

# **NetCache™ 5.4**

## **Guide to Global Request Manager**

Network Appliance, Inc.  
495 East Java Drive  
Sunnyvale, CA 94089 USA  
Telephone: +1 (408) 822-6000  
Fax: +1 (408) 822-4501  
Support telephone: +1 (888) 4-NETAPP  
Information e-mail: [doccomments@netapp.com](mailto:doccomments@netapp.com)  
Information Web: <http://www.netapp.com>

Part number 210-00374  
December 2002

# Copyright and trademark information

---

## Copyright information

Copyright © 1994–2002 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 1998–2001 The OpenSSL Project. All rights reserved.

No part of this book covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Portions of this product are derived from the Berkeley Net2 release and the 4.4-Lite-2 release, which are copyrighted and publicly distributed by The Regents of the University of California.

Copyright © 1980–1995 The Regents of the University of California. All rights reserved.

Portions of this product are derived from NetBSD, which is copyrighted by Carnegie Mellon University.

Copyright © 1994, 1995 Carnegie Mellon University. All rights reserved. Author Chris G. Demetriou.

Permission to use, copy, modify, and distribute this software and its documentation is hereby granted, provided that both the copyright notice and its permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation.

CARNEGIE MELLON ALLOWS FREE USE OF THIS SOFTWARE IN ITS “AS IS” CONDITION. CARNEGIE MELLON DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Software derived from copyrighted material of The Regents of the University of California and Carnegie Mellon University is subject to the following license and disclaimer:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notices, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS

INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions of the software were created by Netscape Communications Corp.

The contents of those portions are subject to the Netscape Public License Version 1.0 (the “License”); you may not use those portions except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/NPL/>.

Software distributed under the License is distributed on an “AS IS” basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is Mozilla Communicator client code, released March 31, 1998.

The Initial Developer of the Original Code is Netscape Communications Corp. Portions created by Netscape are Copyright © 1998 Netscape Communications Corp. All rights reserved.

Software derived from copyrighted material of Network Appliance, Inc. is subject to the following license and disclaimer:

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use and purchase of this product do not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## **Trademark information**

NetApp and the Network Appliance design are registered trademarks of Network Appliance, Inc. in the United States, Canada, and the European Union. Network Appliance is a registered trademark of Network Appliance, Inc. in Monaco and a trademark of Network Appliance, Inc. in the United States and Canada. FAServer is a registered trademark of Network Appliance, Inc. in the United States and the European Union. NetCache is a registered trademark of Network Appliance, Inc. in the European Union and Japan, and a trademark of Network Appliance, Inc. in the United States. SnapCopy is a registered trademark of Network Appliance, Inc. in the European Union and a trademark of Network Appliance, Inc. in the United States. SnapManager is a registered trademark of Network Appliance, Inc. in the United States. WAFL is a registered trademark of Network Appliance, Inc. in the United States, the European Union, and Canada. FilerView and SecureShare are registered trademarks of Network Appliance, Inc. in the United States. Data ONTAP is a trademark of Network Appliance, Inc. in the United States and Canada. Snapshot is a trademark of Network Appliance, Inc. in the United States and the European Union. NetApp—the Network Appliance Company is a registered trademark of Network Appliance, Inc. in the United States and other countries. ApplianceWatch, BareMetal, Center-to-Edge, ContentDirector, DataFabric, MultiStore, NearStore, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapDrive, SnapMirror, SnapRestore, SnapVault, Vfiler, and Web Filer are trademarks of Network Appliance, Inc. in the United States.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries.

Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries.

RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance NetCache is certified RealSystem compatible.

# Table of Contents

---

|                  |                                                                                |     |
|------------------|--------------------------------------------------------------------------------|-----|
|                  | <b>Preface</b> . . . . .                                                       | vii |
| <b>Chapter 1</b> | <b>Understanding Global Request Manager</b> . . . . .                          | 1   |
|                  | <b>Section A: Overview of Global Request Manager.</b> . . . . .                | 2   |
|                  | <b>Section B: Components of a GRM deployment</b> . . . . .                     | 4   |
|                  | Overview of GRM servers . . . . .                                              | 5   |
|                  | About a GRM DNS server . . . . .                                               | 6   |
|                  | About a GRM L7 server—Application-level (L7) redirection . . . . .             | 7   |
|                  | About GRM agent caches and the groups to which they belong . . . . .           | 10  |
|                  | <b>Section C: How a GRM server determines IP addresses to return</b> . . . . . | 12  |
|                  | Factors affecting IP addresses to be returned. . . . .                         | 13  |
|                  | Status information provided to the GRM server . . . . .                        | 16  |
|                  | GRM domains and their use for IP address filtering . . . . .                   | 18  |
|                  | <b>Section D: Load balancing over center and edge caches</b> . . . . .         | 23  |
|                  | <b>Section E: Failover strategies.</b> . . . . .                               | 26  |
| <b>Chapter 2</b> | <b>Examples of GRM Deployments</b> . . . . .                                   | 29  |
|                  | <b>Section A: Deployment with only a GRM DNS server</b> . . . . .              | 30  |
|                  | <b>Section B: Types of deployments with a GRM L7 server.</b> . . . . .         | 32  |
|                  | Requirement for directing client requests to a GRM L7 server . . . . .         | 33  |
|                  | Using a DNS nameserver to locate a GRM L7 server . . . . .                     | 35  |
|                  | Using a GRM DNS server to locate a GRM L7 server. . . . .                      | 38  |
|                  | Using the network infrastructure to locate a GRM L7 server . . . . .           | 40  |
|                  | Using a proxy.pac file with a GRM L7 server . . . . .                          | 42  |
| <b>Chapter 3</b> | <b>Content Handling with GRM.</b> . . . . .                                    | 47  |
|                  | How GRM obtains content . . . . .                                              | 48  |
|                  | Planning for handling cache misses for CDN content . . . . .                   | 49  |

|                  |                                                                               |           |
|------------------|-------------------------------------------------------------------------------|-----------|
|                  | Using a GRM agent cache to access the Internet also . . . . .                 | 51        |
| <b>Chapter 4</b> | <b>DNS Requirements to Support a GRM DNS Server . . . . .</b>                 | <b>53</b> |
|                  | <b>Section A: Planning for interaction with a back-end DNS nameserver . .</b> | <b>54</b> |
|                  | <b>Section B: Planning for supporting queries for CDN content . . . . .</b>   | <b>56</b> |
|                  | Methods for supporting queries for CDN content . . . . .                      | 57        |
|                  | Method 1: GRM provides DNS services for a content provider . . . . .          | 59        |
|                  | Method 2: URL rewriting to direct requests for CDN content . . . . .          | 62        |
| <b>Chapter 5</b> | <b>Planning Summary for GRM. . . . .</b>                                      | <b>65</b> |
|                  | What you need to plan for . . . . .                                           | 66        |
|                  | Determining the number of GRM servers and agent caches needed . . . . .       | 68        |
|                  | Determining the protocols to run on a GRM agent cache . . . . .               | 70        |
|                  | GRM configuration overview . . . . .                                          | 71        |
| <b>Chapter 6</b> | <b>Testing and Troubleshooting the GRM Network . . . . .</b>                  | <b>73</b> |
|                  | Overview of GRM testing . . . . .                                             | 74        |
|                  | Verifying the GRM network . . . . .                                           | 76        |
|                  | Verifying GRM client request resolution. . . . .                              | 82        |
|                  | Troubleshooting the GRM configuration. . . . .                                | 83        |
|                  | <b>Glossary . . . . .</b>                                                     | <b>89</b> |
|                  | <b>Index . . . . .</b>                                                        | <b>95</b> |

# Preface

---

**Purpose of this guide**

This guide describes the NetCache™ 5.4 Global Request Manager (GRM) feature. It provides information for the following supported NetCache platforms: C1100, C1105, C1200, C2100, C3100, and C6100. Information in this guide will help you determine the best deployment strategy for your organization. The guide will also help you plan the configuration of the NetCache appliances that are part of your GRM deployment and the DNS nameservers that support GRM.

**Who this guide is for**

This guide is for managers and system administrators who are deploying GRM in enterprise or telco environments.

**Obtaining additional deployment assistance**

Network Appliance™ sales engineers provide assistance by helping you to evaluate your requirements, determine how many and what type of NetCache appliances you need, and determine how to deploy your NetCache appliances.

You can also find information related to NetCache on the NOW Web site (<http://now.netapp.com/>).



## About this chapter

This chapter provides information to help you understand the purpose of Global Request Manager (GRM), the components of a GRM deployment, and the interaction between them. This chapter also describes the process that GRM uses to determine the NetCache appliances that can serve content to specific clients.

## Topics in this chapter

This chapter contains information about the following topics:

- ◆ Section A, “[Overview of Global Request Manager](#),” on page 2
- ◆ Section B, “[Components of a GRM deployment](#),” on page 4
- ◆ Section C, “[How a GRM server determines IP addresses to return](#),” on page 12
- ◆ Section D, “[Load balancing over center and edge caches](#),” on page 23
- ◆ Section E, “[Failover strategies](#),” on page 26

## Section A: Overview of Global Request Manager

---

### What is the Global Request Manager?

GRM is the NetCache appliance request redirection feature used to redirect client requests to the cache, or caches, closest to them. Whether a CDN exists within an enterprise network or a telco, it requires a number of components specifically geared toward providing premium service to users—including routing, caching, content management, and network management and reporting services.

The NetCache GRM feature works seamlessly with other products in the Network Appliance™ content delivery suite to build robust CDNs. See “[How GRM obtains content](#)” on page 48 for a description of how the Network Appliance ContentDirector™ application can be used in a GRM deployment for pushing content to NetCache appliances and how ContentReporter can be used to track that information.

### Advantages of directing requests to caches close to clients

A request routing scheme is necessary so that clients can find the content. In traditional caching deployments, L4 switches or WCCP routers are deployed on the network, or client browsers are manually configured to direct requests to NetCache appliances. These solutions are not ideal for CDNs, however, for the following reasons:

- ◆ CDNs want to provide premium service.  
In order to provide premium service, CDNs want to serve content as close to clients as possible. The reason is that the closer the device that serves the content is to the client, the faster the delivery. Additionally, the quality of the content served is likely to be higher. The content is less vulnerable to network problems because it travels a shorter distance.
- ◆ Enterprises and telcos typically do not want to reconfigure their networks to deploy CDN service.

GRM provides an alternative to using L4 switches, WCCP routers, or manual configuration of client browsers to direct content requests to NetCache appliances—an alternative that is especially suited to the needs of CDNs. Often, NetCache appliances configured as GRM agent caches can be “dropped” into the customer environment; typically, reconfiguration of a customer’s network environment is not necessary to support GRM. Reconfiguration might be necessary if your GRM deployment is not typical, for example, if you configure GRM agent caches to fetch content from random origin servers on the Internet, as well as from the CDN data centers (not the recommended approach).

**Importance of  
serving content  
close to users**

With the growing demand for streaming media in particular, serving content stored close to users is especially important because network congestion significantly degrades the quality of streaming media. Additionally, streaming media files, which are very large, consume a great amount of network bandwidth. NetCache appliances deployed close to users, working on behalf of the CDN content servers, can store and serve CDN content to clients. A significant amount of bandwidth is saved because the cache can fulfill multiple requests for the same content instead of the content servers at the other end of the WAN.

## Section B: Components of a GRM deployment

---

### Available GRM components

A GRM deployment always includes at least one *GRM server* (typically more than one) and one or more *GRM agent caches*. NetCache appliances configured as GRM servers provide redirection services, redirecting client requests to NetCache appliances configured as GRM agent caches. GRM agent caches, which are located close to clients, cache content and resolve content requests on behalf of clients.

Individual requirements of an organization determine the deployment configuration and the number of GRM servers and GRM agent caches required.

### Topics in this section

This section contains the following topics:

- ◆ [“Overview of GRM servers”](#) on page 5
- ◆ [“About a GRM DNS server”](#) on page 6
- ◆ [“About a GRM L7 server—Application-level \(L7\) redirection”](#) on page 7
- ◆ [“About GRM agent caches and the groups to which they belong”](#) on page 10

## Overview of GRM servers

---

### Redirection methods available with a GRM server

GRM supports two different redirection methods for GRM servers:

- ◆ DNS-based redirection services (GRM DNS server)
- ◆ Application-level (L7) redirection services (GRM L7 server)

When you configure a NetCache appliance as a GRM server, you configure general GRM server options plus specific options for the type of redirection that you need. Depending on your redirection requirements, you might deploy all GRM DNS servers, all GRM L7 servers, or a combination of the two types of GRM servers (*a mixed hierarchy*). Typically, multiple GRM servers are deployed for redundancy.

---

#### Note

In this guide the term “GRM server” is used when the information applies to both GRM DNS servers and GRM L7 servers. The type of GRM server is specified if the information applies to only one type of GRM server.

---

### General GRM server functions

GRM server functions include the following:

- ◆ Maintains authority and responsibility for a set of configured domains called GRM domains. You identify these domains when you configure the server and, if desired, set filtering policies on location and protocols handled.
- ◆ Communicates with GRM agent caches, obtaining proximity and status information from the agents.
- ◆ Determines which GRM agent cache IP addresses to return in response to a DNS query (if the server is a GRM DNS server) or a client request (if the server is a GRM L7 server).

## About a GRM DNS server

---

### How a GRM DNS server performs redirection

A GRM DNS server provides standard DNS redirection plus a layer of intelligence that is not available with standard DNS, including support for proximity measurements, detection of unavailable GRM agent caches, and filtering. As with standard DNS, the client's local DNS nameserver communicates with the GRM DNS server through a DNS request, and the GRM DNS server responds with a DNS response.

The GRM DNS server returns to the client a list of IP addresses of the available GRM agent caches that are *close to the local DNS nameserver* that the client uses. The client browser then selects a GRM agent cache from the list returned by the GRM DNS server. The GRM agent cache that the browser selects might or might not be the closest cache to the client.

See Section A, “[Deployment with only a GRM DNS server](#),” on page 30 and “[Using a GRM DNS server to locate a GRM L7 server](#)” on page 38 for deployment examples.

### Relationship to the back-end nameserver

A GRM DNS server can assume the identity of the existing authoritative DNS nameserver for the CDN name space; that is, the GRM server can assume the IP address that was previously assigned to the authoritative DNS nameserver. The authoritative DNS nameserver is then reconfigured with a new IP address. The authoritative DNS nameserver becomes the back end to the GRM DNS server in this case.

### When you need to include a GRM DNS server

Deploying a GRM DNS server is appropriate in the following circumstances:

- ◆ When local DNS nameservers are close to clients.  
See Section A, “[Deployment with only a GRM DNS server](#),” on page 30 for a deployment example.
- ◆ When you are deploying GRM L7 servers and you want to use a GRM DNS server to direct client requests to the GRM L7 servers.  
See “[Using a GRM DNS server to locate a GRM L7 server](#)” on page 38 for a deployment example.

## About a GRM L7 server—Application-level (L7) redirection

---

### How a GRM L7 server performs redirection

Like the GRM DNS server, the GRM L7 server includes intelligence to support proximity measurements, detection of unavailable GRM agent caches, and filtering. However, the GRM L7 server can provide more precise redirection than the GRM DNS server. The reason is that a GRM L7 server uses Application-level redirection (*L7 redirection*). L7 redirection is based on the URL and not on the header content. The GRM L7 server rewrites the request *based on the client's IP address*, and sends a 302 location changed response to the client. The response includes the IP address of the closest eligible GRM agent cache or caches.

In contrast, the GRM DNS server returns a list of IP addresses of eligible GRM agent caches, and the client browser selects the cache to send the request to.

**Note**

The method of rewriting varies for HTTP and streaming requests. See “[GRM L7 server request rewriting for HTTP and streaming content](#)” on page 8 for details.

---

### When you need to include a GRM L7 server

The following table provides general guidelines for determining when you need to include a GRM L7 server in your deployment.

| If...                                                                                    | Then...                                                                                                                                                                                                                                                       |
|------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The local DNS nameserver and the clients are physically close.                           | You can use all GRM DNS servers (a <i>pure</i> GRM DNS deployment).                                                                                                                                                                                           |
| If the local DNS nameserver and the clients it services are <i>not</i> physically close. | <p>You will want to include one or more GRM L7 servers in your deployment.</p> <p>You must provide a means for client requests to be directed to the GRM L7 servers. See “<a href="#">Requirement to direct requests to the GRM L7 server</a>” on page 8.</p> |

For example, some enterprise CDNs are spread over continents, with clients located on one continent and their local DNS nameserver located on a different continent. If only GRM DNS servers are used in such cases, the GRM agent

cache IP addresses that GRM DNS servers return would be on the same continent as the local DNS nameserver. The reason is that the GRM DNS server sees only the IP address of the local DNS nameserver, not the IP address of the client.

By deploying a GRM L7 server for redirection service to clients, the GRM L7 server will redirect requests based on client IP address. Therefore, the GRM agent cache IP address returned to the client will be truly close to the requesting client.

### **Requirement to direct requests to the GRM L7 server**

When you deploy a GRM L7 server, your deployment must include a way for client requests to locate the GRM L7 server. You can use one of the following methods:

- ◆ A standard DNS nameserver
- ◆ A GRM DNS server
- ◆ A locator solution in your network infrastructure
- ◆ A proxy.pac file

You choose the solution that best meets your requirements. For example, if your network infrastructure already includes a locator device, you can configure that device to send client requests to the GRM L7 server closest to the client.

Section B, “[Types of deployments with a GRM L7 server](#),” on page 32 describes the advantages and disadvantages of each method and provides deployment examples.

---

#### **Note**

If a proxy.pac file is used to direct client requests to a GRM L7 server, the GRM L7 server might return the IP address of more than one GRM agent cache close to the client. If another method is used, the GRM L7 server rewrites the client request with the IP address of only one GRM agent cache.

---

### **GRM L7 server request rewriting for HTTP and streaming content**

The L7 request rewriting process and how you can influence it differ for HTTP and streaming protocols.

**Rewriting of requests for HTTP content:** For HTTP, a GRM L7 server sends a 302 response to redirect a client request and rewrites the URL as

`http://redirected-agent-ip/original_url`

Normally, a GRM L7 server returns the closest GRM agent cache in its 302 response to the client. However, for content that does not require a large amount of bandwidth (for example, HTML files and GIF files), it is probably more efficient for the GRM L7 server to serve the content to the client directly rather than redirect the client to a GRM agent cache close to the client.

However, it is more efficient for a GRM L7 server to redirect requests for content that requires a significant amount of bandwidth (for example, PDF, PPT, DOC, and XLS files) to GRM agent caches close to clients. You use the HTTP Redirection Suffix option on the GRM Server - General page in the Appliance Manager to control the HTTP file extensions for which the GRM L7 server is to redirect content. By default, requests for content with .PDF, .PPT., and .XLS extensions are redirected to GRM agent caches close to the requesting clients. HTTP content that does not match the listed extensions to be redirected is served by the GRM L7 server.

**Note**—  
GRM prepends the GRM agent cache IP address to the regular URL before sending the 302 response to the client. If a user bookmarks the URL for a document served by the GRM agent cache (for example, a URL for a .PDF file) and the cache becomes unavailable, the bookmark will no longer function. This occurs for HTTP content only.

**Rewriting of requests for streaming content:** The GRM L7 server redirects all streaming traffic.

| Protocol | How rewriting occurs                                                                                                                                                                                                                                                                                                                    |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RTSP     | A GRM L7 server sends a 302 response to redirect a client request, and rewrites the URL as<br><br>rtsp://redirected-agent-ip/original_url                                                                                                                                                                                               |
| MMS      | A GRM L7 server rewrites the ASX file with the IP address of the closest GRM agent cache, in the following form<br><br>Original request:<br>http://grm-domain/path/file.asx<br><br>Contents of the ASX file:<br>mms://server/path/file.wmv<br><br>ASX rewrites the content request as<br>mms://redirected-agent-ip/server/path/file.wmv |

## About GRM agent caches and the groups to which they belong

---

### Function of GRM agent caches

Every GRM deployment includes a set of NetCache appliances that are configured as GRM agent caches. GRM agent caches cache content and resolve content requests on behalf of clients.

When you configure a GRM server, you identify each GRM agent cache and the group with which it is associated. A GRM agent cache is categorized as follows:

- ◆ An *edge cache* in an *edge group*
- ◆ A *center cache* in a *center group*

### About GRM edge caches and edge groups

Edge caches are located in groups that are referred to as edge groups. Edge caches in a particular edge group serve content only to the clients associated with that edge group. For example, in a telco CDN, each POP is typically an edge group and each GRM agent cache deployed at the POP is an edge cache. An edge cache in POP A can serve content only to clients in POP A.

Characteristics of edge caches are as follows:

- ◆ They perform health monitoring and inform the GRM server that they are active. (The GRM server considers health information when it determines which caches are eligible to return content to a client.)
- ◆ They do not perform proximity measurements. When you configure your GRM server, you specify the CIDR blocks associated with each edge group. The CIDR block subnets associated with an edge group define which clients can be served by caches within the edge group. See Section C, “[How a GRM server determines IP addresses to return](#),” on page 12.
- ◆ They fail over to center caches if no edge caches in the edge group are available. For details, see “[Failover for GRM servers](#)” on page 26.

### About GRM center caches and center groups

Center caches are located in *center groups*. A center group is equivalent to a telco or enterprise CDN data center.

Characteristics of center caches are as follows:

- ◆ They provide content for clients that are not associated with an edge group. For example, if a CDN provides services for POPs and clients not in POPs, the center caches would serve content for clients not in the POPs.

## Section B: Components of a GRM deployment

- ◆ They perform health monitoring, informing the GRM server that they are active.
- ◆ They measure proximity between themselves and the requesting client (for a GRM L7 server) or the local DNS nameserver (for a GRM DNS server) and return that information to the GRM server.

## Section C: How a GRM server determines IP addresses to return

---

### About this section

This section describes the factors that a GRM server considers when determining which GRM agent cache IP addresses to return. It also describes the GRM domain filtering policies you can configure that affect the IP addresses that the GRM server will return.

### Topics in this section


This section contains the following topics:

- ◆ [“Factors affecting IP addresses to be returned”](#) on page 13
- ◆ [“Status information provided to the GRM server”](#) on page 16
- ◆ [“GRM domains and their use for IP address filtering”](#) on page 18

## Factors affecting IP addresses to be returned

### Information considered when determining IP addresses to return

A GRM server evaluates many factors when determining how to redirect a client request. The following illustration shows what the GRM server considers to determine the IP addresses to return when it receives a DNS query (if the server is a GRM DNS server) or a client request (if the server is a GRM L7 server).



Details about the preceding flow chart are discussed in the following paragraphs.

**Is the request or query from an edge group?** The GRM server first checks the list of CIDR blocks in its configuration to determine whether the client request or DNS query is associated with an edge group.

- ◆ If the GRM server can make a match, it next considers any policies that have been set and other factors that affect eligibility.
- ◆ If the GRM server cannot make a match, it tries to find previous proximity information to enable it to determine the GRM agent caches close to the client.

**Discarding caches not meeting policies and other criteria:** Before a GRM server returns GRM agent cache IP addresses to a client, it eliminates any GRM agent caches that do not meet policies or other controls that have been set. The GRM server evaluates a number of factors to determine whether a cache meets the criteria to be returned to a client, as follows:

- ◆ The GRM server applies optional controls that have been set, for example, the threshold for the CPU activity of a center or edge cache and availability. For more information, see “[Status information provided to the GRM server](#)” on page 16.
- ◆ The GRM server applies any GRM domain filtering criteria that might have been specified for the domain, and it returns IP addresses of eligible center and edge caches to the client.

If the GRM server is not directing content requests for the host or domain, then

- ❖ A GRM DNS server passes the DNS query to the authoritative DNS nameserver for which it is the front end.
- ❖ A GRM L7 server serves the client content request itself;.

See “[GRM domains and their use for IP address filtering](#)” on page 18 for details about domains.

---

**Note**

A GRM DNS server returns the IP addresses of *all* eligible edge and center caches. If a proxy.pac file is used to direct client requests to a GRM L7 server, the GRM L7 server might return the IP address of more than one GRM agent cache close to the client. If another method is used, the GRM L7 server rewrites the client request with the IP address of only one GRM agent cache.

---

**If the request is not from an edge group and no proximity measurement exists:** The process a GRM server follows differs, depending on whether it is a GRM DNS server or GRM L7 server.

- ◆ GRM DNS server

If the GRM DNS server cannot measure proximity, or if it has not determined proximity to a local DNS nameserver previously, the GRM server uses the IANA database to narrow down the list of eligible caches to be returned by matching the client's IP address to a continent.

The GRM server does not have proximity information in its database in the following circumstances:

- ❖ The local DNS nameserver has not previously queried the GRM server, so no proximity measurements were made.
  - ❖ Caches reported that a local DNS nameserver was inaccessible, so the GRM server discarded the proximity data for it.
  - ❖ A local DNS nameserver has not contacted the GRM server recently (within a maximum of 30 days), so the GRM server discarded the DNS entries for the local DNS nameserver.
- ◆ GRM L7 server

When the GRM L7 server receives the first request from a new client, it has no proximity information about that client.

- ❖ If the Force Redirection option is enabled (the default), the GRM L7 server will force redirection of that request to the most appropriate agent, based on the responses the server receives from its agents to its request for proximity measurements. If the GRM L7 server does not receive a response for redirection information from its GRM agent caches within 30 seconds, the GRM L7 server serves the content itself.
- ❖ If the Force Redirection option is disabled, the GRM L7 server automatically serves the content itself. During the process, the GRM L7 server collects the proximity information that it can use the next time it receives a request from the client.

The GRM L7 server serves content to a client directly in the following additional circumstances:

- ❖ No edge cache or center cache that meets the domain filtering criteria is available.
- ❖ The GRM L7 server receives a request for a type of HTTP content that is not listed in the HTTP Redirection Suffix option.

## Status information provided to the GRM server

---

**Type of status information sent by GRM agent caches** GRM agent caches continuously send GRM servers the information shown in the following table. A GRM server considers this information along with other factors in determining the GRM agent cache IP addresses to return to a client.

| Information               | Comments                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| System health             | Center and edge caches automatically report their status every 30 seconds. A GRM server considers a center or edge cache down if it has not received a report from the cache for 60 seconds. A GRM server does not return the IP addresses of any edge and center caches that are down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| CPU load                  | When you configure a GRM server, you set the maximum CPU activity level for each center and edge cache. Each center and edge cache reports its CPU threshold every 30 seconds. The GRM server does not return the IP addresses of center and edge caches that report a CPU load greater than the configured threshold.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Client-to-cache proximity | <p><b>Edge caches:</b> When you configure a GRM server, you enter CIDR blocks associated with each edge group. This configuration enables the GRM server to determine the set of edge caches to associate with a client request (for a GRM L7 server) or with a local DNS query (for a GRM DNS server). No active proximity measurements are needed.</p> <p>For clients associated with an edge group, the GRM server returns only the IP addresses of edge caches in that edge group.</p> <p><b>Center caches:</b> The GRM server checks its proximity tables or sends a probe query for proximity measurements to center caches in order to determine the distance between a center cache and a client that is not associated with an edge group. A GRM DNS server measures the distance between the local DNS nameserver that sent the query and the GRM DNS server. A GRM L7 server measures the distance between the client and its GRM agent caches.</p> <p>Measurement intervals increase geometrically (30 seconds, 1 minute, 2 minutes, and so on).</p> |

### **Mechanisms for measuring proximity**

The GRM server uses a number of mechanisms to determine proximity, changing from one to another if one does not produce results. The GRM server uses ICMP Echo request and response, DNS query, and traceroute. A GRM server might not be able to measure proximity if traceroute provides inconsistent results, which could result from blocking by firewalls on a WAN. In addition, for security reasons, firewalls and routers sometimes block the GRM proximity measurement traffic.

### **Additional information about proximity for a GRM DNS server**

The assumption is that the round-trip time between the GRM DNS server and the local DNS nameserver that sent the query is similar to the actual round-trip time that the client will experience. In the majority of cases, this assumption is valid. However, an alternative deployment scenario might be required in the following circumstances because the GRM proximity measurements might be inaccurate:

- ◆ If the local DNS nameserver is configured for recursive queries (that is, the local DNS nameserver does not take a direct path to the authoritative GRM domain)
- ◆ If the local DNS nameserver is located a great distance from the client

---

**Note**

If the local DNS nameserver is located a great distance from the client, Network Appliance recommends that you deploy a GRM L7 server.

---

### **Communication protocol used for status updates**

The GRM server and the center and edge caches communicate with each other over TCP. Enabling the GRM server or GRM agent caches automatically enables NetCache to listen on TCP port 89 for communication between the GRM server and GRM agent caches.

### **Viewing the IP addresses that were returned**

The GRM server access log shows which center and edge caches the GRM server returned. See the *Guide to Messages and Log File Entries* and the Appliance Manager online Help for more information about the GRM server access log.

## GRM domains and their use for IP address filtering

---

### About this section

This section describes how the GRM server uses GRM domain information. It also provides information to help you decide whether specifying optional filtering criteria would be beneficial for your organization.

---

#### Note

Network Appliance recommends that you do not use domain filtering unless you must handle circumstances such as those listed in [“When you might want to set domain filters”](#) on page 19.

---

### GRM domains defined

One step in configuring a GRM server is to identify GRM domains. GRM domains are fully qualified domain names (FQDNs) or DNS domains for which the GRM server is directing content requests. For a GRM DNS server, these FQDNs and domains are a subset of the authoritative domains maintained by the DNS server that can rewrite the host name. For a GRM L7 server, these FQDNs and domains are the CDN domains for which you want the server to redirect requests.

For each GRM domain that you identify, you can include optional filtering criteria to constrain the IP addresses that the GRM server will return to a client. The GRM server considers the GRM domain filtering criteria when determining which center and edge caches are eligible to handle the content request.

### Multiple CDNs and GRM domains

One GRM server can control multiple CDNs if that server is authoritative for all the domains associated with the CDNs.

### About domain policies

Optionally, you can set the following types of policies associated with GRM domains to constrain the IP addresses that the GRM server returns:

- ◆ GRM domain filters based on protocols, location, and group  
Domain filtering is needed only if it is necessary to limit the content served by a particular group or from a particular location. See [“When you might want to set domain filters”](#) on page 19.

- ◆ Host-level policies

By specifying FQDNs for specific hosts within domains that the GRM server handles, you can set up the GRM server to evaluate policies that are set at the host level before evaluating the policies that are set at the GRM domain level.

## **Domain filters you can set**

Setting domain filters is optional. You can set GRM domain filters for the following:

- ◆ Protocols that are running on caches

You can set up protocol filtering for your domains for one or more of the following protocols:

- ❖ HTTP

- ❖ RTSP

- ❖ MMS

- ◆ Location of center and edge caches

You can use the following location filters, by themselves or in combination, to set up filtering on location:

- ❖ Country

- ❖ Region

- ❖ Continent or island

- ◆ Group

Protocol, location, and group filters work together, as described in the following examples:

- ◆ If you specify protocols as part of your domain filtering, the GRM server returns only the caches whose protocol services exactly match the domain filter. For example, if an edge cache handles only HTTP but your domain filter is set for HTTP and RTSP, the GRM server does not return the IP address of that edge cache.

- ◆ If you set both protocol and location filters, the GRM server returns IP addresses of only the caches that match both the protocol and location filters that are set for the GRM domain.

## **When you might want to set domain filters**

You might want to set domain filtering criteria in the circumstances discussed in the following paragraphs.

**Binding a GRM domain to specific protocols:** Assume that an edge group in the domain includes some edge caches that provide service only for streaming protocols (RTSP and MMS) and other edge caches that provide service only for HTTP.

If no protocol filtering is set then

- ◆ A GRM DNS server would return the IP addresses of all edge caches in the domain that are associated with the edge group. The problem is that even if the client requests only streaming content, the client browser might select, through round robin, the IP address of an edge cache configured only for Web content.
- ◆ A GRM L7 server might return the IP address of an edge cache that can handle only Web content even if the client request is only for streaming content.

Creating one GRM domain that filters on streaming protocols and another GRM domain that filters on Web protocols would ensure that only the IP addresses of caches that can handle the requested protocol would be returned.

**Filtering according to location:** If you need to restrict the locations to which content requests can be sent, for example, because of service agreements or political issues, you could create a GRM domain with filtering criteria on location. Then, when the GRM DNS server receives a DNS query for that domain, it returns only the IP addresses of caches that match the filtering criteria that you set up. Similarly, when the GRM L7 server receives a client request for the domain, it returns only IP addresses of caches that match the filtering criteria you set up.

**Filtering according to group:** When you set your domain filter, if you need to restrict the data centers or edge groups to which content requests can be sent, you can identify the center groups, edge groups, or both to which the requests can be sent.

**GRM domain  
filtering examples**

The following table shows examples of filters set in the GRM domain configuration and whether the filters match characteristics identified for a particular center or edge cache.

| <b>Filter on...</b>                      | <b>Domain settings (only filters are specified)</b>                                                                        | <b>Center group or edge group settings (all must be specified)</b>                                                                                                                           | <b>Match?</b> |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| Region                                   | Country: None specified<br>Region: West<br>Continent: None specified<br>Group: None specified<br>Protocols: None specified | <b>Settings for cache Jupiter</b><br>Country: United States<br>Region: West<br>Continent: North America<br>Group: Data Center 1 U. S.<br>Protocols handled by this cache:<br>HTTP, RTSP, MMS | Yes           |
| Country, region, and continent           | Country: France<br>Region: West<br>Continent: Europe<br>Group: None specified<br>Protocols: None specified                 | <b>Settings for cache Bordeaux</b><br>Country: France<br>Region: East<br>Continent: Europe<br>Group: Data Center France<br>Protocols handled by this cache:<br>HTTP, RTSP, MMS               | No            |
| Protocol, country, region, and continent | Country: France<br>Region: West<br>Continent: Europe<br>Group: None specified<br>Protocols: MMS, HTTP                      | <b>Settings for cache Riesling</b><br>Country: France<br>Region: West<br>Continent: Europe<br>Group: Edge Group 1 France<br>Protocols handled by this cache:<br>HTTP                         | No            |

| <b>Filter on...</b>             | <b>Domain settings (only filters are specified)</b>                                                                                      | <b>Center group or edge group settings (all must be specified)</b>                                                                                                               | <b>Match?</b> |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| Protocol, region, and continent | Country: None specified<br>Region: West<br>Continent: Europe<br>Group: None specified<br>Protocols: MMS, HTTP                            | <b>Settings for cache Port</b><br>Country: France<br>Region: West<br>Continent: Europe<br>Group: Edge Group 2 France<br>Protocols handled by this cache: MMS, HTTP               | Yes           |
| Group                           | Country: None specified<br>Region: None specified<br>Continent: None specified<br>Group: Data Center 2 U.S.<br>Protocols: None specified | <b>Settings for cache Mars</b><br>Country: United States<br>Region: East<br>Continent: North America<br>Group: Data Center 2 U. S.<br>Protocols handled by this cache: MMS, HTTP | Yes           |

**Note**


---

For center and edge caches, NetCache automatically determines the protocols that the cache handles. You do not have to enter protocol information when you define your center and edge caches in the GRM server configuration.

---

## Section D: Load balancing over center and edge caches

---

### Rotation of IP addresses for load balancing—GRM L7 server

When the GRM L7 server rewrites a client request, it rotates the IP addresses of qualifying GRM agent caches it returns among multiple client requests so that no one cache is returned consecutively, even to different clients. In this way, GRM L7 performs load balancing in rewriting requests. “[Mechanisms for load balancing—GRM L7 server](#)” on page 23 describes the rotation scheme that the GRM L7 server uses.

The GRM L7 server also provides load balancing through rotation of IP addresses when a proxy.pac file is used to direct client requests to a GRM L7 server and the GRM L7 server returns multiple addresses in an IP address list. Typically, when multiple IP addresses are returned in a proxy.pac file, client browsers select the first IP address in the list. If all clients were sent the same list of IP addresses, in the same order, all content requests would be sent to the same center or edge cache. The GRM L7 solution provides load balancing by rewriting the proxy.pac file in such a way that IP addresses in the lists returned to clients are themselves rotated.

### Mechanisms for load balancing—GRM L7 server

The following table shows an example of the IP address rotation scheme that the GRM L7 server uses to generate the IP addresses returned in response to client requests. This scheme ensures that content requests are distributed over all caches in a group. The rotation order does not apply across groups.

For example, assume that three eligible center caches exist—A, B, and C. Each time the GRM L7 server receives a new client request, it rotates the IP address it returns to clients, as shown in the following table.

| Client request | Order of IP addresses that the GRM server returns |
|----------------|---------------------------------------------------|
| First          | A                                                 |
| Second         | B                                                 |
| Third          | C                                                 |
| Fourth         | A                                                 |

If client requests are directed to the GRM L7 server by a proxy.pac file, multiple IP addresses can be returned. If multiple IP addresses are returned, the rotation scheme is like that shown for the GRM DNS server in [“Rotation of IP addresses for load balancing—GRM DNS server”](#) on page 24.

**Rotation of IP addresses for load balancing—GRM DNS server**

Both the local DNS nameserver and the GRM DNS server include schemes that support load balancing of content requests over caches.

After the GRM DNS server returns a list of IP addresses of available caches to the local DNS nameserver, the local DNS nameserver caches the IP address list returned by the GRM server for one minute—the Time To Live (TTL) set by the GRM DNS server. During the one-minute time period, the local DNS nameserver rotates the list of IP addresses it received until the TTL set by the GRM DNS server expires. After the TTL expires, the local DNS nameserver requeries the GRM DNS server.

The GRM DNS server rotates the IP addresses it returns to local DNS nameservers based on the scheme described in [“Mechanisms for load balancing—GRM DNS server”](#) on page 24. As a result, even if the list of IP addresses that the GRM DNS server returns to different local DNS nameservers includes the same IP addresses, the order of the IP addresses would be different.

The rotation schemes provided by the local DNS nameservers and the GRM DNS server are beneficial because most client browsers have no mechanism for load balancing. Typically, when multiple IP addresses are returned in a DNS response, client browsers select the first IP address in the list of IP addresses. If all clients were sent the same list of IP addresses in the same order, all content requests would be sent to the same center or edge cache.

**Mechanisms for load balancing—GRM DNS server**

The following table shows an example of the IP address rotation scheme that the GRM DNS server uses to generate the IP addresses returned in response to DNS queries. This scheme ensures that content requests are distributed over all caches. For example, assume that three center caches exist—A, B, and C. Each time the GRM DNS server receives a new DNS query, it rotates the list of IP addresses it returns to clients, as shown in the following table. The rotation order does not apply across groups.

| DNS lookup | Order of IP addresses in the list the GRM server returns |
|------------|----------------------------------------------------------|
| First      | A, B, C                                                  |
| Second     | B, C, A                                                  |
| Third      | C, A, B                                                  |
| Fourth     | A, B, C (The pattern is repeated, first through third.)  |

In most cases, the GRM DNS server mechanism for load balancing is sufficient to balance the load of content requests across caches.

## Section E: Failover strategies

---

### **Do you need a failover strategy?**

NetCache appliances are extremely reliable, with a proven track record of 99.9 percent availability. However, if you want to ensure uninterrupted content delivery service, Network Appliance recommends that you have a failover strategy to handle the unlikely occurrence of your GRM servers or center or edge caches failing.

### **Failover for GRM servers**

To avoid interruption of your content delivery service, Network Appliance recommends that you deploy multiple GRM servers of the same type. Then, for example, if one GRM L7 server becomes unavailable, the other GRM L7 server can handle the traffic that would ordinarily be handled by the unavailable server.

All center and edge caches send status updates to all GRM servers listed in the cache's configuration. A GRM server that was temporarily unavailable begins receiving status information from the center and edge caches as soon as that GRM server becomes available again.

### **Failover from edge caches to center caches**

GRM includes automatic failover to center caches if all edge caches in an edge group become unavailable at the same time. If multiple edge caches are deployed for an edge group, a client first fails over to the other edge caches returned by the GRM server as eligible caches. A client fails over to a center cache only if no edge caches are available.

If your CDN services both clients inside edge groups and clients not associated with edge groups, an automatic failover strategy will already be in place. The reason is that you will need to deploy at least one center cache to provide content delivery service for the clients outside of edge groups. That center cache automatically provides failover for the edge groups.

If your CDN does not service clients that are outside of edge groups, you should deploy a center cache specifically for failover if it is important for your CDN to provide uninterrupted service. A center cache can provide failover for multiple edge groups, possibly eliminating the need for deploying an extra edge cache in an edge group so that the traffic load can be handled even if one edge cache becomes unavailable.

**If no center or edge caches are available**

If no center or edge caches are available, the following occurs:

- ◆ If the server is a GRM L7 server  
The GRM L7 server can serve content to clients.
- ◆ If the server is a GRM DNS server  
By default, when no center or edge caches are available to serve the requested content, the GRM DNS server does not return any IP addresses to the local DNS nameserver. However, you can use the Origin Server Failover option in the GRM DNS server configuration to force the GRM server to query the back-end DNS nameserver and return the IP address of the origin server if a center or edge cache cannot be found to serve the requested content.



**About this chapter**

This chapter provides several GRM deployments to illustrate the use of GRM DNS servers and GRM L7 servers. Each deployment describes how request redirection works with that deployment.

**Topics in this chapter**

This chapter contains information about the following sections:


- ◆ Section A, “[Deployment with only a GRM DNS server](#),” on page 30
- ◆ Section B, “[Types of deployments with a GRM L7 server](#),” on page 32

## Section A: Deployment with only a GRM DNS server

### When GRM DNS servers are used

The following illustration shows a deployment in which only a GRM DNS server is used for redirection services. Deploying one or more GRM DNS servers is appropriate in the following circumstances:

- ◆ When clients are physically close to their local DNS nameservers (as in this example), you can deploy only GRM DNS servers to provide redirection services (a *pure GRM server deployment*).
- ◆ You choose to redirect client requests to a GRM L7 server by using a GRM DNS server. See [“Requirement for directing client requests to a GRM L7 server”](#) on page 33 for more information.



The preceding illustration shows the resolution of a DNS query from a client in a POP. Both the GRM DNS server and the back-end nameserver for the GRM DNS server are located in the CDN data center. A back-end nameserver can be in a location different from the GRM DNS server, and the GRM DNS server can point to it, however. For simplicity, the example shows only one GRM DNS server and one back-end DNS nameserver. Typically, multiple GRM DNS servers and multiple back-end DNS nameservers would be deployed for redundancy.

A GRM DNS server can assume the identity of the existing authoritative DNS nameserver for the CDN name space; that is, the GRM DNS server can assume the IP address that was previously assigned to the authoritative DNS nameserver. The authoritative DNS nameserver is then reconfigured with a new IP address. The authoritative DNS nameserver becomes the back-end to the GRM DNS server in this case.

If the GRM DNS server returns multiple IP addresses of GRM agent caches, the client browser selects one GRM agent cache from the list returned and sends the content request to that cache. As the previous illustration shows, if a GRM agent cache does not have the requested content (a cache miss), the cache fetches it from the CDN origin server.

## Section B: Types of deployments with a GRM L7 server

---

### When GRM L7 servers are used

For organizations whose networks span countries or continents, GRM L7 deployments provide a way of redirecting client requests to caches truly close to those clients. This is not always possible when only DNS is used. The reason is that when DNS is used, proximity is measured in relation to the client's *local DNS nameserver*. In some organizations, however, the local DNS nameserver is far from clients it serves.

### Topics in this section

This section describes several types of GRM L7 server deployments and the advantages and disadvantages of the different deployments. This section contains the following topics:

- ◆ [“Requirement for directing client requests to a GRM L7 server”](#) on page 33
- ◆ [“Using a DNS nameserver to locate a GRM L7 server”](#) on page 35
- ◆ [“Using a GRM DNS server to locate a GRM L7 server”](#) on page 38
- ◆ [“Using the network infrastructure to locate a GRM L7 server”](#) on page 40
- ◆ [“Using a proxy.pac file with a GRM L7 server”](#) on page 42

# Requirement for directing client requests to a GRM L7 server

**Providing a means of locating a GRM L7 server**

If you are using a GRM L7 server, you must provide a way for the client requests to reach the GRM L7 server that is closest to the client. The following table shows the methods that are available. This chapter includes a deployment example for each method.

| Methods for directing requests to a GRM L7 server | Advantages and disadvantages                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Use a standard DNS nameserver                     | <p><b>Advantage</b></p> <ul style="list-style-type: none"><li>◆ This is a simple solution.</li></ul> <p><b>Disadvantages</b></p> <ul style="list-style-type: none"><li>◆ There is no way to guarantee that the GRM L7 server to which the DNS nameserver redirects requests is the closest one to the client.</li></ul> <p>For example, if one GRM L7 server is in Australia and another is in the United States, the DNS nameserver round-robin process might direct requests from clients in the United States to the GRM L7 server in Australia.</p> <ul style="list-style-type: none"><li>◆ The GRM L7 server can become overloaded if all client requests are sent to the same GRM L7 server</li></ul> <p>See <a href="#">“Using a DNS nameserver to locate a GRM L7 server”</a> on page 35.</p> |
| Deploy a GRM DNS server                           | <p><b>Advantages</b></p> <ul style="list-style-type: none"><li>◆ The GRM DNS server acts as an intelligent means of locating the closest GRM L7 server to the clients.</li><li>◆ The GRM DNS server can also be used to monitor the health of the GRM L7 servers.</li></ul> <p><b>Disadvantages</b></p> <ul style="list-style-type: none"><li>◆ This solution is more complicated than using only a DNS nameserver.</li><li>◆ This solution requires additional NetCache appliances.</li></ul> <p>See <a href="#">“Using a DNS nameserver to locate a GRM L7 server”</a> on page 35.</p>                                                                                                                                                                                                              |

| <b>Methods for directing requests to a GRM L7 server</b> | <b>Advantages and disadvantages</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Use a locator solution in your network infrastructure    | <p><b>Advantage</b></p> <ul style="list-style-type: none"> <li>◆ An existing locator device in the network infrastructure can be used. Some locator devices provide sophisticated features, for example, features for load balancing over GRM L7 servers.</li> </ul> <p><b>Disadvantages</b></p> <ul style="list-style-type: none"> <li>◆ You must rely on the existing network infrastructure.</li> <li>◆ You must have devices in your network infrastructure to support this function.</li> </ul> <p>See <a href="#">“Using the network infrastructure to locate a GRM L7 server”</a> on page 40.</p> |
| Use a proxy.pac file                                     | <p><b>Advantage</b></p> <ul style="list-style-type: none"> <li>◆ This solution is inexpensive, scalable, and easy to implement.</li> </ul> <p><b>Disadvantages</b></p> <ul style="list-style-type: none"> <li>◆ Client browsers must be manually configured to point to the proxy.pac file on the GRM L7 server that is closest to the clients.</li> <li>◆ Load balancing must be programmed into the proxy.pac file (for example, hash on URL).</li> </ul> <p>See <a href="#">“Using a proxy.pac file with a GRM L7 server”</a> on page 42.</p>                                                         |

## Using a DNS nameserver to locate a GRM L7 server

---

### Using a DNS nameserver to locate the GRM L7 server

A simple GRM L7 deployment is to use a DNS nameserver to direct client requests to one or more GRM L7 servers. This deployment is best for organizations for which a simple configuration is the highest priority. This type of deployment has the following disadvantages:

- ◆ GRM L7 server overload


If all client requests are sent to a single GRM L7 server, the server can easily become overloaded when the load of client requests is high. To handle overloading, you can deploy multiple GRM L7 servers and allow the DNS server to return their IP addresses to clients in a round-robin fashion.

- ◆ Poor response time

There is no way to guarantee that the GRM L7 server to which the DNS nameserver redirects requests is the closest one to that client. (The DNS nameserver uses the round-robin method of load balancing.)

For example, a DNS nameserver could send a request from a client in Australia to a GRM L7 server in the United States rather than to a closer GRM L7 server in Australia. In this case, the GRM L7 server in the United States would find the GRM agent cache closest to the client. However, sending a request to the GRM L7 server in the United States takes longer than sending a request to the GRM L7 server in Australia.

The following illustration shows a deployment with a DNS nameserver directing client requests to the GRM L7 server.



The following table describes the path of the request in the previous illustration.

| Stage | Description                                                                                                                                                               |
|-------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1     | The client queries its local DNS nameserver (LDNS) for the IP address of the origin server of www.netcache.com.                                                           |
| 2     | The LDNS forwards the request to the authoritative DNS nameserver for resolution.                                                                                         |
| 3     | The authoritative DNS server (responsible for the netcache.com domain) resolves the request and returns the IP address of the GRM L7 server (135.105.52.101) to the LDNS. |
| 4     | The LDNS returns the IP address of the GRM L7 server to the client in its response to the client's DNS request.                                                           |
| 5     | The client directs its request to the GRM L7 server whose IP address was returned by the LDNS.                                                                            |


| Stage | Description                                                                                                                                                                                                                  |
|-------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6     | The GRM L7 server applies its policies and other criteria to determine eligible caches and determines which GRM agent cache is closest to the client. The GRM L7 server then sends the rewritten request back to the client. |
| 7     | The client sends the request to the GRM agent cache to which it was redirected.                                                                                                                                              |
| 8     | The GRM agent cache serves the content to the requesting client.                                                                                                                                                             |

# Using a GRM DNS server to locate a GRM L7 server

## Using a GRM DNS server to locate the GRM L7 server

You can use one or more GRM DNS servers to direct client requests to GRM L7 servers close to the clients. (A deployment in which you are using both GRM DNS servers and GRM L7 servers is referred to as a *mixed hierarchy*.)

For simplicity, the following illustration shows a mixed hierarchy with only one GRM L7 server. Typically, multiple GRM DNS servers and multiple GRM L7 servers would be used.



In the previous illustration, the GRM L7 server is also configured as a GRM agent cache. The reason is that the GRM DNS agent on the GRM L7 server can report proximity measurements and health status to the GRM DNS server at the top of the hierarchy. (GRM servers cannot report this information to each other.) To determine the closest GRM L7 server to a client, the GRM DNS server uses the information reported by the agents on the GRM L7 servers to determine the distance between each GRM L7 server and the requesting client's local DNS nameserver.


The following table describes the path of the request in the previous illustration.

| Stage | Description                                                                                                                                                                                                                                        |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1     | The client queries its local DNS nameserver (LDNS) for the IP address of the origin server of <a href="http://www.netcache.com">www.netcache.com</a> .                                                                                             |
| 2     | The LDNS forwards the request for resolution to the GRM DNS server that has taken over the IP address of the authoritative DNS nameserver.                                                                                                         |
| 3     | The GRM DNS server (also acting as the authoritative DNS server responsible for the <a href="http://netcache.com">netcache.com</a> domain) resolves the request and returns to the LDNS the IP address of the closest GRM L7 server to the client. |
| 4     | The LDNS returns the IP address of the GRM L7 server to the client in its response to the client's DNS request.                                                                                                                                    |
| 5     | The client directs its request to the GRM L7 server whose IP address was returned by the LDNS.                                                                                                                                                     |
| 6     | The GRM L7 server applies its policies and other criteria to determine eligible caches and determines which GRM agent cache is closest to the client. The GRM L7 server then sends the rewritten request back to the client.                       |
| 7     | The client sends the request to the GRM agent cache to which it was redirected.                                                                                                                                                                    |
| 8     | The GRM agent cache serves the content to the requesting client.                                                                                                                                                                                   |

# Using the network infrastructure to locate a GRM L7 server

## Using a device in the networking infrastructure to locate the GRM L7 server

If your network has in place an intelligent means of selecting the closest GRM L7 server, you can use your network infrastructure to direct client requests to the GRM L7 servers, as the following example shows.



In this type of deployment, all GRM L7 servers are authoritative for the same set of domains, and they all maintain proximity information for the agent caches within those domains. Transport time is saved because the device in the network infrastructure selects the GRM L7 server that is close to the client. For example, the request from a client in Asia is sent to the GRM L7 server in Asia, not to the GRM L7 server in the United States or Australia.

The following table describes the path of the request in the previous illustration.

| Stage | Description                                                                                                                                                                                                                  |
|-------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1     | The client makes a request (www.netapp.com in the example), which is received by the networking infrastructure.                                                                                                              |
| 2     | The networking infrastructure identifies the closest GRM L7 server to the client. It then redirects the client to that GRM L7 server. (The method used to identify the closest GRM L7 server varies by device.)              |
| 3     | The client sends the request to the closest GRM L7 server.                                                                                                                                                                   |
| 4     | The GRM L7 server applies its policies and other criteria to determine eligible caches and determines which GRM agent cache is closest to the client. The GRM L7 server then sends the rewritten request back to the client. |
| 5     | The client sends the request to the GRM agent cache to which it was redirected.                                                                                                                                              |
| 6     | The cache delivers the content to the client.                                                                                                                                                                                |

## Using a proxy.pac file with a GRM L7 server

---

### Using a proxy.pac file to locate GRM agent caches

An alternative to using a device to direct client requests to a GRM L7 server is to configure client browsers to point to a GRM L7 server with an automatic proxy configuration file that supports L7 redirection. The automatic proxy configuration file (typically called a *proxy.pac file*) is used to direct requests to the GRM agent caches close to clients.

A proxy.pac file consists of a JavaScript function. You create the file by entering the JavaScript function in the file. Each time a client requests a URL, the JavaScript function is executed, and it returns the IP addresses that should be used for retrieving the resource.

For each GRM L7 server that you deploy, you create a proxy.pac file and include in it specific keywords that the GRM L7 server recognizes. You then upload that file to the GRM L7 server.

To enable client requests to reach a specific GRM L7 server, you must determine which GRM L7 server is closest to each client. You must then manually configure each client browser to point to the GRM L7 server that is closest to that client.

#### Note

---

For general information about automatic proxy configuration files, see the *Deployment Guide*.

---

### How the GRM L7 server uses the proxy.pac file

When the GRM L7 server responds to a client request, it makes a copy of the proxy.pac file and dynamically replaces the GRM-specific keywords in the proxy.pac file with one or more lists of IP addresses of qualifying GRM agent caches. To determine the list of qualifying caches, the GRM L7 server applies its rules to determine eligibility, for example, its edge cache and domain and policy filtering rules.

### How load balancing works with a proxy.pac file

You can implement load balancing over GRM agent caches by including multiple lists of GRM-specific keywords (as described in [“What you need to modify in the proxy.pac file for the GRM L7 server”](#) on page 43). In this way, you control the number of lists of GRM agent caches that the GRM L7 server sends to the client by the values you include in the proxy.pac file.

The GRM L7 server includes a mechanism to support load balancing over GRM agent caches when a proxy.pac file is used. The GRM L7 server rotates the order of IP addresses in the proxy.pac file keywords that it replaces. The reason is that client browsers always begin with the first IP address in the list and proceed to the second IP address only if the first cache is unavailable. If the majority of clients received the same list (in the same order), the first GRM agent cache would be selected by all the browsers unless the cache was unavailable.

**What you need to modify in the proxy.pac file for the GRM L7 server**

To use a proxy.pac file to direct client requests to a GRM L7 server, you must replace certain standard lines in the proxy.pac file with keywords that the GRM L7 server recognizes.

Before uploading the proxy.pac file to a GRM L7 server, replace all instances of the following lines in the file:

PROXY IP1:port; PROXY IP2:port

with

PROXY GRM-IP-LIST:port

The following proxy.pac file example shows the file after the instances of PROXY IPx:port were changed to PROXY GRM-IP-LIST:port.

```
function FindProxyForURL(url, host)
{
    var hash = 0;

    hash = (host.length %2);

    if (url.substring(0,5) == "https")
        return "DIRECT";

    if (url.substring(0,5) == "snews")
        return "DIRECT";

    if (hash == 0) {
        return "PROXY GRM-IP-LIST:8080; DIRECT";
    } else {
        return "PROXY GRM-IP-LIST:8080; DIRECT";
    }
}
```

The DIRECT statement sets up the file so that if all GRM agent caches whose IP addresses are returned are unavailable, the client attempts to fetch the content from the origin server directly.

When the GRM L7 server sends the copy of the proxy.pac file to the client in response to a client request, the GRM L7 server rewrites the keyword lines PROXY GRM-IP-LIST:port with the IP addresses of qualifying GRM agent caches. The following example shows the same proxy.pac file as in the previous example, but with the two PROXY IP1:port; PROXY IP2:port lines overwritten by the GRM L7 server. Each line includes the list of two IP addresses for qualifying center caches, and each list ends with the DIRECT statement.

```
function FindProxyForURL(url, host)
{
    var hash = 0;

    hash = (host.length %2);

    if (url.substring(0,5) == "https")
        return "DIRECT";

    if (url.substring(0,5) == "snews")
        return "DIRECT";

    if (hash == 0) {
        return "PROXY 10.56.18.123:8080; PROXY 10.56.18.125:8080;
            DIRECT";
    } else {
        return "PROXY 10.56.18.125:8080; PROXY 10.56.18.123:8080;
            DIRECT";
    }
}
```

When it rewrites the file, the GRM L7 server rotates the IP addresses 10.56.18.123:8080 and 10.56.18.125:8080 between clients so that each list begins with a different address. This process enables load balancing over GRM agent caches.

### **How the GRM server determines the number of IP addresses to return**

The number of IP addresses that the GRM L7 server writes to each of the replaced PROXY GRM-IP-LIST:port lines in the proxy.pac file is determined by the value of the Maximum Number of Caches option, which defaults to 4. If the Maximum Center Caches option is set to 2 and there are two qualifying center caches (as is the case in the previous example), the GRM L7 server rewrites each of the two lines in the file to include the two IP addresses. To ensure load balancing, the GRM L7 server reverses the order in which it writes the IP addresses for each line when it responds to different client requests.

The Maximum Center Caches option is on the GRM Server - General page in the Appliance Manager.

**Configuring client browsers to access a proxy.pac file**

The proxy.pac configuration is supported by both Netscape Navigator and Microsoft Internet Explorer browsers. For example, to configure the Microsoft Internet Explorer browser for proxy.pac file use, you configure the following Internet Options > LAN settings:

- ◆ In the “Automatic configuration” section
  - ❖ Check the “Use automatic configuration script” box.
  - ❖ In the “Address” field, enter the IP address of the GRM L7 server (appliance) that is closest to the client.
- ◆ In the “Proxy server” section
  - ❖ Check “Use a proxy server box.”

**Uploading a proxy.pac file to the GRM L7 server**

For details about uploading the proxy.pac file to an appliance, see the Appliance Manager online Help.

**About this chapter**

This chapter describes how GRM handles content and provides information to help you determine the best method for your organization to use for handling cache misses.

**Topics in this chapter**

This chapter contains the following topics:

- ◆ [“How GRM obtains content”](#) on page 48
- ◆ [“Planning for handling cache misses for CDN content”](#) on page 49
- ◆ [“Using a GRM agent cache to access the Internet also”](#) on page 51

## How GRM obtains content

---

### How center and edge caches obtain content

In typical GRM deployments, GRM center and edge caches handle only CDN content in order to provide premium service to subscribers. CDN content can be

- ◆ Pulled on demand in response to client requests
- ◆ Pushed through the ContentDirector™ software

For an alternative deployment, see [“Using a GRM agent cache to access the Internet also”](#) on page 51.

### Content pushed through ContentDirector

ContentDirector is a Network Appliance software product that enables you to store content in a central location, such as a data center, and push that content to NetCache appliances located throughout the network. ContentDirector can be integrated with GRM to enable the GRM server to send content requests to the closest GRM agent cache that contains the requested content. This redirection combines the speed of prepopulated content and higher data integrity of geographic proximity.

With ContentDirector, you can target specific content for particular locations and make the content available before peak demand. For example, if you are providing a specific set of movie clips for a CDN, you can push those clips to GRM agent caches at off-hours, when demand on the network is low.

ContentDirector configuration provides the GRM server with information about domains, center groups, and center caches—information that administrators without ContentDirector who set up a GRM server must configure.

See the ContentDirector documents on the NOW Web site (<http://now.netapp.com>) for more information about ContentDirector and how to set up ContentDirector and GRM to work together.

---

#### Note

You can use the ContentReporter application from WebWasher Corporation to plan for and bill for content. ContentReporter gathers log files from each NetCache appliance and records the information in the ContentReporter database. You can then use third-party applications to create reports that help you with billing and with planning content to push. ContentReporter is available from Network Appliance sales representatives and from WebWasher Corporation.

---

# Planning for handling cache misses for CDN content

---

## Available methods for handling cache misses

If a center or edge cache to which a client sends a request for CDN content does not have the requested content (that is, when a cache miss occurs), the cache fetches it from the CDN content servers directly. You must determine the method to use for changing (rewriting) the host name in client requests so that looping does not occur with cache misses for CDN content. When host name rewriting is set up, any upstream transparent caches that intercept the request will not redirect traffic back through the CDN. Therefore, looping is avoided. This looping would occur, for example, if a Web accelerator were deployed in front of an enterprise's origin server.

## Determining the appropriate method

Available methods are

- ◆ Method 1: DNS CNAME rewriting
- ◆ Method 2: Configuration of NetCache accelerator rules

These methods are applicable for deployments using GRM DNS servers, GRM L7 servers, or both.

| Use DNS CNAME rewriting if...                                                                                                                              | Configure NetCache accelerator rules on center and edge caches if...                                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| You are redirecting URLs for a large number of Web sites. In this case, using DNS CNAME rewriting is more efficient than using NetCache accelerator rules. | You are redirecting URLs for a few Web sites only. In this case, it is simpler to use NetCache acceleration rules. |
| You are redirecting URLs for multiple customers.                                                                                                           | You are redirecting requests for a single customer with one origin server.                                         |

### Note

If the GRM DNS server takes over a content provider's existing DNS, do not use accelerator rules. If you do so, you will duplicate your effort.

**About method 1:  
DNS CNAME  
rewriting**

A CDN hoster can configure its DNS nameserver to return a canonical name (CNAME) with the associated IP address. Returning the CNAME along with the IP address avoids looping that might occur otherwise. The cache can then fetch content from the origin server by using the canonical name and store the content under the original name.

For example, assume that the CDN hoster asked a content provider to rewrite URLs to `images.enterprise.cdnhoster.com`. No origin server actually exists that can serve content for `images.enterprise.cdnhoster.com` if a cache must resolve a cache miss. Therefore, the CDN hoster adds to its DNS database a CNAME record, `images.enterprise.com`, that corresponds to the fully qualified domain name (FQDN) `images.enterprise.cdnhoster.com`. Subsequently, when a cache miss occurs, the cache uses the CNAME to rewrite the host header and fetches the content from the content provider's origin server (DNS CNAME rewriting).

**About method 2:  
configuring  
accelerator rules on  
GRM center caches  
or edge caches**

Using NetCache accelerator rules for host-name rewriting eliminates the need for the DNS nameserver that can rewrite the host name and configuring DNS CNAME rewriting. On each center or edge cache, you must configure standard NetCache accelerator rules with origin server addresses to force the cache to serve content from the origin servers only.

A disadvantage of using accelerator rules is that this method is not very scalable.

If you are using accelerator rules, Network Appliance recommends that you use the NetCache multiple-cache management feature. Multiple-cache management simplifies distribution of the accelerator rules among the caches.

## Using a GRM agent cache to access the Internet also

---

### Requirements for a cache to access the Internet

A NetCache appliance configured as a GRM agent can be used to access content on the Internet, as well as to handle CDN content. However, Network Appliance recommends that you use your GRM agent caches to handle only CDN content delivery, which provides your customers with the most preferential service possible.

If you must use the same NetCache appliance to provide both CDN content delivery and Internet service, you must do the following in addition to GRM agent configuration to redirect requests for content on the Internet to your NetCache appliances:

- ◆ Deploy an L4 or L7 switch or a WCCP router in front of the center or edge cache
- ◆ Configure the center or edge cache for transparent proxying

### For a deployment with a GRM DNS server

A local DNS nameserver sends DNS queries for domains for which the GRM DNS server is authoritative to the GRM DNS server. The GRM DNS server then returns the IP addresses of center and edge caches to the client.

A local DNS nameserver sends DNS queries for servers not hosted by the CDN up the hierarchy of DNS nameservers until a match occurs. The L4 or L7 switch WCCP router redirects the client request to one of the caches. If a cache miss occurs, the cache fetches the requested content from the appropriate origin server on the Internet.

### For a deployment with a GRM L7 server

After a client request is directed to a GRM L7 server (through one of the methods described in “[About a GRM L7 server—Application-level \(L7\) redirection](#)” on page 7), the GRM L7 server rewrites the request. The L4 or L7 switch or WCCP router redirects the client request to one of the caches. If a cache miss occurs, the cache fetches the requested content from the appropriate origin server on the Internet.



**About this chapter**

This chapter describes the required DNS setup on the GRM server and your DNS nameservers to enable the GRM server to perform DNS resolution for CDN and non-CDN clients.

**Topics in this chapter**

This chapter contains information about the following topics:

- ◆ Section A, “[Planning for interaction with a back-end DNS nameserver](#),” on page 54
- ◆ Section B, “[Planning for supporting queries for CDN content](#),” on page 56

## Section A: Planning for interaction with a back-end DNS nameserver

---

### Configuring the GRM DNS server to resolve non-CDN DNS queries

The GRM DNS server must proxy some DNS queries. The GRM DNS server resolves local DNS queries only for FQDNs associated with the CDN network (GRM domains). Queries for FQDNs not serviced by the GRM DNS server and zone transfers are proxied by the GRM DNS server to the back-end DNS nameserver.

To enable the GRM DNS server to proxy these DNS requests, you must configure it to use the back-end DNS nameserver as its nameserver. You do so either by listing the back-end DNS nameserver in the Nameservers option or the Domain-to-Nameserver Mapping option on the DNS - General page in the Appliance Manager.

### GRM DNS server acting as the authoritative DNS server for content provider domains

A GRM DNS server can assume the identity of the existing authoritative DNS nameserver for the CDN name space, that is, the GRM DNS server can assume the IP address that was previously assigned to the authoritative DNS nameserver. The authoritative DNS nameserver is then reconfigured with a new IP address.

If you want the GRM DNS server to act as the authoritative DNS server for content provider domains, you must do one of the following:

- ◆ You must configure the GRM DNS server to take over the IP address that was previously assigned to the authoritative DNS nameserver. You must reconfigure the authoritative DNS nameserver with a new IP address.
- ◆ You must change the DNS hierarchy to point to the GRM DNS server.

If multiple back-end nameservers exist for a GRM DNS server, the GRM DNS server can assume the identity of the back-end nameservers (with IP address aliasing). If you cannot use IP address aliasing, you can change the IP address of the back-end nameserver to that of the GRM DNS server in the DNS hierarchy.

---

#### Note

Some GRM deployments do not require the GRM DNS server to handle all DNS requests for the CDN namespace. See Section B, “[Planning for supporting queries for CDN content](#),” on page 56 for more information.

---

### **Assigning an additional IP address with the Appliance Manager:**

Assigning an additional IP address to an interface (called an *alias*) enables you to administer the GRM DNS server by its standard host name, instead of by the host name assigned to the back-end DNS nameserver. See the online Help for information about how to assign an alias to an interface.

### **Supporting multiple back-end nameservers in different domains**

Multiple DNS nameservers can serve as back-end nameservers for one GRM DNS server. This feature enables you to store content from different domains on different back-end nameservers. For example, one back-end nameserver for a GRM DNS server could store content for domain abc.com, and another back-end nameserver could store content for lmn.com.

If your deployment includes multiple back-end nameservers for a GRM DNS server and those nameservers store content from different domains, you must configure the GRM DNS server to associate specific back-end nameservers with particular domains. The GRM DNS server can then provide GRM agent caches with information to enable them to determine which back-end DNS nameserver should receive the content miss.

---

#### **Note**

A back-end nameserver for a GRM DNS server can be in the same location as the GRM server. Alternatively, a back-end nameserver can be in a different location, and the GRM DNS server can point to it.

---

### **DNS setup to support proximity measurements**

If you are using a GRM DNS server, make sure that you have configured DNS so that DNS queries are sent directly to the authoritative DNS nameserver instead of relying on recursion as a way for DNS queries to reach the authoritative DNS nameserver. Otherwise, the GRM DNS server cannot obtain accurate proximity information.

## Section B: Planning for supporting queries for CDN content

---

### About this section

One of your planning tasks is to determine how to redirect requests for CDN content for each content provider for which you provide CDN services. This section describes the methods that are available.

### Topics in this section

This section includes the following topics:

- ◆ [“Methods for supporting queries for CDN content”](#) on page 57
- ◆ [“Method 1: GRM provides DNS services for a content provider”](#) on page 59
- ◆ [“Method 2: URL rewriting to direct requests for CDN content”](#) on page 62

# Methods for supporting queries for CDN content

---

## Available methods

In general, a GRM DNS server can be deployed to receive DNS queries for

- ◆ The entire namespace for a content provider. In this case, the GRM DNS server takes over the DNS services for a content provider.
- ◆ Only specific domains within the CDN namespace. In this case, URL rewriting is used.

You need to negotiate with each content provider for which you provide CDN services to determine the method the content provider wants to use to redirect requests for CDN content. Some content providers might prefer to have the GRM DNS server handle all queries for an entire domain. Other content providers might prefer to use URL rewriting to selectively redirect requests to CDN content.

For example, a content provider just beginning service with a CDN hoster might want to rewrite only one URL initially to test the quality of service provided before committing to full service. Some content providers might be reluctant to expose their DNS configuration information to an outside organization.

---

### Note

You must also set up host-name rewriting to handle cache misses for CDN content that is not cached in a GRM agent cache and, therefore, must be fetched from a CDN origin server. See “[Planning for handling cache misses for CDN content](#)” on page 49 for more information.

---

## Relationship between GRM domain list entries and the content handled

The entries in the GRM domain list reflect the content handled by the CDN, as shown in the following table.

| For this content-handling method... | The GRM domain list contains these types of entries...                                                                                                          |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL rewriting                       | Domain names. Typically, FQDNs would not be listed.<br><br>The URLs that the content provider has rewritten determine the content that is sent through the CDN. |

| For this content-handling method...                     | The GRM domain list contains these types of entries... |
|---------------------------------------------------------|--------------------------------------------------------|
| The GRM server taking over DNS for the content provider | FQDNs                                                  |

## Method 1: GRM provides DNS services for a content provider

---

### DNS requirements when a GRM DNS server takes over DNS services

For deployments in which the GRM DNS server will handle all DNS queries for a content provider namespace, the GRM DNS server must assume the identity of the existing authoritative DNS nameserver. Masquerading as the authoritative DNS nameserver ensures that client DNS queries are ultimately resolved by the GRM DNS server. To achieve this configuration, you must assign the GRM DNS server the public IP address that was previously assigned to the authoritative DNS nameserver, or change the DNS hierarchy to point to the GRM DNS server IP address. The authoritative DNS nameserver must then be reconfigured with a new IP address.

After the GRM DNS server has assumed the identity of the back-end DNS nameserver, all DNS queries for domains associated with the CDN are resolved by the GRM DNS server. DNS queries for non-CDN domains and other DNS operations are proxied to the back-end DNS nameserver.

For example, consider a content provider that wants to redirect all requests for `images.foo.com` to the GRM DNS server. To achieve the redirection, the content provider defines `images.foo.com` in the GRM domain list on the GRM DNS server.

If the GRM DNS server receives a local DNS (LDNS) query for `movies.foo.com`, it determines whether `movies.foo.com` should be redirected. Because `movies.foo.com` has not been defined in the GRM domain list, the GRM DNS server proxies the query to the back-end DNS nameserver, which responds to the LDNS query.

### Determining the best location of the GRM DNS server

When the GRM DNS server takes over the DNS services for a content provider, the number of content providers that the CDN hoster services is a determining factor for the best location for the GRM DNS server.

#### ◆ Scenario 1: The CDN hoster provides DNS services for a single content provider

When the CDN hoster provides DNS services for only one content provider, the best solution is to deploy the GRM DNS server in the content provider's data center with the content provider's DNS nameserver.

The GRM DNS server is responsible for querying the appropriate back-end nameserver if the GRM DNS server does not redirect a particular FQDN.

The GRM DNS server also informs the GRM agent caches about back-end DNS nameservers to use for resolving cache misses.

If a cache miss occurs, a GRM agent cache queries the same back-end nameservers if host-name rewriting is through DNS CNAME rewriting. Alternatively, if host-name rewriting is through NetCache accelerator rules, the GRM agent cache queries origin servers directly.

◆ **Scenario 2: The CDN hoster provides DNS services for multiple content providers**

In this scenario, a content provider’s DNS nameservers and the GRM DNS server are not collocated. The GRM DNS server is in the CDN hoster’s data center. Each content provider continues to “own” its data. By pointing the GRM DNS server to a DNS nameserver in the content provider’s location, the content provider’s primary DNS nameserver does not need to send updates to the slave nameservers (the GRM DNS servers) in the data center in order to keep files synchronized.

The GRM DNS server is responsible for querying the appropriate back-end nameserver if the GRM DNS server does not redirect a particular FQDN. The GRM DNS server also informs the GRM agent caches about back-end DNS nameservers to use for resolving cache misses.

If a cache miss occurs, a GRM agent cache queries the same back-end nameservers if host-name rewriting is through DNS CNAME rewriting. Alternatively, if host-name rewriting is through NetCache accelerator rules, the GRM agent cache queries origin servers directly.

---

**Note**—

Enabling the Iterative DNS feature and using it in combination with the Nameserver-to-Domain Mapping option reduces the query load on the back-end nameservers. These options are on the DNS - General page in the Appliance Manager (Setup tab > DNS > General).

---

**Configuration overview**

The following table provides of an overview of the DNS-related steps required to complete GRM-related configuration if the GRM DNS server is taking over DNS for a content provider.

| Step | Action                                                                                                                                      |
|------|---------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Assign an additional, private IP address to the current authoritative DNS nameserver. The new IP address should not be publicly accessible. |

|   |                                                                                                                                                                                                                                                                            |
|---|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2 | Run the setup program on the GRM DNS server and assign a permanent IP address when prompted.                                                                                                                                                                               |
| 3 | Install the GRM server license on the GRM DNS server.<br>See the Appliance Manager online Help for instructions.                                                                                                                                                           |
| 4 | Set up a method for host-name rewriting to handle cache misses for CDN content.<br><br>See <a href="#">“Planning for handling cache misses for CDN content”</a> on page 49.                                                                                                |
| 5 | Remove the public IP address from the authoritative DNS nameserver.                                                                                                                                                                                                        |
| 6 | Assign the public IP address previously assigned to the authoritative DNS nameserver to the GRM DNS server, or change the DNS hierarchy to use the GRM DNS server’s IP address.<br><br>The authoritative DNS nameserver is now referred to as the back-end DNS nameserver. |
| 7 | Test the connectivity to the back-end DNS nameserver.                                                                                                                                                                                                                      |
| 8 | Complete the GRM server configuration.<br><br>See the Appliance Manager online Help for instructions.                                                                                                                                                                      |

## Method 2: URL rewriting to direct requests for CDN content

---

### URL rewriting defined

URL rewriting is a technique used to change links in Web pages to direct requests to a different location. With URL rewriting, a new domain is created for a content provider. The GRM DNS server does not take over the DNS services for an existing authoritative DNS nameserver, as is the case with method 1.

URL rewriting is used by CDNs that host content for a content provider so that cacheable content flows through the CDN. In this case, instead of the content provider changing its DNS to support content flowing through the CDN, the CDN asks the content provider to change the URLs in the Web pages to URLs over which the CDN has control. The CDN can then configure its DNS with the authoritative zones that it needs.

### Configuration overview

The following table provides an overview of the DNS-related steps required to complete GRM configuration if the GRM server handles only specific domains within the CDN namespace; that is, the content provider uses URL rewriting to direct client requests for CDN content.

| Step | Action                                                                                                                                                                        |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Set up the DNS nameserver that is to rewrite the host name to be authoritative for the domains controlled by the CDN.                                                         |
| 2    | Run the setup program on the GRM server and assign a permanent IP address when prompted.                                                                                      |
| 3    | Install the GRM license on the GRM server.<br>See the Appliance Manager online Help for instructions.                                                                         |
| 4    | Set up a method for host-name rewriting to handle cache misses for CDN content.<br><br>See “ <a href="#">Planning for handling cache misses for CDN content</a> ” on page 49. |
| 5    | Test the connectivity to the back-end DNS nameserver.                                                                                                                         |

| Step | Action                                                                                     |
|------|--------------------------------------------------------------------------------------------|
| 6    | Continue the GRM configuration.<br>See the Appliance Manager online Help for instructions. |



**About this chapter**

Previous chapters have described how GRM works and have included information to help you plan for and configure the components of a GRM deployment. This chapter summarizes the planning activities for a GRM deployment and provides some additional information to help you complete GRM planning before you begin configuration.

**Topics in this chapter**

This chapter contains the following topics:

- ◆ [“What you need to plan for”](#) on page 66
- ◆ [“Determining the number of GRM servers and agent caches needed”](#) on page 68
- ◆ [“Determining the protocols to run on a GRM agent cache”](#) on page 70
- ◆ [“GRM configuration overview”](#) on page 71

## What you need to plan for

The following table lists some issues to consider while planning GRM deployment for your CDN.

| Considerations                                                                                                                                                                                                                                                                                                                                                                                                    | For information to help you plan, see...                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>GRM servers</b>                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                                                                                        |
| Plan the type and number of GRM servers you need.<br>Consider the type of redirection that you need to support your organization's requirements.                                                                                                                                                                                                                                                                  | Section B, " <a href="#">Components of a GRM deployment</a> ," on page 4; Chapter 2, " <a href="#">Examples of GRM Deployments</a> ," on page 29; " <a href="#">Number of GRM servers</a> " on page 68 |
| Determine whether you need GRM domains. <ul style="list-style-type: none"><li>◆ If you are deploying GRM with ContentDirector 2.0.1 and later<br/>ContentDirector can automatically detect domains. See the ContentDirector documentation for more information.</li><li>◆ If you are not deploying GRM with ContentDirector 2.0.1 and later<br/>You must add GRM domains to the NetCache configuration.</li></ul> | " <a href="#">GRM domains and their use for IP address filtering</a> " on page 18                                                                                                                      |
| <b>GRM agent caches</b>                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                        |
| Plan the number of center and edge caches you need and where to locate them.                                                                                                                                                                                                                                                                                                                                      | " <a href="#">Determining the types of GRM agent caches you need</a> " on page 68; Section E, " <a href="#">Failover strategies</a> ," on page 26                                                      |
| Determine the protocols that each GRM agent cache will handle.                                                                                                                                                                                                                                                                                                                                                    | " <a href="#">Determining the protocols to run on a GRM agent cache</a> " on page 70                                                                                                                   |
| <b>Content management</b>                                                                                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                        |
| Decide whether you want to use ContentDirector to push content to your caches.                                                                                                                                                                                                                                                                                                                                    | The chapter on pushing content in the <i>Deployment Guide</i> and the ContentDirector product documentation (available at <a href="http://now.netapp.com">http://now.netapp.com</a> ).                 |

| Considerations                                                                                                                                                                           | For information to help you plan, see...                                                         |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Plan how you want to set up host-name rewriting to handle cache misses for CDN content that is not cached on a GRM agent cache and, therefore, must be fetched from a CDN origin server. | <a href="#">“How GRM obtains content”</a> on page 48                                             |
| <b>DNS requirements for GRM DNS servers</b>                                                                                                                                              |                                                                                                  |
| Decide how to set up the interaction between the DNS back-end nameserver and the GRM DNS server.                                                                                         | Section A, <a href="#">“Planning for interaction with a back-end DNS nameserver,”</a> on page 54 |
| Determine the method that you need to use to support DNS queries for CDN content.                                                                                                        | Section B, <a href="#">“Planning for supporting queries for CDN content,”</a> on page 56         |

# Determining the number of GRM servers and agent caches needed

---

**Number of GRM servers**

Network Appliance recommends that you deploy multiple GRM servers to avoid interruption of your content delivery service if a GRM server becomes unavailable.

**Determining the types of GRM agent caches you need**

When you determine the types of GRM agent caches that you need, consider the reasons for deploying center and edge caches, which are summarized in the following table.

| You need...                                                                  | If...                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A center cache (deployed in a center group, that is, in a data center)       | <ul style="list-style-type: none"><li>◆ You need to accommodate clients that are not members of edge groups.</li><li>◆ You need a failover mechanism for content delivery service if all edge caches in an edge group fail.</li><li>◆ You are deploying edge caches in stages and need to have a center cache handle content delivery service until edge caches are deployed.</li></ul> |
| An edge cache (deployed in an edge group)                                    | <ul style="list-style-type: none"><li>◆ You want NetCache appliances in your edge groups as part of your global request routing scheme.</li></ul>                                                                                                                                                                                                                                       |
| A single GRM agent cache identified as both a center cache and an edge cache | <ul style="list-style-type: none"><li>◆ You need the capabilities of both a center cache and an edge cache to provide failover from one edge group to another.</li></ul>                                                                                                                                                                                                                |

Your Network Appliance sales engineer can help you determine how many NetCache appliances you need for your GRM deployment and the models you need. No limit exists for the number of center and edge caches that a GRM server can support.

## **Licensing**

Each GRM server and center and edge cache requires a license. Contact your Network Appliance sales engineer or the Installed Base Group at [ibg@netapp.com](mailto:ibg@netapp.com) to purchase a license.

## Determining the protocols to run on a GRM agent cache

---

### Recommendations

You need to determine whether to run one or multiple protocol services on a specific center or edge cache. For deployments with heavy traffic loads, Network Appliance recommends that you configure a GRM agent cache for a single service, such as streaming media or Web service. Deploying dedicated GRM agent caches improves hit rates and bandwidth savings. Smaller deployments with lighter traffic loads might be able to run multiple services on a single GRM agent cache without affecting performance.

If you plan to constrain GRM domains, as described in “[GRM domains and their use for IP address filtering](#)” on page 18, ensure that you consider the protocols that are running on each center and edge cache. If, for example, you set a domain filter for HTTP and RTSP, the GRM server will return only the IP addresses of the GRM agent caches that are configured for *both* HTTP and RTSP.

**Tip:** After you have configured your GRM agent caches, you can see information about the protocols they support on the GRM Server - Caches page in the Appliance Manager (Data tab > GRM Server > Caches). This information will help you to configure your GRM domains if you plan to constrain the IP addresses that the GRM server is to return in response to a request.

# GRM configuration overview

---

**GRM configuration overview**

The following table summarizes the actions required to configure GRM components and to set up DNS to support GRM DNS servers.

| Step | Action                                                                                                                                                     |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Plan your GRM network. Document the identities, properties, and physical location of all GRM servers, domains, center caches, and edge caches.             |
| 2    | Set up DNS support for a GRM DNS server.<br><br>See Chapter 4, “ <a href="#">DNS Requirements to Support a GRM DNS Server</a> ,” on page 53.               |
| 3    | Configure the GRM server.<br><br>See the Appliance Manager online Help.                                                                                    |
| 4    | Configure protocol service and the GRM agent software on the NetCache appliances acting as GRM agent caches.<br><br>See the Appliance Manager online Help. |
| 5    | Test the GRM network.<br><br>See Chapter 6, “ <a href="#">Testing and Troubleshooting the GRM Network</a> ,” on page 73.                                   |



**About this chapter**      This chapter provides GRM testing procedures, lists common configuration problems, and provides troubleshooting tips for resolving these problems.

**Topics in this chapter**      This chapter contains information about the following topics:

- ◆ “[Overview of GRM testing](#)” on page 74
- ◆ “[Verifying the GRM network](#)” on page 76
- ◆ “[Verifying GRM client request resolution](#)” on page 82
- ◆ “[Troubleshooting the GRM configuration](#)” on page 83

# Overview of GRM testing

---

## Required steps for GRM testing

The following table lists the steps required to test GRM.

| Step                                           | Action                                                                                            | See...                                                                                                                                                                                                                          |
|------------------------------------------------|---------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Verifying the GRM network</b>               |                                                                                                   |                                                                                                                                                                                                                                 |
| <b>1</b>                                       | If you are deploying a GRM DNS server, test the connectivity to the back-end DNS nameserver.      | <a href="#">“Testing connectivity to the back-end DNS nameserver”</a> on page 76                                                                                                                                                |
| <b>2</b>                                       | Verify that all center and edge caches are reporting in to the GRM server.                        | <a href="#">“Verifying that center and edge caches are reporting in”</a> on page 76                                                                                                                                             |
| <b>3</b>                                       | Verify that clients in an edge group are directed to an edge cache.                               | <a href="#">“Verifying edge group redirection for a GRM DNS server”</a> on page 76<br><a href="#">“Verifying edge group redirection for a GRM L7 server”</a> on page 78                                                         |
| <b>4</b>                                       | Verify that clients from outside configured edge groups are directed to the closest center cache. | <a href="#">“Verifying GRM redirection for clients outside of edge groups for a GRM DNS server”</a> on page 78<br><a href="#">“Verifying GRM redirection for clients outside of edge groups for a GRM L7 server”</a> on page 80 |
| <b>5</b>                                       | Verify edge group failover.                                                                       | <a href="#">“Verifying edge cache failover for a GRM DNS server”</a> on page 80<br><a href="#">“Verifying edge cache failover for a GRM L7 server”</a> on page 81                                                               |
| <b>Verifying GRM client request resolution</b> |                                                                                                   |                                                                                                                                                                                                                                 |
| <b>6</b>                                       | Verify that CDN clients can receive requested content.                                            | <a href="#">“Verifying GRM client request resolution”</a> on page 82                                                                                                                                                            |

## Using the GRM server access log for test verification

If the GRM server is configured as a DNS server, the GRM server access log provides DNS query information from interactions between DNS nameservers and the GRM server. If the GRM server is configured as an L7 server, the log reports the URL redirection information between the client and the GRM server (client IP address, the requested URL, and the agents to be returned).

You can use the GRM server access log as follows to verify that GRM is functioning as expected:

- ◆ Use the `x-grm-answer` field to determine the IP addresses of returned edge and center caches.
- ◆ Examine the `x-grm-method` field in the GRM server access log to determine why the GRM server returned each edge cache and center cache.

The `x-grm-method` header has the following possible values:

- ❖ `edge`: Only an edge cache was returned.
- ❖ `closest-center`: The closest center cache was returned.
- ❖ `random-center`: The proximity was unknown, so a random center cache was returned.

To access the GRM server access log from the Appliance Manager, select Data tab > Log Files > GRM Server.

# Verifying the GRM network

---

## Testing connectivity to the back-end DNS nameserver

If you are deploying a GRM DNS server, you must verify connectivity to the back-end DNS nameserver after you assign an authoritative DNS nameserver’s IP address to the GRM server.

At a remote console, enter the following command to initiate a zone transfer:

```
dig axfr domain @server
```

If the zone transfer is successful, the back-end DNS nameserver is receiving requests proxied to it by the GRM server.

## Verifying that center and edge caches are reporting in

After configuring the GRM options, you must verify that all center and edge caches are reporting in to the GRM server. To verify that all center and edge caches are reporting in, view the GRM server access log in the Appliance Manager. The GRM server access log lists all the center and edge caches and reports whether they are reporting their status to the GRM server.

To view the GRM server access log from the Appliance Manager, select Data tab > GRM Server > Caches.

If center or edge caches are not reporting in, see [“Troubleshooting GRM configuration issues”](#) on page 83.

## Verifying edge group redirection for a GRM DNS server

For a GRM DNS server, to verify that CDN clients in an edge group are directed to an edge cache assigned to that edge group, complete the following steps.

| Step | Action                                                                                                                                                |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Use the telnet command to access a client within the edge group.                                                                                      |
| 2    | Enter the following command:<br><b>nslookup</b><br><br>After you enter the preceding command, the nslookup prompt is displayed. For example:<br><br>> |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | Force the client to use the GRM DNS server as its local DNS nameserver by entering the following at the nslookup prompt:<br><br><b>&gt; server grm_ip_address</b>                                                                                                                                                                                                                                                                                                                                             |
| 4    | At the nslookup prompt, initiate a DNS lookup to a domain controlled by the GRM DNS server. For example:<br><br><b>&gt; images.cdn.company.com</b><br><br>The IP address of each edge cache returned by the GRM DNS server is displayed.                                                                                                                                                                                                                                                                      |
| 5    | Verify that the IP addresses displayed at the nslookup prompt are for edge caches assigned to the correct edge group.                                                                                                                                                                                                                                                                                                                                                                                         |
| 6    | Examine the x-grm-method field in the GRM server access log to determine why the GRM DNS server returned each edge cache.<br><br>See <a href="#">“Using the GRM server access log for test verification”</a> on page 75 for information about accessing and using the GRM server access log.<br><br>If the IP addresses returned by the GRM DNS server are not the IP addresses that you expected, see <a href="#">“Troubleshooting GRM configuration issues”</a> on page 83 for troubleshooting information. |
| 7    | Return to the nslookup session, exit nslookup, and restart it.<br><br>Exiting nslookup causes the client to revert to its local DNS nameserver.                                                                                                                                                                                                                                                                                                                                                               |
| 8    | Repeat Steps 4 and 5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 9    | Examine the GRM server access log to verify that the IP address of the local DNS nameserver matches the IP address listed for the DNS lookup in Step 4.<br><br>If the IP address returned by the GRM DNS server does not match or if edge caches were not returned, see <a href="#">“Troubleshooting GRM configuration issues”</a> on page 83 for troubleshooting information.                                                                                                                                |

### Verifying edge group redirection for a GRM L7 server

For a GRM L7 server, to verify that CDN clients in an edge group are directed to an edge cache assigned to that edge group, complete the following step.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Examine the <code>x-grm-method</code> field in the GRM server access log to determine why the GRM L7 server returned each edge cache.</p> <p>See <a href="#">“Using the GRM server access log for test verification”</a> on page 75 for information about accessing and using the GRM server access log.</p> <p>If the IP addresses returned by the GRM L7 server are not the IP addresses that you expected, see <a href="#">“Troubleshooting GRM configuration issues”</a> on page 83 for troubleshooting information.</p> |

### Verifying GRM redirection for clients outside of edge groups for a GRM DNS server

For a GRM DNS server, to verify that clients not associated with edge groups are directed to the closest center caches, complete the following steps.

| Step | Action                                                                                                                                                                                |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Use the <code>telnet</code> command to access a client that is not part of an edge group.</p>                                                                                      |
| 2    | <p>Enter the following command:</p> <p><b>nslookup</b></p> <p>After you enter the preceding command, the <code>nslookup</code> prompt is displayed. For example:</p> <p>&gt;</p>      |
| 3    | <p>Force the client to use the GRM DNS server as its local DNS nameserver by entering the following at the <code>nslookup</code> prompt:</p> <p>&gt; <b>server grm_ip_address</b></p> |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 4    | <p>At the <code>nslookup</code> prompt, initiate a DNS lookup to a domain controlled by the GRM DNS server. For example:</p> <pre>&gt; <i>images.cdn.company.com</i></pre> <p>The IP address of each center cache returned by the GRM DNS server is displayed.</p>                                                                                                                                                                                                                                                             |
| 5    | <p>Verify that the IP addresses displayed at the <code>nslookup</code> prompt are for the closest center caches.</p>                                                                                                                                                                                                                                                                                                                                                                                                           |
| 6    | <p>Examine the <code>x-grm-method</code> field in the GRM server access log to determine why the GRM DNS server returned each center cache.</p> <p>See <a href="#">“Using the GRM server access log for test verification”</a> on page 75 for information about accessing and using the GRM server access log.</p> <p>If the IP addresses returned by the GRM DNS server are not the IP addresses you expected, see <a href="#">“Troubleshooting GRM configuration issues”</a> on page 83 for troubleshooting information.</p> |
| 7    | <p>Return to the <code>nslookup</code> session, exit <code>nslookup</code>, and restart it.</p> <p>Exiting <code>nslookup</code> causes the client to revert to its local DNS nameserver.</p>                                                                                                                                                                                                                                                                                                                                  |
| 8    | <p>Repeat Steps 4 and 5.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 9    | <p>Check the GRM server access log to verify that the IP address of the local DNS nameserver matches the IP address listed for the DNS lookup in Step 4.</p> <p>If the IP address returned by the GRM DNS server does not match, see <a href="#">“Troubleshooting GRM configuration issues”</a> on page 83 for troubleshooting information.</p>                                                                                                                                                                                |

## Verifying GRM redirection for clients outside of edge groups for a GRM L7 server

For a GRM L7 server, to verify that clients not associated with edge groups are directed to the closest center caches, complete the following step.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Examine the <code>x-grm-method</code> field in the GRM server access log to determine why the GRM L7 server returned each center cache.</p> <p>See <a href="#">“Using the GRM server access log for test verification”</a> on page 75 for information about accessing and using the GRM server access log.</p> <p>If the IP addresses returned by the GRM L7 server are not the IP addresses you expected, see <a href="#">“Troubleshooting GRM configuration issues”</a> on page 83 for troubleshooting information.</p> |

## Verifying edge cache failover for a GRM DNS server

For a GRM DNS server, to verify that center caches are returned when all edge caches in an edge group are unavailable, complete the following steps.

| Step | Action                                                                                                                                                                          |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Log on to a workstation located within the edge group.                                                                                                                          |
| 2    | Halt all of the edge caches in the edge group.                                                                                                                                  |
| 3    | Wait one minute to ensure that the GRM DNS server identifies that all edge caches have halted.                                                                                  |
| 4    | <p>Enter the following command:</p> <p><b>nslookup</b></p> <p>After entering the preceding command, the <code>nslookup</code> prompt is displayed. For example:</p> <p>&gt;</p> |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 5    | <p>At the <code>nslookup</code> prompt, initiate a DNS lookup to a domain controlled by the GRM DNS server. For example:</p> <pre>&gt; <i>images.cdn.company.com</i></pre> <p>After the preceding command has been executed, the IP addresses of center caches returned by the GRM DNS server should be displayed. No edge cache IP addresses should be returned.</p> <p>If no center caches are returned or if edge caches are returned, see <a href="#">“Troubleshooting the GRM configuration”</a> on page 83.</p> |

### Verifying edge cache failover for a GRM L7 server

For a GRM L7 server, to verify that center caches are returned when all edge caches in an edge group are unavailable, complete the following step.

| Step | Action                                                                                                                                                                                                                                                                                         |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Examine the <code>x-grm-method</code> field in the GRM server access log to verify that center caches were returned.</p> <p>See <a href="#">“Using the GRM server access log for test verification”</a> on page 75 for information about accessing and using the GRM server access log.</p> |

# Verifying GRM client request resolution

---

## About GRM client testing

After you have completed the GRM testing procedures described in “[Verifying the GRM network](#)” on page 76, you should test GRM from the client browser to ensure that clients are redirected to the correct center or edge caches. The results of these tests should be identical to the test results obtained previously.

## Testing GRM functionality for CDN clients

To test GRM functionality for CDN clients, you must test the same GRM functionality described in “[Verifying the GRM network](#)” on page 76. However, in this case, you use a client browser to verify GRM functionality.

Complete the following tests:

- ◆ Verify edge group redirection.  
Access a client browser from within an edge group and request content from a domain controlled by the GRM server. Use the GRM server access log to verify that the IP addresses of one or more edge caches in the edge group were returned. For information about using the GRM server access log, see “[Using the GRM server access log for test verification](#)” on page 75.
- ◆ Verify redirection for clients outside of edge groups.  
Access a client browser that is not in an edge group and request content from a domain controlled by the GRM server. Use the GRM server access log to verify that the IP addresses of one or more of the closest center caches were returned.
- ◆ Verify center cache failover.  
Halt all edge caches and access a client browser from within an edge group. (Ensure that the client browser supports failover.) Request content from a domain controlled by the GRM server, and use the GRM server access log to verify that the IP addresses of one or more center caches were returned.

For troubleshooting information, see “[Troubleshooting the GRM configuration](#)” on page 83.

## Troubleshooting the GRM configuration

---

### Troubleshooting GRM configuration issues

The following table lists common GRM configuration problems and provides information about how to troubleshoot them.

| If...                                                                                                                                                              | Then...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| You have deployed a GRM DNS server and the IP address in the GRM server access log is the IP address of the local DNS nameserver but no edge caches were returned. | Verify that the IP address of the local DNS nameserver is covered by the CIDR blocks specified in the edge group definition specified on the GRM server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| You have deployed a GRM DNS server and the IP address in the GRM server access log is different from the IP address of the local DNS nameserver.                   | Troubleshoot the local DNS infrastructure to determine why the local DNS nameserver does not contact the authoritative DNS nameserver (GRM server) directly.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| One or more of the center or edge caches are not reporting status to the GRM server.                                                                               | <p>Verify that</p> <ul style="list-style-type: none"><li>◆ Port 89 (the port on which the GRM server and GRM agent caches communicate) has not been blocked by any firewalls that are located between the agents and the GRM server.</li><li>◆ The GRM server software has been enabled on the GRM server.</li><li>◆ The center or edge cache has a GRM license installed.</li><li>◆ The GRM agent software has been enabled on the center or edge cache.</li><li>◆ The center or edge cache has been identified by a host name that the GRM server is able to resolve.</li></ul> <p>If you think that the GRM server cannot resolve the host name, reconfigure the center or edge cache information in the GRM server so that the center or edge cache is identified by an IP address.</p> |

| If...                                                                                                                                    | Then...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>You expected the GRM server to return the IP addresses of edge caches, but none was returned.</p>                                     | <ol style="list-style-type: none"> <li>1. Ensure that the edge group definition (specified on the GRM server) contains the correct CIDR blocks.</li> <li>2. Verify that at least one edge cache belongs to the specified edge group.</li> <li>3. Check the GRM Server - Caches page to verify that the edge cache was not excluded because it was overloaded. To access this page, select Data Tab &gt; GRM Server &gt; Caches.</li> </ol>                                                                             |
| <p>You defined an edge group to include center caches for failover, but the IP addresses of center caches were not returned.</p>         | <p>Check the Edge Groups and Caches page to verify that the edge group was defined to include center caches. To access this page, select Setup tab &gt; GRM Server &gt; Edge Groups and Caches.</p>                                                                                                                                                                                                                                                                                                                    |
| <p>You expected the GRM server to return the IP addresses of center caches, but the IP addresses of center caches were not returned.</p> | <ol style="list-style-type: none"> <li>1. Check the GRM Server - Caches page to verify that the center caches are reporting in. Also look for connection errors. To access this page, select Data Tab &gt; GRM Server &gt; Caches.</li> <li>2. Check the GRM Server - Caches page to verify that the center cache was not excluded because it was overloaded. To access this page, select Data Tab &gt; GRM Server &gt; Caches.</li> </ol>                                                                             |
| <p>You expected the GRM server to return a close center cache, but a random center cache was returned instead.</p>                       | <p>Use the <code>grm</code> CLI command to verify that the GRM server had sufficient time to make proximity measurements to the IP address.</p> <p>For example:</p> <pre><b>grm -ip IPaddress</b></pre> <p>If no proximity measurements are reported, continue checking until proximity measurements are displayed and repeat the test.</p> <p>If no proximity measurements are reported, the center caches might not be able to reach the client. Contact your Network Appliance service engineer for assistance.</p> |

| If...                                                                                                                                                                                                                                                                                                                         | Then...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>If you are deploying a GRM DNS server or are using a proxy.pac file to direct requests to the GRM L7 server, and you expected the GRM server to return the IP addresses of more center caches.</p> <p>For example, if you defined an edge group to return three center caches, but only one center cache was returned.</p> | <p>Use the <code>grm</code> CLI command to verify the center caches have not been excluded because of their proximity to the local DNS nameserver.</p> <p>For example:</p> <p><b><code>grm -ip LDNS IP address</code></b></p> <p>The preceding command lists the proximity of center and edge caches.</p> <p><b>Example:</b> Assume that you have set the Proximity for Center Caches option to 10 milliseconds and the CLI output shows the following:</p> <pre>center1.company.com: proximity 20 center2.company.com: proximity 35</pre> <p>The closest center cache, center1.company.com, is 20 milliseconds from the local DNS nameserver. Because the proximity match is set to 10 milliseconds, only center caches that are 20 and 30 milliseconds away from the local DNS nameserver will be returned.</p> <p>To have more center caches returned, modify your proximity settings accordingly.</p> <p><b>Note</b>_____</p> <p>If the proximity range indicates that more center caches should have been returned, check the GRM Server - Caches page to verify that the center cache was not excluded because it was overloaded. To access this page, select Data tab &gt; GRM Server &gt; Caches.</p> |
| <p>You shut down all edge caches to test center cache failover, but the GRM server is still returning the IP addresses of edge caches.</p>                                                                                                                                                                                    | <p>Check the Cache Status entry on the GRM Server - Caches page to verify that all edge caches in the edge group are displayed as Down. To access this page, select Data Tab &gt; GRM Server &gt; Caches.</p> <p>Wait until all edge caches are marked as Down and retest.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

| If...                                                                                                                                                  | Then...                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>If you are deploying a GRM DNS server and the GRM Server - Caches page shows all edge caches marked as Down, but no center caches are returned.</p> | <p>Using a DNS tool, such as <code>dig</code>, identify the DNS lookup Time To Live (TTL) to determine if the local DNS nameserver is caching stale DNS information. If the TTL is too high, the local DNS nameserver configuration might need to be altered. Contact your DNS system administrator for more information.</p> |

### Viewing connection status for GRM servers and agent caches

The GRM Server - Caches page and the GRM Agents - Connection Status page on the Data tab report errors that could result in a connection problem. Through these two pages, you can detect errors from the GRM server to specific GRM agent caches and from specific GRM agent caches to the GRM server.

## Troubleshooting GRM functionality for CDN clients

If you successfully completed the testing procedures described in “[Verifying the GRM network](#)” on page 76, the client tests you performed in “[Testing GRM functionality for CDN clients](#)” on page 82 should have been successful. If your tests were unsuccessful, complete the following steps.

| Step | Action                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------|---------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Check the proxy settings on the browser to ensure that no proxy is set.         |                                                                                                                                                                                                                                                                                                                                                                                                                  |
| 2    | View the GRM server access log to determine if the GRM server is being queried. |                                                                                                                                                                                                                                                                                                                                                                                                                  |
| 3    | <b>If...</b>                                                                    | <b>Then...</b>                                                                                                                                                                                                                                                                                                                                                                                                   |
|      | The GRM server is <i>not</i> being queried                                      | If the GRM server is a GRM DNS server or a GRM DNS server is being used to direct requests to a GRM L7 server, use a DNS tool, such as <code>dig</code> , to determine if the TTL on the local DNS nameserver is causing it to cache stale DNS information. If the TTL is too high, the local DNS nameserver configuration might need to be altered. Contact your DNS system administrator for more information. |

| Step | Action                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------|----------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | The GRM server is <i>not</i> being queried (continued)                                                               | <p>If the GRM server is an L7 server and a GRM DNS server is not directing requests to it, troubleshoot the method used to direct the requests to the GRM L7 server.</p> <p>For example, if you are using a device in the network infrastructure to direct requests to the GRM L7 server, check for problems with the device. If you are using a proxy.pac file, check that the proxy.pac file is set up correctly on the GRM L7 server and that the client browsers are configured to point to the correct GRM L7 server.</p> |
|      | The GRM server <i>is</i> being queried                                                                               | <p>Review the configuration settings on the GRM agent cache. Determine if the desired protocol services are enabled and that the correct ports are open.</p> <p>On the GRM server, review the information on the Data tab &gt; GRM Server - Caches page to see whether all configured protocols are being reported to the GRM server.</p>                                                                                                                                                                                      |
| 4    | Perform any of the relevant tasks described in “ <a href="#">Troubleshooting the GRM configuration</a> ” on page 83. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

# Glossary

---

## **authoritative DNS nameserver**

In standard DNS, an authoritative DNS nameserver is one that has knowledge of the DNS domain to which the origin server belongs.

In GRM deployments, GRM DNS servers interact with authoritative DNS nameservers to resolve DNS queries for unrecognized domains and to perform other functions, such as zone transfers.

In the context of GRM, the GRM DNS server can *appear* to be the authoritative DNS server for a content provider if a GRM DNS server takes over DNS for the existing CDN namespace. In this case, the GRM DNS server assumes the identity of the authoritative DNS nameserver (hereafter referred to as the *back-end* DNS nameserver), and subsequently resolves all DNS queries for domains associated with the CDN namespace. DNS queries for non-CDN domains and other DNS operations are proxied through the GRM DNS server to the back-end DNS nameserver.

## **back-end DNS nameserver**

The authoritative DNS nameserver for the CDN. The GRM DNS server must use the back-end DNS nameserver to resolve DNS queries for non-CDN domains and to perform other DNS operations. These operations are proxied through the GRM DNS server to the back-end DNS nameserver.

## **cache hit**

The process by which a NetCache appliance receives a request for an object and can resolve the request from its cache rather than having to fetch the object from the origin server.

## **cache miss**

The process by which a NetCache appliance receives a request for an object but that object is not in the cache. The NetCache appliance then attempts to fetch the object from the origin server.

## **canonical name**

A canonical name (CNAME) is the host name of the device as specified in a DNS A record. Specifying a canonical name allows multiple aliases to be associated with a single name. *See also* CNAME record.

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>center caches</b>                         | Network appliances on which the GRM agent software is enabled that have been identified in the GRM server configuration as a center cache. Edge caches are located in <i>edge groups</i> , such as POPs, close to end users. Edge caches in a particular edge group serve only the clients associated with that edge group. That is, when a client in a specific edge group makes a request or DNS query, the GRM server does not return IP addresses of edge caches in a different edge group.       |
| <b>CIDR</b>                                  | <i>See</i> Classless Inter-Domain Routing.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Classless Inter-Domain Routing (CIDR)</b> | <p>Classless Inter-Domain Routing (CIDR) eliminates the concept of class A, B, and C networks and replaces them with a generalized IP prefix. A CIDR address includes the standard 32-bit IP address and also information on how many bits are used for the network prefix. For example:</p> <p>172.200.0.0/16</p> <p>In the preceding example, the /16 indicates that the first 16 bits of the IP address are used to identify the unique network—the remaining bits identify the specific host.</p> |
| <b>CNAME record</b>                          | <p>Canonical Name Record. A CNAME record creates an alias of a canonical name. The CNAME alias gains all properties of the original, including IP addresses. Because of this, no other record should have the same name that appears in a CNAME record.</p> <p><i>See also</i> canonical name.</p>                                                                                                                                                                                                    |
| <b>Content Delivery Network (CDN)</b>        | A network service provider that distributes content and services to optimize content delivery and charges a fee to do so.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>dig</b>                                   | A command line tool for querying DNS servers. Similar to nslookup. <i>See also</i> nslookup.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Domain Name Service (DNS)</b>             | Domain Name Service (DNS) is the standard mechanism on the Internet for advertising and accessing a variety of information about computer hosts. DNS handles mapping between host names, which humans find convenient, and                                                                                                                                                                                                                                                                            |

Internet addresses, which computers require. DNS makes host information available throughout the Internet, distributed among many sites and organizations, not just in a single central site.

**edge caches**

Network appliances on which the GRM agent software is enabled that have been identified in the GRM server configuration as an edge cache. Edge caches are located in *edge groups*, such as POPs, close to end users. Edge caches in a particular edge group serve only the clients associated with that edge group. That is, when a client in a specific edge group makes a request or DNS query, the GRM server does not return IP addresses of edge caches in a different edge group.

**edge group**

A group, such as a POP, in which edge caches are located. Edge groups are identified in the GRM server configuration.

Edge caches in a particular edge group serve only the clients associated with that edge group. That is, when a client in a specific edge group makes a request or DNS query, the GRM server does not return IP addresses of edge caches in a different edge group.

**GRM agent caches**

A generic term to refer to NetCache appliances on which the GRM agent software is enabled. In the GRM server configuration, a cache on which the GRM agent software is installed can be identified as a center cache or an edge cache.

**GRM server**

A NetCache appliance located in a CDN data center that is configured as a GRM server. You can configure a GRM server as a GRM DNS server (for DNS-based redirection services) or a GRM L7 server (for application-level (L7) redirection services).

**L4 or L7 switch**

Switches typically include multiple methods for handling traffic. When this guide refers to L4 or L7 switches, it refers to a feature on a switch that can support transparent proxying; that is, a feature that operates on the L4 or L7 layer of the OSI model. Switch names vary by vendor.

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Microsoft Media Streaming (MMS)</b>     | The streaming media control protocol used in Windows Media 3 & 4 streams.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>MMS</b>                                 | <i>See</i> Microsoft Media Streaming.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>nslookup</b>                            | Standard tool for querying nameservers in command line operating systems, for example, UNIX and LINUX.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Real Time Streaming Protocol (RTSP)</b> | An open standard for delivery of real-time media over the Internet. <i>See also</i> streaming media.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>RTSP</b>                                | <i>See</i> Real Time Streaming Protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>streaming media</b>                     | A term used to describe media files that begin playing while they are being transmitted over the network to the media player on the client computer. In contrast, conventional Web files must be downloaded entirely before the user can view them. Commonly requested types of streaming media are video and audio. Streaming media also includes interactive media, cartoon-like animations, panoramic data, and more.                                                                                                         |
| <b>streaming media cache</b>               | A specially configured NetCache appliance that intercepts client requests for streaming media (RTSP, MMS, or both). When a streaming media cache is deployed between clients and a streaming media server, streaming requests that would otherwise have been sent directly to the streaming server are sent to the cache. For live streaming media presentations, bandwidth savings are realized when multiple clients request the same unique stream. NetCache caches on-demand media streams. <i>See also</i> streaming media. |
| <b>transparent proxying</b>                | A method for providing access to NetCache appliances in which users do not have to manually configure Web browsers or media players for requests to be sent to the NetCache appliances. Transparent proxying is typically deployed with an L4 switch, and L7 switch, or a WCCP router. Less commonly, it is deployed with a policy-based router.                                                                                                                                                                                 |

**URL**

Uniform Resource Locator. To find a Web resource, you enter a name called a URL in the Location field of your Web browser, for example, *http://www.netapp.com/*.

**WCCP router**

A Cisco router running WCCP 2.0 software.

**Web cache**

A NetCache appliance configured as an intermediary server that accepts requests from clients and forwards them to a Web server or, if appropriate, services the requests from its own cache. A Web cache acts as an agent for the client browser. A NetCache appliance configured as a Web cache handles any or all of HTTP, FTP, Gopher, and Tunnel (for example, HTTPS and SSL) requests.

**zone transfer**

A special type of query in which the entire contents of a zone are requested from a nameserver. (DNS namespaces are divided up into zones, which store name information about one or more DNS domains.) Secondary DNS nameservers update locally stored zone information by requesting zone transfers from primary DNS nameservers.



# Index

---

## Numerics

302 location changed response 7

302 redirection response 9

## A

accelerator rules

when to use for content handling 49

alias for appliance interface 55

application-level redirection, defined 5

ASX file

how L7 redirection occurs with 9

authoritative DNS nameserver

defined 89

GRM DNS server interaction with 6

assuming identity of 31, 54

authoritative for CDN name space 54

passing query to nameserver 14

proxying DNS requests to 54

taking over for 54, 59, 60

recursive DNS queries and 17

testing connectivity to 76

automatic proxy configuration file

browsers that support 46

use with GRM L7 server 8, 42

example 43

load balancing with 42

uploading to appliance 46

what you modify 43

## B

back-end DNS nameserver

associating with GRM domains 55

defined 89

GRM DNS server proxying requests for 54

GRM DNS server taking over for 59, 60

testing connectivity to 76

billing for content 48

binding GRM domains to protocols 20

## C

cache hit, defined 89

cache miss

defined 89

how GRM handles 49

canonical name. *See* CNAME

CDNs

defined 90

goals for CDN service 2

requirements that GRM fulfills 2

strategy when spread over continents 7

center caches

as edge caches also 68

characteristics 10

defined 10

failover for edge caches 26

license requirement for 68

load balancing over 23, 24

planning for 68

protocol services on 70

proximity data provided 16

status to GRM server

CPU load 16

proximity 16

system health 16

testing client redirection to 78, 82

when needed 68

center groups, defined 10

CIDR address

association with an edge group 16

defined 90

GRM server evaluation of 14

CIDR blocks

association with an edge group 10

defined 90

GRM server evaluation of 14

Classless Inter-Domain Routing. *See* CIDR  
address, CIDR blocks

CNAME defined 89

CNAME records

defined 90

when URL rewriting is used 50

- CNAME rewriting
  - advantage of 50
  - defined 50
  - DNS configuration to support 50
  - when to use 49

- configuration overview, GRM 71

- content handling with GRM 47
  - accelerator rules
    - when to use 49
  - billing for content 48
  - DNS CNAME rewriting
    - when to use 49
  - fetching Internet content 51
  - how content is obtained 48
  - methods 49, 57
  - when GRM L7 server serves 15

- ContentDirector
  - planning for use with GRM 66
  - product description 48
  - pushing content to GRM caches 48

- ContentReporter
  - billing for content 48
  - product description 48
  - where to obtain 48

- CPU load, GRM server evaluation of 14

## D

- dig utility
  - defined 90
  - identifying DNS lookup TTL 86
  - initiating a zone transfer 76

- DNS
  - defined 90
  - GRM handling specific domains 62
  - requirements to support GRM DNS server 53
    - configuration overview 62
  - routing with GRM 29

- DNS CNAME rewriting
  - advantage of 50
  - defined 50
  - DNS configuration to support 50
  - when to use 49

- DNS nameserver
  - authoritative

- IP address to GRM DNS server 55
- configuring for URL rewriting 50
- redirection to GRM L7 server
  - disadvantages of 35

- DNS queries
  - non-CDN, GRM DNS server resolving 54
  - proxying by GRM DNS server 54
  - recursive 17

- DNS-based redirection 5

- DOC files
  - how L7 redirection works with 9

- Domain Name Service. *See* DNS

- domains, GRM
  - and multiple CDNs 18
  - DNS domains and 18
  - filters for 19, 20

## E

- edge caches
  - as center caches also 68
  - characteristics 10
  - defined 10, 90, 91
  - failover to center caches 10, 26
  - license requirement for 68
  - load balancing over 23, 24
  - planning for 68
  - protocol services on 70
  - status to GRM server
    - CPU load 16
    - proximity 16
    - system health 16
  - testing
    - client redirection to 82
    - reporting to GRM server 76
    - when needed 68

- edge groups
  - CIDR address associated with 16
  - defined 10, 91
  - testing client redirection to 76, 78
- enterprises spread over continents
  - deployment strategy 7

## F

- failover, GRM

- center cache, testing 82
- deployment strategies 26
- edge caches to center caches 10
- edge group, testing 80, 81
- if no GRM agent caches exist 27
- method provided 26
- reasons for using 26
- filtering for GRM domains 19
  - group filter 19
  - location filter 19
  - protocol filter 19
- firewalls
  - blocking GRM proximity messages 17
  - troubleshooting for GRM 83
- Force Redirection option, described 15
- FQDNs
  - and CNAME rewriting 50
  - using as domain filters 19
  - with GRM domains 18
- fully qualified domain names. *See* FQDNs

## G

### GRM

- advantage over transparent proxying 2
- communication port 17
- components
  - center caches 10
  - edge caches 10
  - GRM agent caches 4
  - GRM servers 4
- configuration overview 71
- content handling with 47
  - caches misses 49
  - how content is obtained 48
  - methods 57
  - taking over DNS 57
  - URL rewriting 57
- deployments
  - components 4
  - when to deploy a GRM DNS server 6
  - when to deploy a GRM L7 server 7
- DNS configuration to support 53
- failover 26
  - if no GRM agent caches exist 27

- methods 26
  - reasons for using 26
- how GRM works 31
- licensing requirements 69
- planning summary 65
- port for communication 17
- redirection methods 5
- testing 73
- troubleshooting 73
  - configuration 83
  - with CDN clients 87

### GRM agent caches

- defined 4, 91
- determining how many needed 68
- determining types needed 68
- fetching Internet content 51
  - requirements 51
- license 68
- load balancing over 23, 24
- protocols to run on 70
- status to GRM server 16
- types 10

### GRM DNS servers

- assigning IP address aliases 55
- content handling with
  - GRM domain list entries 57
  - handling all DNS queries 59
  - URL rewriting 62
- deployment example 30
- determining location of 59
- directing requests to GRM L7 servers
  - advantages and disadvantages 33
- DNS requirements
  - configuration to support 60, 62
- DNS requirements for 53
- failover 26
- load balancing mechanisms 24
- multiple back-end DNS servers with 54
- proximity measurements with 16
- proxying of DNS queries 54
- queries for CDN content
  - methods to support 57
- queries for non-CDN DNS queries, resolving 54
- redirection method 6

- redirection to GRM L7 servers 38
  - testing configuration 74
  - testing connectivity to a DNS server 76
  - when IANA database is used for proximity 15
  - when to deploy 6
  - GRM domains
    - and multiple CDNs 18
    - defined 18
    - DNS domains and 18
    - filters for 19
      - examples 20
      - FQDNs and 19
      - group 20
      - how filters work together 19
      - location 20
      - protocols 20
      - when to set 19
    - FQDNs with 18
    - identifying to GRM server 18
    - multiple back-end nameservers, association with 55
  - GRM L7 servers
    - deployment example 32
      - DNS nameserver for redirection to 35, 36
      - GRM DNS server for redirection to 38
      - network infrastructure for redirection to 40
      - proxy.pac file for redirection to 42
    - failover 26
    - load balancing mechanisms 23
    - overload with DNS nameserver redirection 35
    - proximity measurement with 16
    - redirection method 7
    - redirection of requests for large files 9
    - redirection to
      - DNS nameserver 33
      - methods 8, 33
      - using a GRM DNS server 33
      - using a proxy.pac file 34
      - using the network infrastructure 34
    - request rewriting, how it works
      - 302 request redirection response 9
      - with HTTP content 8
      - with streaming content 9
    - when they serve content directly 15
    - when to deploy 7
  - GRM server access log
    - checking after configuration 75
    - IP addresses returned 17
  - GRM servers
    - See also* GRM DNS servers, GRM L7 servers
    - center caches
      - mechanisms to determine proximity 17
      - proximity measurements 17
    - consideration of proximity data 14
    - defined 31, 91
    - factors determining IP addresses returned 13
    - failover 26
    - functions 5
    - GRM domains
      - filtering 19
      - role of 18
    - handling non-GRM content 14
    - identifying GRM domains 18
    - license requirement for 68
    - location of 68
    - number needed 68
    - proximity measurements with 16
    - status from GRM agent caches
      - communication protocol used 17
    - types 5
      - GRM DNS server 5
      - GRM L7 server 5
  - GRM testing
    - center cache failover 82
    - clients outside edge groups 78
    - connectivity to back-end nameserver 76
    - edge group failover 80, 81
    - edge group redirection 76, 78
    - overview 74
    - reporting of GRM agent caches 76
- ## H
- Host header 50
  - host-name rewriting
    - methods for 49
    - using accelerator rules 50
    - using CNAME rewriting 50
  - HTTP Redirection Suffix option, use of 15

HTTP, L7 request rewriting with 8

## I

IANA database

- use for proximity measurements 15

Installed Base Group, license purchase 69

interface aliases 55

IP addresses

- authoritative DNS nameserver

- to GRM DNS server 6, 31, 54, 55

- returned by GRM server

- how determined 12, 13

- viewing 17

## L

L4 switch versus GRM 2

L7 redirection, defined 5

L7 request rewriting

- with HTTP content 8

- with streaming content 9

license

- center cache 68

- edge cache 68

- GRM agent caches 68

licensing requirements for GRM 69

load balancing

- automatic proxy configuration file 23

- GRM DNS server mechanism 24

- GRM L7 server mechanism 23

- how GRM server handles 23, 24

- proxy.pac file 23

load, CPU 16

local DNS nameservers

- IP address list from GRM DNS server

- how long cached 24

- proximity measurements to

- DNS configuration to support 55

- obstacles 15

- recursive queries

- problem with 17

- sending queries to GRM server 51

locator device in network

- directing requests to GRM L7 server 34

- transparent proxying

how it differs from GRM 8

log files

- GRM server

- viewing 17

- GRM server access log

- configuration testing 75

log files, GRM server

- creating reports with 48

## M

Microsoft Media Streaming. *See* MMS

mixed hierarchy, defined 5

MMS, defined 92

## N

nslookup, defined 92

## O

Origin Server Failover option 27

## P

PDF files

- how L7 redirection works with 9

planning for GRM 65

- configuration planning 65

- ContentDirector use 48

- ContentReporter use 48

- deployment considerations 65

port for GRM communication 17

PPT files

- how L7 redirection works with 9

protocols

- binding GRM domains to 20

- running on GRM agent caches 70

proximity data

- client-to-edge cache

- how determined 16

- GRM DNS servers

- DNS configuration to support 55

- GRM server consideration of 14

- how center caches provide 16

- if GRM DNS server has none 15

- if GRM L7 server has none 15
- measurements
  - problem with LDNS recursive queries 17
  - with center caches 16
- mechanisms to determine proximity 17
- proxy.pac file
  - browsers that support 46
  - use with GRM L7 server 8, 42
    - example 43
    - load balancing with 42
    - uploading to appliance 46
    - what you modify 43
- pure GRM server deployment, defined 30

## Q

- queries for CDN content
  - support needed for GRM DNS server 56, 57

## R

- recursive queries with LDNS
  - problem for GRM DNS server 17
- redirection methods available with GRM 5
- request routing
  - for CDNs spread over continents 7
  - with L4 switches 2
- RTSP
  - defined 92
  - L7 request rewriting with 9

## S

- status sent to GRM server
  - communication protocol used 17
  - CPU load 16
  - proximity data 16
  - system health 16
- streaming media
  - defined 92
  - L7 request rewriting with 9
  - streaming media cache, defined 92
- system health, sent to GRM server 16

## T

- TCP port 89
  - firewalls and 83
  - used for GRM communication 17
- testing GRM 73
  - center cache failover 82
  - client to GRM agent cache redirection 82
  - clients outside edge groups 78
  - connectivity to authoritative DNS nameserver 76
  - edge group failover 80, 81
  - edge group redirection 76, 78
  - network 76
  - overview 74
  - reporting of GRM agent caches 76
- Time-To-Live. *See* TTL
- transparent proxying
  - defined 92
  - how it differs from GRM 2
- troubleshooting GRM 73
  - configuration 83
  - with CDN clients 87
- TTL, LDNS
  - caching and 24

## U

- URL rewriting
  - defined 57, 62
  - GRM domain listings for 57
  - returning CNAME 50
  - when used 62
- URL, defined 93

## W

- WCCP routers, defined 93
- Web cache, defined 93
- WebWasher ContentReporter
  - billing for content 48
  - described 48
  - where to obtain 48

## **X**

XLS files

how L7 redirection works with 9

## **Z**

zone transfer, defined 93

