registrar when the backup is complete.

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

Database Recovery Strategy

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

Table 22: Cisco Prime Network Registrar Databases for Recovery

Subdirectory	Cluster	Type	Description
mcd	local	CNRDB	MCD change log data. Only exists for upgrades from pre 8.0 databases as long as there is MCD change log history that has not been trimmed.
ccm	local, regional	CNRDB	Central Configuration Management database. Stores local centrally managed cluster and the SNMP server data.
dns	local	CNRDB	DNS database. Stores zone state information, names of protected RRs, and zone configuration data for the DNS server.
cdns	local	CNRDB	Caching DNS database. Stores the initial DNSSEC root trust anchor and root hints.
dhcp ²	local	CNRDB	DHCP database. Stores lease state data for the DHCP server.
dhcpeventstore	local		Queue that Cisco Prime Network Registrar maintains to interact with external servers, such as for LDAP and DHCPv4 DNS Update interactions. Recovery is not necessary.
tftp	local		Default data directory for the TFTP server. Recovery is not necessary.
replica	regional	CNRDB	Stores replica data for the local clusters.
lease6hist	regional	CNRDB	DHCPv6 lease history database.

Subdirectory	Cluster	Type	Description
leasehist	regional	CNRDB	DHCPv4 lease history database.
subnetutil	regional	CNRDB	Subnet utilization database.

² Restoring the DHCP database (.../data/dhcp/ndb) from a backup is NOT RECOMMENDED. This is because, this data is constantly changing as the DHCP server is running (because of client activity and lease expirations either on this server or its partner). Therefore, restoring the DHCP ndb database would set the clock back in time for the server, but not for clients. Hence, it is best to retain the DHCP server database rather than recovering it, or if recovery is needed, delete it and recover the current leases from the partner via failover (see [Restoring DHCP Data from a Failover Server, on page 177](#)).

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

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

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

Recovering CNRDB Data from Backups

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

Step 1 Stop the Cisco Prime Network Registrar server agent.

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

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

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

Step 4 Rename the files.

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

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

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

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

Step 6 Restart the server agent.

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

After a successful database recovery, initiate an immediate backup and archive the files using the `cnr_shadow_backup` utility (see [Performing Manual Backups, on page 165](#)).

Recovering All CNRDBs Using tar or Similar Tools

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

- Step 1** Shut down Cisco Prime Network Registrar. Run `/etc/init.d/nwreglocal stop` to ensure that Cisco Prime Network Registrar is down.
- Step 2** Rename the active data directory (such as `mv data old-data`).
- Note** You must have sufficient disk space for twice the size of the data directory (and all the files in it and its subdirectories). If you do not have sufficient disk space, move the active data directory to another drive.
- Step 3** Create a new data directory and then untar or recover the backed up directory.
- We recommend that you run the CNRDB directory and recovery tools to ensure that the databases are good.
- Step 4** Start Cisco Prime Network Registrar.
- Note** Technically the restores do not need to include the `*.bak` directories (and subdirectories of those directories) as those contain nightly shadow backups. However, unless your available storage space is severely limited, we recommend a full restore of the entire data directory (and subdirectories) including the shadow backups.
-

Recovering Single CNRDB from tar or Similar Tools

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

- Step 1** Shut down Cisco Prime Network Registrar. Run `/etc/init.d/nwreglocal stop` to ensure that Cisco Prime Network Registrar is down.
- Step 2** Rename the active data directory (such as `mv data old-data`).
- Note** You must have sufficient disk space for twice the size of the data directory (and all the files in it and its subdirectories). If you do not have sufficient disk space, move the active data directory to another drive.
- Step 3** Create a new data directory and then untar or recover only the files in that directory (and its subdirectories) from the backup.
- We recommend that you run the CNRDB integrity and recovery tools to ensure that the CNRDB are good.
- Step 4** Repeat **Step 2** to **Step 3** for other DBs that have to be recovered.
- Step 5** Start Cisco Prime Network Registrar.
-

Virus Scanning While Running Cisco Prime Network Registrar

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

Troubleshooting Databases

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

Related Topics

[Using the `cnr_exim` Data Import and Export Tool, on page 170](#)

[Using the `cnrdb_recover` Utility, on page 172](#)

[Using the `cnrdb_verify` Utility, on page 173](#)

[Using the `cnrdb_checkpoint` Utility, on page 174](#)

[Using the `cnrdb_util` Utility, on page 174](#)

[Restoring DHCP Data from a Failover Server, on page 177](#)

Using the `cnr_exim` Data Import and Export Tool

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

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

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



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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Table 23: cnr_exim Options

Option	Description
<code>-a</code>	<p>Allows exporting and importing of protected or unprotected RRs. Valid <i>values</i> are: protectedRR, unprotectedRR, and none</p> <p>Export:</p> <p>All RRs are exported by default, so you must explicitly specify the export of protected or unprotected RRs using the option "<code>-a protectedRR</code>", "<code>-a unprotectedRR</code>", or "<code>-a none</code>". If this option is not specified, all RRs are exported.</p> <p>Import:</p> <p>All RRs are imported by default, so you must explicitly specify the import of protected or unprotected RRs using the option "<code>-a protectedRR</code>" or "<code>-a unprotectedRR</code>". If this option is not specified, all RRs are imported.</p>

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

Using the `cnrdb_recover` Utility

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

Table 24: `cnrdb_recover` Options

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

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

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

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

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

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

Using the `cnrdb_verify` Utility

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

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

Table 25: cnrdb_verify Options

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

Using the cnrdb_checkpoint Utility

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

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

Using the cnrdb_util Utility

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

The utility is located on the following directory:

- **Windows** — *(installation directory)*\bin\cnrdb_util.bat
- **Linux** — *(installation directory)*/userbin/cnrdb_util



Important

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

The cnrdb_util utility runs in two modes.

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

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

```
./cnrdb_util -h
```

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

Table 26: cnrdb_util Operations

Operation	Description
-d	Dump one or all CPNR databases.
-l	Load one or all CPNR databases.
-b	Create shadow backup of all CPNR databases.
-r	Recover one or all CPNR databases from shadow backup.
-c	Cleanup sleepycat log files in one or all CPNR databases.
-h	Display help text for the supported options.



Important

You can perform only one operation at a time.

Table 27: cnrdb_util Options

Option	Description
-m { local regional }	Specifies the CPNR installation mode. If not specified, this information is read from the cnr.conf file. If the file is not found, local mode is used by default.
-prog path	Specifies the path to the dump, load, or shadow backup executable. If not specified, this will be derived from the CPNR installation path.
-db db-path	Specifies the source database path for the '-d' dump, '-l' load, or '-r' recover operation. If not specified, the datadir entry from the cnr.conf file, or the current directory is used. This option is not applicable for the '-b' backup operation.

Option	Description
-n { ccm dhcp dns mcd leasehist lease6hist replica subnetutil all }	Specifies the name of the source database for the '-d' dump, '-l' load, or '-r' recover operation. If not specified, the operation will be performed on all databases present in database path. This option is not applicable for the '-b' backup operation. <ul style="list-style-type: none"> Valid database names for local mode are { ccm dhcp dns mcd all } Valid database names for regional mode are { ccm dns leasehist lease6hist replica subnetutil all }
-s	Specifies that this program should attempt to stop the CPNR Server Agent, if it is running.
-out path	Specifies the destination path for output files. If not specified, the source db path is used. This option is not applicable for the '-b' backup and '-c' cleanup operations.
-db_pagesize number	Specifies the size of database pages (in bytes) to be used when creating new databases. <p>The minimum page size is 512 bytes and the maximum page size is 64K bytes, and must be a power of two. If no page size is specified, a page size is selected based on the underlying filesystem I/O block size. (A page size selected in this way has a lower limit of 512 bytes and an upper limit of 16K bytes.)</p> <p>Usually the default is appropriate. However, large page sizes may not have good performance. 4096 and 8192 are typically good sizes. You can determine the page size of the database by using the cnrdb_stat utility.</p>

**Important**

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

**Note**

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

Restoring DHCP Data from a Failover Server

You can restore DHCP data from a failover server that is more current than the result of a shadow backup. Be sure that the failover partner configurations are synchronized. Also, ensure that the following steps are run on the bad failover partner (i.e., the one whose database is bad) and that you want to restore to.

On Windows

1. Set the default path; for example:

```
SET PATH=%PATH%;.;C:\PROGRA~1\NETWOR~1\LOCAL\BIN
```

2. Stop the server agent:

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

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

```
del C:\NetworkRegistrar\Local\data\dhcpeventstore\*.*
del C:\NetworkRegistrar\Local\data\dhcp\ndb\dhcp.ndb
del C:\NetworkRegistrar\Local\data\dhcp\ndb\logs\*.*
```

4. Restart the server agent:

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

On Linux

1. Stop the server agent:

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

2. Determine the processes running:

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

3. Kill the remaining processes:

```
kill -9 pid
```

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

```
rm /var/nwreg2/data/dhcpeventstore/*.*
```

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

5. Restart the server agent:

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




CHAPTER 10

Managing Reports

This chapter explains how to manage the Cisco Prime Network Registrar address space reporting tool, which is available from a regional cluster by using the web UI. Before you proceed with this chapter, become familiar with the concepts in the previous chapters of this part of the User's Guide.

- [ARIN Reports and Allocation Reports, on page 179](#)
- [Managing ARIN Reports, on page 179](#)
- [Managing IPv4 Address Space Utilization Reports, on page 183](#)
- [Managing Shared WHOIS Project Allocation and Assignment Reports, on page 184](#)
- [Managing BYOD Reports, on page 184](#)
- [Registered Devices, on page 185](#)
- [Scopes/Prefix , on page 185](#)

ARIN Reports and Allocation Reports

Using the Cisco Prime Network Registrar web UI, you can generate:

- American Registry of Internet Numbers (ARIN) reports, including:
 - Organization and point of contact (POC) reports
 - IPv4 address space utilization reports
 - Shared WHOIS project (SWIP) allocation and assignment reports
- Allocation reports that show how addresses are deployed across the routers and router interfaces of your network, including:
 - Allocation by owner reports
 - Allocation by router interface or by network reports

Managing ARIN Reports

ARIN, which is one of the five Regional Internet Registries (RIRs), manages IP resources in Canada, the United States of America, and many Caribbean and North Atlantic islands.

ARIN allocates blocks of IP addresses to Internet Service Providers (ISPs), which, in turn, reassign blocks of address space to their customers. ARIN distinguishes between *allocating* IP address space and *assigning* IP address space. It allocates address space to smaller IRs for subsequent distribution to the IRs' members and

customers. It assigns address space to an ISP, or other organization, for use only within the network of that organization and only for the purposes documented in its requests and reports to ARIN.

**Note**

ARIN manages IP address resources under the auspices of the Internet Corporation for Assigned Names and Numbers (ICANN). In other geographies, ICANN has delegated authority for IP resources to different regional Internet Registries. Cisco Prime Network Registrar does not currently support the reports that these registries might require, nor does it now support IPv6 reports or autonomous system (AS) numbers.

ARIN maintains detailed documentation about its policies and guidelines on its website.

<http://www.arin.net>

Be sure that you are familiar with these policies and guidelines before proceeding with ARIN reports.

The three options that you can specify for ARIN reports are:

- **New**—For a newly added POC or organization.
- **Modify**—Includes changed POC or organization data, such as phone numbers and addresses.
- **Remove**—Signals that you want to remove the POC or organization from the ARIN database.

Related Topics

[Managing Point of Contact and Organization Reports, on page 180](#)

[Managing IPv4 Address Space Utilization Reports, on page 183](#)

[Managing Shared WHOIS Project Allocation and Assignment Reports, on page 184](#)

Managing Point of Contact and Organization Reports

Cisco Prime Network Registrar provides reports that can submit Points of Contact (POC) and organizational information to ARIN. After you fill in these reports, you need to e-mail the information to ARIN. Submit the POC report (also called a template) to ARIN before preparing other reports.

Each POC is uniquely identified by a name called a POC handle and is associated with one or more Organization Identifiers (Org IDs) or resource delegations, such as an IP address space allocation or assignment. A POC handle, which ARIN assigns, can represent either an individual or a role.

The Organization report creates an Org ID and associates POC records with it. Create the Organization report after you create the POC report.

To manage POC and organization reports, log into the Cisco Prime Network Registrar regional web UI as a member of an administrator group assigned to the regional-addr-admin role.

Related Topics

[Creating a Point of Contact Report, on page 181](#)

[Registering a Point of Contact, on page 181](#)

[Editing a Point of Contact Report, on page 181](#)

[Creating an Organization Report, on page 182](#)

[Registering an Organization, on page 182](#)

[Editing an Organization Report, on page 183](#)

Creating a Point of Contact Report

You create POCs so that managers can interact with ARIN to request and administer IP resources and so that network professionals can manage network operation issues.

Regional Web UI

- Step 1** From the **Administration** menu, choose **Contacts** under the **Settings** submenu to open the List/Add ARIN Points of Contact page.
 - Step 2** Click the **Add Contact** icon in the Contacts pane on the left, to open the Add Point of Contact page.
 - Step 3** Enter data in the fields on the page:
 - **Name**—A unique identifier for the POC (required).
 - **First Name**—The first name of the point of contact (required).
 - **Last Name**—The last name of the point of contact (required).
 - **Type**—From the drop-down list, choose Person or Role (optional, with preset value Person).
 - Step 4** Click **Add Point of Contact**.
-

Registering a Point of Contact

You must register the POC with ARIN to receive a POC handle.

Regional Web UI

- Step 1** From the **Administration** menu, choose **Contacts** under the **Settings** submenu to open the List/Add ARIN Points of Contact page.
 - Step 2** Click the required contact in the Contacts pane on the left.
 - Step 3** Click the **Register Report** tab to view the ARIN template file.
 - Step 4** Copy and paste the template file into an e-mail and send the file to ARIN.
-

Editing a Point of Contact Report

Edit a POC report after ARIN returns a POC handle to your organization or if your POC has changed.

Regional Web UI

- Step 1** From the **Administration** menu, choose **Contacts** under the **Settings** submenu to open the List/Add ARIN Points of Contact page.
- Step 2** Click the required contact in the Contacts pane on the left. The Edit Point of Contact page opens.
- Step 3** Enter values for Middle Name, Handle, and Description (optional)

- Step 4** In the Poc Emails field:
- Enter the e-mail address for the POC.
 - Click **Add Email Address** to add additional e-mail addresses.
- Step 5** In the Poc Phones field:
- Enter a phone number and extension, if applicable, then choose a type (Office, Mobile, Fax, or Pager) from the drop-down list,
 - Click **Add Phone** to add additional telephones.
- Step 6** Miscellaneous Settings. Add these additional attributes as strings or lists of text.
- Step 7** After making the changes, click **Save**.
-

Creating an Organization Report

Each organization is represented in the ARIN WHOIS database by a unique Org ID, consisting of an organization name, its postal address, and its POCs. While organizations may have more than one Org ID, ARIN recommends consolidating IP address resources under a single Org ID.

If you do not have an Org ID with ARIN, or you are establishing an additional Org ID, you must first create and submit a POC report. When ARIN confirms it has received your POC information, use Cisco Prime Network Registrar to complete an Organization form and submit that information.

Regional Web UI

-
- Step 1** From the **Administration** menu, choose **Organizations** under the **Settings** submenu to open the List/Add ARIN Organizations page.
- Step 2** Click the **Add Organization** icon in the Organizations pane on the left, to open the Add Organization page.
- Step 3** Enter data in the fields on the page:
- **Organization Name**—Name of the organization that you want to register with ARIN.
 - **Description**—A text description of the organization.
 - **Organization Admin POC**—From the drop-down list, choose the POC who administers IP resources from the drop-down list.
 - **Organization Technical Points Of Contact**—From the drop-down list, choose one or more POCs who manage network operations, or click **Add Point of Contact** to add new contact information.
- Step 4** Click **Add Organization**. This opens the Edit Organization page where you can add more details.
-

Registering an Organization

You must register your Organization with ARIN to receive an Organization ID.

Regional Web UI

-
- Step 1** From the **Administration** menu, choose **Organizations** under the **Settings** submenu to open the List/Add ARIN Organizations page.

- Step 2** Click the required organization in the Organizations pane on the left.
- Step 3** Click the **Register Report** tab to view the ARIN template file.
- Step 4** Copy and paste the template file into an e-mail and send the file to ARIN.
-

Editing an Organization Report

You might need to change organizational information that you have registered with ARIN.

Regional Web UI

- Step 1** From the **Administration** menu, choose **Organizations** under the **Settings** submenu to open the List/Add ARIN Organizations page.
- Step 2** Click the required organization in the Organizations pane on the left.
- Step 3** Enter or change data in the fields.
- **Miscellaneous Settings**—Add these additional attributes as strings or lists of text.
 - **Organization Abuse Points of Contact**—From the drop-down list, choose one or more POCs who handle network abuse complaints, or click **Add Point of Contact** to add new contact information.
 - **Organization NOC Points of Contact**—From the drop-down list, choose one or more POCs in network operations centers, or click **Add Point of Contact** to add new contact information.
- Step 4** Click **Save**.
- Step 5** Submit the updated report to ARIN as described in [Registering an Organization, on page 182](#).
-

Managing IPv4 Address Space Utilization Reports

Address space utilization reports serve two purposes:

- To make an initial request for IPv4 address space after you receive a POC handle and an Org ID.
- To support a request for an additional allocation of IPv4 addresses when your business projections show that you are running out of IP addresses.



Note The ARIN website contains extensive information about how it initially allocates address space and its threshold criteria for requesting additional address space. In general, for a single-homed organization, the minimum allocation from ARIN is a /20 block of addresses. For a multihomed organization, the minimum allocation is a /22 block of addresses. ARIN recommends that an organization requiring a smaller block of addresses contact an upstream ISP to obtain addresses.

The Cisco Prime Network Registrar utilization report corresponds to the ARIN ISP Network Request template (ARIN-NET-ISP-3.2.2).

Regional Web UI

- Step 1** From the **Operate** menu, choose **ARIN Address Space Usage** under the **Reports** submenu to open the Select Address Space Report page.
- Step 2** In the Select the Report Type field, choose **Utilization** from the drop-down list. The Select the Filter Type field is updated with the value, *by-owner*. The browser redisplay the Select Address Space Report page with two new fields: Network Name and Network Prefix Length.
- Step 3** In the Select Owner field, choose the owner of this address block from the drop-down list.
- Step 4** Enter values for the Network Name and Network Prefix Length.
- Step 5** Click **Generate Report**. The browser displays an ARIN template file (ARIN-NET-ISP-3.2.2).
Several sections of the report require that you manually enter data because the information is generated and maintained outside the Cisco Prime Network Registrar application.
- Step 6** Click **Save Report**. The browser displays the Address Space Utilization Report as an unformatted text file.
- Step 7** Copy the Address Space Utilization Report to a text editor to manually enter the data that Cisco Prime Network Registrar does not generate.
- Step 8** Copy and paste the edited report into an e-mail and send the file to ARIN.
-

Managing Shared WHOIS Project Allocation and Assignment Reports

The ARIN shared WHOIS project (SWIP) provides a mechanism for finding contact and registration information for resources registered with ARIN. The ARIN database contains IP addresses, autonomous system numbers, organizations or customers that are associated with these resources, and related POCs.

The ARIN WHOIS does not locate any domain- or military-related information. Use whois.internic.net to locate domain information, and whois.nic.mil for military network information.

The regional web UI also provides two allocation and assignment report pages:

- View ARIN SWIP Reallocated Report
- View ARIN SWIP Reassigned Report

Managing BYOD Reports

There are two types of BYOD reports:

- Registered Devices
- Scopes/Prefix

Registered Devices

Registered Device report displays the list of devices that are registered through BYOD web server. The report can be exported in the csv format. Only an admin user is allowed to delete a device using the Registered Device Report page.

Registered Devices Report

To access the Registered Devices Report:

Regional Advanced or Expert Web UI

From the **Operate** menu, choose **BYOD Registered Devices** under the **Reports** submenu to access the report in the List BYOD Registered Devices page.

LDAP server(s) configured to the BYOD web server or local DHCP server(s) or failover pairs, associated with the regional server will be listed in the clusters pane. All the registered devices in the LDAP server or devices registered in the local DHCP servers or failover pairs through the BYOD web server will be displayed in the List BYOD Registered Devices page.



Note You must select the desired server from the cluster pane to view the corresponding registered devices report in the List BYOD Registered Devices page.

Scopes/Prefix

Scope/Prefix report displays the list of scopes and prefixes that are used for BYOD. The report can be exported in the csv format.

Scope/Prefix Report

To view the Scope/Prefix Report:

Regional Advanced or Expert Web UI

From the **Operate** menu, choose **BYOD Scopes/Prefix** under the **Reports** submenu to view the report in the List BYOD Scope/Prefix page.

Local DHCP server(s) or failover pairs associated with the regional server will be listed in the clusters pane. All the scopes and prefixes created in the local DHCP servers or failover pairs for the BYOD web server will be displayed in the List BYOD Scope/Prefix page.



Note You must select the desired server from the cluster pane to view the corresponding scopes and prefixes created during BYOD setup in the List BYOD Scope/Prefix page.



PART **III**

Virtual Appliance

- [Introduction to Cisco Prime Network Registrar Virtual Appliance, on page 189](#)



CHAPTER 11

Introduction to Cisco Prime Network Registrar Virtual Appliance

The Cisco Prime Network Registrar virtual appliance aims at reducing the installation, configuration, and maintenance costs associated with running Cisco Prime Network Registrar on a local system. It also guarantees portability and thus reduces the risk in moving Cisco Prime Network Registrar from one machine to another.

You must get a license for Cisco Prime Network Registrar and download the virtual appliance from Cisco.com. Every Cisco Prime Network Registrar local cluster must be connected to a regional cluster which contains the licenses for the DHCP or DNS services provided by the local cluster. All licenses are loaded into the regional cluster, and local clusters are registered with the regional cluster at the time of their first installation. Cisco Prime Network Registrar will then be up and running, available to be configured.

This is different from just downloading a copy of Cisco Prime Network Registrar and installing it on a server or virtual machine provided by the customer, in that the operating system on which Cisco Prime Network Registrar runs is also provided in the virtual appliance.

The Cisco Prime Network Registrar virtual appliance is supported on VMware ESXi 5.5 or later platforms, CentOS/RHEL 7.2 Hypervisor, and an OpenStack installation running on CentOS/RHEL 7.2.

To know about the difference between vApp and a virtual appliance, see the *User's Guide to Deploying vApps and Virtual Appliances*.

- [How the Cisco Prime Network Registrar Virtual Appliance Works, on page 189](#)
- [Invoking Cisco Prime Network Registrar on the Virtual Appliance, on page 190](#)
- [Monitoring Disk Space Availability on VMware, on page 190](#)
- [Increasing the Size of the Disk on VMware, on page 190](#)
- [Increasing the Size of the Disk on a KVM Hypervisor, on page 191](#)
- [Troubleshooting, on page 192](#)

How the Cisco Prime Network Registrar Virtual Appliance Works

The virtual appliance consists of a virtual machine, which contains a runnable guest OS (CentOS 7.2) and Cisco Prime Network Registrar installed on that OS. When the virtual appliance is installed, Cisco Prime Network Registrar is already installed and is started by the virtual machine power-up.

Invoking Cisco Prime Network Registrar on the Virtual Appliance

You can invoke the Cisco Prime Network Registrar application directly by using the URL `http://hostname:8080`. The secure `https` connection is also available via the URL `https://hostname:8443`.

Monitoring Disk Space Availability on VMware

To determine how much space is available to use for increasing the size of a virtual appliance's disk, do the following:

-
- Step 1** In the vSphere Client window, select the host/server on which the virtual Cisco Prime Network Registrar appliance resides.
 - Step 2** Click **Storage Views** to see the list of the machines hosted by the server and the details about the space currently used by each machine.
Also, you can go to the Virtual Machines tab to view both the **Provisioned Space** and the **Used Space** by machine.
 - Step 3** Click **Summary**.
The **Resources** area of the Summary tab, displays the capacity of the disk and the CPU and memory used.
 - Step 4** Select the virtual machine and click the **Summary** tab.
The **Resources** area of the Summary tab displays the disk space details for the machine.
-

Monitoring Disk Space Availability in Use by the Virtual Appliance

To determine how much free space is left on the disk in use by the virtual appliance, as an aid to determine if you should increase the size of the virtual appliance's disk, do the following:

-
- Step 1** Select the virtual machine in the vSphere Client window and either click the **Console** tab on the right pane or right-click the virtual machine name and choose **Open Console**.
 - Step 2** Log in as root and type `df -k`. The disk space details are displayed.
If the disk space on the disk mounted is not enough, then you should increase the size of the disk (see [Increasing the Size of the Disk on VMware, on page 190](#)).
-

Increasing the Size of the Disk on VMware

If you need a bigger disk, do the following:

Glossary

A	
A record	DNS Address resource record (RR). Maps a hostname to its address and specifies the Internet Protocol address (in dotted decimal form) of the host. There should be one A record for each host address.
access control list (ACL)	DHCP mechanism whereby the server can allow or disallow the request or action defined in a packet. <i>See also</i> transaction signature (TSIG) .
address block	Block of IP addresses to use with DHCP subnet allocation that uses on-demand address pools.
admin	Default name of the superuser or global administrator.
administrator	User account to adopt certain functionality, be it defined by role, constrained role, or group.
alias	Pointer from one domain name to the official (canonical) domain name.
allocation priority	An alternate method of control over allocating addresses among scopes other than the default round-robin method.
ARIN	American Registry of Internet Numbers, one of several regional Internet Registries (IRs), manages IP resources in North America, parts of the Caribbean, and subequatorial Africa. Cisco Prime Network Registrar provides an address space report for this registry.
Asynchronous Transfer Mode (ATM)	International standard for cell relay in which multiple service types (such as voice, video, or data) are conveyed in fixed-length (53-byte) cells.
authoritative name server	DNS name server that possesses complete information about a zone.
AXFR	Full DNS zone transfer. <i>See also</i> zone transfer and IXFR .
B	
Berkeley Internet Name Domain (BIND)	Implementation of the Domain Name System (DNS) protocols. <i>See also</i> DNS .
binding	Collection of DHCP client options and lease information, managed by the main and backup DHCP servers. A binding database is a collection of configuration parameters associated with all DHCP clients. This database holds configuration information about all the datasets.
BOOTP	Bootstrap Protocol. Used by a network node to determine the IP address of its Ethernet interfaces, so that it can affect network booting.
C	

cable modem termination system (CMTS)	Cable modem termination system. Either a router or bridge, typically at the cable head end.
cache	Data stored in indexed disk files to reduce the amount of physical memory.
caching name server	Type of DNS server that caches information learned from other name servers so that it can answer requests quickly, without having to query other servers for each transaction.
canonical name	Another name for an alias DNS host, inherent in a CNAME resource record (RR).
case sensitivity	Values in Cisco Prime Network Registrar are not case-sensitive, with the exception of passwords.
Central Configuration Management (CCM) database	Main database for the Cisco Prime Network Registrar web-based user interface (web UI).
chaddr	DHCP client hardware (MAC) address. Sent in an RFC 2131 packet between the client and server.
change logs, changesets	A change log is a group of changesets made to the Cisco Prime Network Registrar databases due to additions, modifications or deletions in the web UI. A changeset is a set of changes made to a single object in the database.
ciaddr	DHCP client IP address. Sent in an RFC 2131 packet between the client and server.
class of address	Category of an IP address that determines the location of the boundary between network prefix and host suffix. Internet addresses can be A, B, C, D, or E level addresses. Class D addresses are used for multicasting and are not used on hosts. Class E addresses are for experimental use only.
client-class	Cisco Prime Network Registrar feature that provides differentiated services to users that are connected to a common network. You can thereby group your user community based on administrative criteria, and then ensure that each user receives the appropriate class of service.
cluster	In Cisco Prime Network Registrar, a group of DNS, DHCP, and TFTP servers that share the same database.
CNAME record	DNS Canonical Name resource record (RR). Used for nicknames or aliases. The name associated with the resource record is the nickname. The data portion is the official or canonical name.
CNRDB	Name of one of the Cisco Prime Network Registrar internal databases. The other is changeset database.
constraint	Assigned limitation on the role or allowable functionality of an administrator.
D	

Data Over Cable Service Interface Specification (DOCSIS)	Data Over Cable Service Interface Specification. Standard created by cable companies in 1995 to work toward an open cable system standard and that resulted in specifications for connection points, called interfaces.
delegation	Act of assigning responsibility for managing a DNS subzone to another server, or of assigning DHCP address blocks to local clusters.
DHCP	Dynamic Host Configuration Protocol. Designed by the Internet Engineering Task Force (IETF) to reduce the amount of configuration that is required when using TCP/IP. DHCP allocates IP addresses to hosts. It also provides all the parameters that hosts require to operate and exchange information on the Internet network to which they are attached.
Digital Subscriber Line (DSL)	Public network technology that delivers high bandwidth over conventional copper wiring at limited distances.
DNS	Domain Name System. Handles the growing number of Internet users. DNS translates names, such as www.cisco.com, into Internet Protocol (IP) addresses, such as 192.168.40.0, so that computers can communicate with each other.
DNS update	Protocol (RFC 2136) that integrates DNS with DHCP.
domain	Portion of the DNS naming hierarchy tree that refers to general groupings of networks based on organization type or geography. The hierarchy is root, top- or first-level, and second-level domain.
domain name	DNS name that can be either absolute or relative. An absolute name is the fully qualified domain name (FQDN) and is terminated with a period. A relative name is relative to the current domain and does not end with a period.
dotted decimal notation	Syntactic representation of a 32-bit integer that consists of four eight-bit numbers written in base 10 with dots separating them for a representation of IP addresses. Many TCP/IP application programs accept dotted decimal notation in place of destination machine names.
E	
expression	Construct commonly used in the Cisco Prime Network Registrar DHCP implementation to create client identities or look up clients. For example, an expression can be used to construct a scope from a template.
extension and extension point	In Cisco Prime Network Registrar, element of a script written in TCP, C, or C++ that customizes handling DHCP packets as the server processes them, and which supports additional levels of customizing DHCP clients.
F	

failover	Cisco Prime Network Registrar feature (as described in RFC 2131) that provides for multiple, redundant DHCP servers, whereby one server can take over in case of a failure. DHCP clients can continue to keep and renew their leases without needing to know or care which server is responding to their requests.
forwarder	DNS server designated to handle all offsite queries. Using forwarders relieves other DNS servers from having to send packets offsite.
forwarding, DHCP	Mechanism of forwarding DHCP packets to another DHCP server on a per-client basis. You can achieve this in Cisco Prime Network Registrar by using extension scripting.
FQDN	Fully qualified domain name. Absolute domain name that unambiguously specifies a host location in the DNS hierarchy.
G	
giaddr	DHCP gateway (relay agent) IP address. Sent in an RFC 2131 packet between the client and server.
glue record	DNS Address resource record that specifies the address of a subdomain authoritative name server. You only need glue records in the server delegating a domain, not in the domain itself.
group	Associative entity that combines administrators so that they can be assigned roles and constrained roles.
H	
High-Availability (HA) DNS	DNS configuration in which a second primary server can be made available as a hot standby that shadows the main primary server.
HINFO record	DNS Host Information resource record (RR). Provides information about the hardware and software of the host machine.
hint server	See root hint server .
host	Any network device with a TCP/IP network address.
I	
IEEE	Institute of Electrical and Electronics Engineers. Professional organization whose activities include developing communications and network standards.
in-addr.arpa	DNS address mapping domain with which you can index host addresses and names. The Internet can thereby convert IP addresses back to hostnames. See also reverse zone .
IP address	Internet Protocol address. For example, 192.168.40.123.

IP history	Cisco Prime Network Registrar tool that records the lease history of IP addresses in a database.
IPv6	New IP standard involving 128-bit addresses. Cisco Prime Network Registrar provides a DHCPv6 implementation.
ISP	Internet Service Provider. Company that provides leased line, dialup, and DSL (Point-to-Point over Ethernet and DHCP) access to customers.
iterative query	Type of DNS query whereby the name server returns the closest answer to the querying server.
IXFR	Incremental zone transfer. Standard that allows Cisco Prime Network Registrar to update a slave (secondary) server by transferring only the changed data from the primary server.
L	
lame delegation	Condition when DNS servers listed in a zone are not configured to be authoritative for the zone.
LDAP	Lightweight Directory Access Protocol. Method that provides directory services to integrate Cisco Prime Network Registrar client and lease information.
lease	IP address assignment to a DHCP client that also specifies how long the client can use the address. When the lease expires, the client must negotiate a new one with the DHCP server.
lease grace period	Length of time the lease is retained in the DHCP server database after it expires. This protects a client lease in case the client and server are in different time zones, their clocks are not synchronized, or the client is not on the network when the lease expires.
link group	Groups the links to accommodate CMTS Prefix Stability. The <i>group-name</i> attribute is used to specify the name of the group to which the link should belong.
lease history	A report that can be generated to provide a historical view of when a client was issued a lease, for how long, when the client or server released the lease before it expired, and if and when the server renewed the lease and for how long.
lease query	Process by which a relay agent can request lease (and reservation) data directly from a DHCP server in addition to gleaning it from client/server transactions.

link type	There are three different link types: topological, location-independent, and universal. Topological links means a client is allocated leases based on the network segment it is connected to. While the location-independent link type lets a subscriber, that is moved from one CMTS to another within a central office, to retain a delegated prefix, the universal link type lets the subscriber moving from one central office to another to retain the delegated prefix.
local cluster	Location of the local Cisco Prime Network Registrar servers. <i>See also</i> regional cluster .
localhost	Distinguished name referring to the name of the current machine. Localhost is useful for applications requiring a hostname.
loopback zone	DNS zone that enables the server to direct traffic to itself. The host number is almost always 127.0.0.1.
M	
MAC address	Standardized data link layer address. Required for every port or device that connects to a LAN. Other devices in the network use these addresses to locate specific ports on the network and to create and update routing tables and data structures. MAC addresses are six bytes long and are controlled by the IEEE. Also known as a hardware address, MAC layer address, and physical address. A typical MAC address is 1,6,00:d0:ba:d3:bd:3b.
mail exchanger	Host that accepts electronic mail, some of which act as mail forwarders. <i>See also</i> MX record .
master name server	Authoritative DNS name server that transfers zone data to secondary servers through zone transfers.
maximum client lead time (MCLT)	In DHCP failover, a type of lease insurance that controls how much ahead of the backup server lease expiration the client lease expiration should be.
multinetting	State of having multiple DHCP scopes on one subnet or several LAN segments.
Multiple Service Operator (MSO)	Provides subscribers Internet access using cable or wireless technologies.
multithreading	Process of performing multiple server tasks.
MX record	DNS Mail Exchanger resource record (RR). Specifies where mail for a domain name should be delivered. You can have multiple MX records for a single domain name, ranked in preference order.
N	
nameserver	DNS host that stores data and RRs for a domain.

NAPTR	DNS Naming Authority Pointer resource record (RR). Helps with name resolution in a particular namespace and is processed to get to a resolution service. Based on proposed standard RFC 2915.
negative cache time	Memory cache the DNS server maintains for a quick response to repeated requests for negative information, such as "no such name" or "no such data." Cisco Prime Network Registrar discards this information at intervals.
network ID	Portion of the 32-bit IP address that identifies which network a particular system is on, determined by performing an AND operation of the subnet mask and the IP address.
NOTIFY	Standard (RFC 1996) whereby DNS master servers can inform their slaves that changes were made to their zones, and which initiates a zone transfer.
nrcmd	Cisco Prime Network Registrar command line interface (CLI).
O	
on-demand address pool	Wholesale IP address pool issued to a client (usually a VPN router or other provisioning device), from which it can draw for lease assignments. Also known as DHCP subnet allocation.
option, DHCP	DHCP configuration parameter and other control information stored in the options field of a DHCP message. DHCP clients determine what options get requested and sent in a DHCP packet. Cisco Prime Network Registrar allows for creating option definitions as well as the option sets to which they belong.
Organization report	One of the reports to be submitted to ARIN, POC being the other report. <i>See also</i> ARIN and POC report .
Organizationally Unique Identifier (OUI)	Assigned by the IEEE to identify the owner or ISP of a VPN. <i>See also</i> IEEE and virtual private network (VPN) .
owner	Owners can be created as distinguishing factors for address blocks, subnets, and zones. In the context of DNS RRs, an owner is the name of the RR.
P	
ping	Packet Internet Groper. A common method for troubleshooting device accessibility that uses a series of Internet Control Message Protocol (ICMP) Echo messages to determine if a remote host is active or inactive, and the round-trip delay in communicating with the host.
POC report	Point of Contact report. One of the reports to be submitted to ARIN, Organization being the other report. <i>See also</i> ARIN and Organization report .

policy	Group of DHCP attributes or options applied to a single scope or group of scopes. Embedded policies can be created for scopes and other DHCP objects.
polling	Collection of subnet utilization or lease history data over a certain regular period.
prefix allocation groups	Groups prefixes in order to facilitate the prioritization of prefix allocation.
prefix stability	Clients can retain the delegated prefix when they change their location, that is even when they move from one CMTS to another (CMTS Prefix Stability) or move within an address space (Universal Prefix Stability).
primary master	DNS server from which a secondary server receive data through a zone transfer request.
provisional address	Address allocated by the DHCP server to unknown clients for a short time, one-shot basis.
PTR record	DNS Pointer resource record. Used to enable special names to point to some other location in the domain tree. Should refer to official (canonical) names and not aliases. <i>See also</i> in-addr.arpa .
pulling and pushing objects	The Cisco Prime Network Registrar regional cluster provides functions to pull network objects from the replica database of local cluster data, and push objects directly to the local clusters.
R	
recursive query	DNS query where the name server asks other DNS server for any nonauthoritative data not in its own cache. Recursive queries continue to query all name servers until receiving an answer or an error.
refresh interval	Time interval in which a secondary DNS server checks the accuracy of its data by sending an AXFR packet to the primary server.
region	Regions can be created as distinguishing factors for address blocks, subnets, and zones. A region is distinct from the regional cluster.
regional cluster	Location of the regional Cisco Prime Network Registrar CCM server. <i>See also</i> local cluster .
relay agent	Device that connects two or more networks or network systems. In DHCP, a router on a virtual private network that is the IP helper for the DHCP server.
replica database	CCM database that captures copies of local cluster configurations at the regional cluster. These configurations can be pulled to the regional cluster so that they can be pushed to other local clusters.

Request for Comments (RFC)	TCP/IP set of standards.
reservation	IP address or lease that is reserved for a specific DHCP client.
resolution exception	Selectively forwarding DNS queries for specified domains to internal servers rather than recursively querying Internet root name and external servers.
resolver	Client part of the DNS client/server mechanism. A resolver creates queries sent across a network to a name server, interprets responses, and returns information to the requesting programs.
resource record (RR)	DNS configuration record, such as SOA, NS, A, CNAME, HINFO, WKS, MX, and PTR that comprises the data within a DNS zone. Mostly abbreviated as RR. <i>See the "Resource Records" section in Cisco Prime Network Registrar 9.0 Authoritative and Caching DNS User Guide</i>
reverse zone	DNS zone that uses names as addresses to support address queries. <i>See also in-addr.arpa.</i>
role, constrained role	Administrators can be assigned one or more roles to determine what functionality they have in the application. A constrained role is a role constrained by further limitations. There are general roles for DNS, host, address block, DHCP, and CCM database administration. You can further constrain roles for specific hosts and zones. Some roles have distinguishing subroles, such as the database subrole.
root hint server	DNS name server at the top of the hierarchy for all root name queries. A root name server knows the addresses of the authoritative name servers for all the top-level domains. Resolution of nonauthoritative or uncached data must start at the root servers. Sometimes called a hint server.
round-robin	Action when a DNS server rearranges the order of its multiple same-type records each time it is queried.
routed bridge encapsulation (RBE)	Process by which a stub-bridged segment is terminated on a point-to-point routed interface. Specifically, the router is routing on an IEEE 802.3 or Ethernet header carried over a point-to-point protocol, such as PPP, RFC 1483 ATM, or RFC 1490 Frame Relay.
S	
scavenging	Action of periodically scanning dynamic updates to the DNS server for stale resource records and purging these records.
scope	Administrative grouping of TCP/IP addresses on a DHCP server. Required for lease assignments.

secondary master	DNS name server that gets its zone data from another name server authoritative for the zone. When a secondary master server starts up, it contacts the primary master, from which it receives updates.
secondary subnet	A single LAN might have more than one subnet number applicable to the same LAN or network segment in a router. Typically, one subnet is designated as primary, the others as secondary. A site might support addresses on more than one subnet number associated with a single interface. You must configure the DHCP server with the necessary information about your secondary subnets.
selection tags	Mechanisms that help select DHCPv4 scopes and DHCPv6 prefixes for clients and client-classes.
siaddr	IP address of the server to use in the next step of the DHCP boot process. Sent in an RFC 2131 packet between the client and server.
slave forwarder	DNS server that behaves like a stub resolver and passes most queries on to another name server for resolution. <i>See also</i> stub resolver .
slave servers	DNS server that always forwards queries it cannot answer from its cache to a fixed list of forwarding servers instead of querying the root name servers for answers.
SNMP notification	Simple Network Management Protocol messages that warn of server error conditions and problems. <i>See also</i> trap .
SOA record	DNS Start of Authority resource record (RR). Designates the start of a zone.
SRV record	Type of DNS resource record (RR) that allows administrators to use several servers for a single host domain, to move services from host to host with little difficulty, and to designate some hosts as primary servers for a service and others as backups.
staged edit mode	dhcp or dns edit mode in which the data is stored on the CCM server, but not live on the protocol server. <i>See also</i> synchronous edit mode .
stub resolver	DNS server that hands off queries to another server instead of performing the full resolution itself.
subnet allocation, DHCP	Cisco Prime Network Registrar use of on-demand address pools for entire subnet allocation of IP addresses to provisioning devices.
subnet mask	Separate IP address, or part of a host IP address, that determines the host address subnet. For example, 192.168.40.0 255.255.255.0 (or 192.168.40.0/24) indicates that the first 24 bits of the IP address are its subnet, 192.168.40. In this way, addresses do not need to be divided strictly along network class lines.
subnet pool	Set of IP addresses associated with a network number and subnet mask, including secondary subnets.

subnet sorting	Attribute of the Cisco Prime Network Registrar DNS server. By enabling it, the server checks the network address of the client before responding to a query.
subnet utilization	A report that can be generated to determine how many addresses in the subnet were allocated and what the free address space is.
subnetting	Action of dividing any network class into multiple subnetworks.
subscriber limitation	Limitation to the number of addresses service providers can determine for the DHCP server to give out to devices on customer premises, handled in Cisco Prime Network Registrar by DHCP option 82 definitions.
subzones	Partition of a delegated domain, represented as a child of the parent node. A subzone always ends with the name of its parent. For example, boston.example.com. can be a subzone of example.com.
subzone delegation	Dividing a zone into subzones. You can delegate administrative authority for these subzones, and have them managed by people within those zones or served by separate servers.
supernet	Aggregation of IP network addresses advertised as a single classless network address.
synchronization	Synchronization can occur between the regional cluster and local clusters, the CCM and other protocol servers, failover servers, HA DNS servers, and routers.
synchronous edit mode	dhcp or dns edit mode in which the data is live on the protocol server. <i>See also</i> staged edit mode .
T	
TAC	Cisco Technical Assistance Center. Cisco Prime Network Registrar provide a cnr_tactool utility to use in reporting issues to the TAC.
TCP/IP	Suite of data communication protocols. Its name comes from two of the more important protocols in the suite: the Transmission Control Protocol (TCP) and the Internet Protocol (IP). It forms the basis of Internet traffic.
template	DNS zones and DHCP scopes can have templates to create multiple objects with similar properties.
transaction signature (TSIG)	DHCP mechanism that ensures that DNS messages come from a trusted source and are not tampered with. <i>See also</i> access control list (ACL) .
trap	Criteria set to detect certain SNMP events, such as to determine free addresses on the network. <i>See also</i> SNMP notification .

trimming and compacting	Trimming is periodic elimination of old historical data to regulate the size of log and other files. Compacting is reducing data older than a certain age to subsets of the records.
Trivial File Transfer Protocol (TFTP)	Protocol used to transfer files across the network using UDP. <i>See also</i> User Datagram Protocol (UDP) .
U	
Universal Time (UT)	International standard time reference that was formerly called Greenwich Mean Time (GMT), also called Universal Coordinated Time (UCT).
update configuration, DNS	Defines the relationship of a zone with its main and backup DNS servers for DNS update purposes.
update map, DNS	Defines an update relationship between a DHCP policy and a list of DNS zones.
update policy, DNS	Provide a mechanism in DHCP for managing update authorization at the DNS RR level.
User Datagram Protocol (UDP)	Connectionless TCP/IP transport layer protocol.
V	
virtual channel identifier (VCI) and virtual path identifier (VPI)	16-bit field in the header of an ATM cell. The VCI, together with the VPI, identifies the next destination of a cell as it passes through a series of ATM switches on its way to its destination. ATM switches use the VPI/VCI fields to identify the next network VCL that a cell needs to transit on its way to its final destination. The function of the VCI is similar to that of the DLCI in Frame Relay.
virtual private network (VPN)	Protocol over which IP traffic of private address space can travel securely over a public TCP/IP network. A VPN uses tunneling to encrypt all information at the IP level. <i>See also</i> VRF .
VRF	VPN Routing and Forwarding instance. Routing table and forwarding information base table, populated by routing protocol contexts. <i>See also</i> virtual private network (VPN) .
W	
well-known port	Any set of IP protocol port numbers preassigned for specific uses by transport level protocols, for example, TCP and UDP. Each server listens at a well-known port so clients can locate it.
WKS record	DNS Well Known Service resource record (RR). Used to list the services provided by the hosts in a zone. Common protocols are TCP and UDP.
Y	

yiaddr	"Your" client IP address, or address that the DHCP server offers (and ultimately assigns) the client. Sent in an RFC 2131 packet between the client and server.
Z	
zone	Delegation point in the DNS tree hierarchy that contains all the names from a certain point downward, except for those names that were delegated to other zones. A zone defines the contents of a contiguous section of the domain space, usually bounded by administrative boundaries. Each zone has configuration data composed of entries called resource records. A zone can map exactly to a single domain, but can also include only part of a domain, with the remainder delegated to another subzone.
zone distribution	Configuration that simplifies creating multiple zones that share the same secondary zone attributes. The zone distribution requires adding one or more predefined secondary servers.
zone of authority	Group of DNS domains for which a given name server is an authority.
zone transfer	Action that occurs when a secondary DNS server starts up and updates itself from the primary server. A secondary DNS server queries a primary name server with a specific packet type called AXFR (transfer all) or IXFR (incrementally transfer) and initiates a transfer of a copy of the database.

