Intel® NetStructure™ 1520 Cache Appliance

Administrator's Guide

Intel Corporation
5200 N. E. Elam Young Parkway
Hillsboro, Oregon 97124-6497

July 2000 A29914-001

Contents

List of Procedures

Preface

This manual describes how to use and configure an Intel® NetStructure™ Cache Appliance system (referred to as “appliance” in this manual) either as a single node or as a cluster of nodes.

The manual covers the following topics:

- *Chapter 1* contains an overview of the appliance and an overview of this guide.
- *Chapter 2* through *Chapter 1* contain procedural information about starting, monitoring, and configuring the appliance.
- *Chapter 6* contains information to help you troubleshoot problems you might have with the appliance.
- *Appendix A* contains background information about the appliance’s main components and features of the appliance.
- *Appendix B* provides error information.

Who should read this manual

This manual is intended for system administrators who configure, run, and administer Intel NetStructure Cache Appliance systems. Consequently, the information in the manual was written with the assumption that the reader has experience in Web server administration and configuring TCP/IP networking.

Conventions used in this manual

This manual uses the following conventions.

Convention	Purpose
italics	*Represent emphasis and introduce terms, for example, "the management cluster."*
bold	*Represents graphical user interface options and menu names, for example, "**Reset**"*
`monospaced font`	*Represents commands, file names, file content, computer input, and output, for example, "use the `reconfigure` command."*
`monospaced bold`	*Represents commands that you should enter literally, for example, type **`reboot`**.*
`monospaced italic`	*Represents variables for which you should substitute a value, for example, "enter a `filename`."*
brackets `[]`	*Represent optional command arguments in command syntax, for example, **`add pathname`** `[size]`*

Chapter 1

Introduction

The Intel® NetStructure™ Cache Appliance is a carrier-class caching appliance that offers high performance, high availability, and simple centralized management. The appliance automatically and efficiently copies network documents and images, bringing them closer and serving them faster to your users.

When placed strategically in a network, the appliance can serve user requests for objects from its cache or the caches of neighboring appliances rather than have requests served from an origin server. This relief results in improved network performance, and a perceived higher quality of service. At the same time, the appliances reduce Internet bandwidth usage by eliminating redundant requests for popular documents.

This chapter provides the following overviews:

What is an Intel® NetStructure™ Cache Appliance?

Internet users request billions of documents each day all over the world. Unfortunately, global data networking has become difficult for professionals as they struggle with overloaded servers trying to keep pace with society's growing data demands.

The Intel NetStructure Cache Appliance family provides you with a turnkey, scalable solution you can place in your network to deliver industry-leading caching capabilities. Your system is designed for fast and reliable caching for Internet Service Providers (ISPs), backbone providers, and large intranets.

Why use this caching appliance?

Caching can significantly reduce pressure on busy networks and servers by storing copies of popular documents near their users. Instead of making multiple requests for the same document across congested networks to overloaded servers, users access copies from the caching appliance's large, fast local cache. This reduces backbone congestion, provides faster response, and improves the quality of service.

The following design features make the Intel NetStructure Cache Appliance a carrier-class caching product:

- ✔ Speed (the ability to handle thousands of simultaneous user connections)
- ✔ Scalability (you can easily add nodes to a management cluster as needed)
- ✔ Fault tolerance (redundant boot images)
- ✔ Secure single-point administration (you can configure many nodes at once)

See *Intel NetStructure Cache Appliance features, on page 3* for more information about these features.

Flexible cache architecture

You can use the appliance alone or with other enterprise software, including other caching products. Here are some examples of ways to use the appliance.

Web proxy cache

User requests go to the appliance on the way to the destined web server. If the cache contains the requested document, the appliance serves the requested document directly. If the cache does not have the desired document, the appliance acts as a proxy, fetching the document from the web server on the user's behalf, and keeps a copy to satisfy future requests.

Server accelerator

The appliance can be configured as a web server to accelerate slower traditional web servers. Documents stored in cache are served at high speed, while documents not in cache are requested on demand from slower, traditional web servers. This *server accelerator* feature is also called *reverse proxy.*

Part of an HTTP cache hierarchy

The appliance can participate in flexible cache hierarchies, where Internet requests not fulfilled in one cache can be routed to other regional caches, taking advantage of the of nearby caches.

ICP sibling

The appliance supports the standard Internet Cache Protocol (ICP) to interoperate with existing ICP cache hierarchies. The appliance can send ICP queries to neighboring caches as part of an ICP cache hierarchy.

NNTP news cache

The appliance caches and serves NNTP news articles and can accept news feeds for designated news groups.

Intel NetStructure Cache Appliance features

The appliance provides a rich set of features to ensure high performance and superior stability and to offer broad flexibility. The following list provides a brief overview of the appliance's primary features. For a more exhaustive list and description of features, refer to *Carrier-class architecture, on page 143*.

Scalability

The appliance scales from a single node into multiple-node *clusters*, allowing you to improve system performance and reliability simply by adding more nodes to your cluster. Support exists for two types of clusters: *soft clustering* and *management-only clustering*. For more information on clustering, see *Clustering, on page 140*.

Boot Image Redundancy

The appliance features both a primary and secondary boot image on separate hard drives. When a drive with a boot image fails, a system administrator can detect and replace the faulty hard drive. This feature helps maximize the time your system is up and running uninterrupted.

Multithreading process support

The appliance is the first commercial caching proxy server to aggressively implement multithreading, breaking down large transactions into small, efficient tasks. The appliance processes multiple outstanding requests simultaneously and efficiently, even under peak loads.

High-speed caching

The cache consists of a high speed object database stored on raw disk. Objects are stored and indexed according to their URL and associated headers. This enables the appliance to store, retrieve, and serve not only web pages, but parts of web pages, providing optimum bandwidth savings.

Broad protocol support

The Intel NetStructure Cache Appliance supports the following protocols:

- ✔ HTTP versions 0.9 through 1.1
- ✔ FTP
- ✔ NNTP
- ✔ ICP
- ✔ SSL encryption
- ✔ WCCP 2.0

HTTP cache hierarchy support

In a hierarchy of proxy servers, the appliance can act either as a parent or child cache, either to other Intel NetStructure Cache Appliances, or to other caching products.

Web server acceleration

Through reverse proxy, the appliance can act as a web server accelerator, handling requests for and relieving stress from web servers.

Transparency option

With transparent interception of user traffic, user requests are automatically injected into the cache on their way to the eventual destination. Users request Internet data as usual without any browser configuration, and the appliance automatically serves their requests.

Secure, single-point administration

The appliance offers two administration alternatives to suit the needs of different environments:

- ✔ **Browser-based interface**: The Manager User Interface (UI) offers password-protected, single-point administration for an entire cluster.
- ✔ **Command-line interface**: The command-line interface lets you configure the system's network addresses and lets you control, configure, and monitor the appliance.

SNMP Network Management

The appliance can be monitored and managed through SNMP network management facilities. The appliance supports two management information bases (MIBs). The first, MIB-2 is a well known standard MIB. The second, the proprietary Intel NetStructure Cache Appliance MIB provides more specific node and cluster information.

Performance reporting

You can get performance statistics at a glance from the Manager UI or from the command-line interface.

How to use this guide

The rest of this guide contains three parts: background information, procedural chapters, and reference appendixes.

	To find out about …	See …
Procedures	*how to get started*	*Starting the system for the first time, on page 8*
	how to use the Manager UI	*Accessing the Manager UI, on page 12*
	how to monitor and configure the appliance using the Manager UI	*Using Monitor and Configure mode, on page 13*
	how to use the command line interface	*Accessing the command-line interface, on page 15*
	how to upgrade software	*Installing a new version of the appliance software, on page 56*
	how to troubleshoot system problems	*Chapter 6, Troubleshooting Problems*
Appendices	*background information including web proxy caching, transparent proxy caching, server acceleration, cache hierarchies, news article caching, and carrier-class architecture*	*Appendix A, Caching Solutions and Performance*
	error messages	*Appendix B, Error Messages*

Chapter 2

Getting Started

This chapter contains the following sections:

Starting the system for the first time

Before you can start the Intel NetStructure Cache Appliance, make sure it is physically connected properly. Connections include:

- ✔ Connecting to the network through the primary network interface.
- ✔ Connecting a Terminal Emulator or Concentrator to the appliance's COM1 port using the serial cable that came packaged with the appliance.
- ✔ Attaching the supplied power cord to the appliance and plugging the cord into an approved receptacle.

You can find instructions on how to physically set up your system in the *Intel NetStructure Cache Appliance Quick Start*.

Note Safety regulations and warranty require that the front bezel mounts and panel must be in place during operation of the appliance.

Once you have made the physical connections, you can initially configure your appliance and start it up.

▼ Initially configuring and starting your system

1 From the Terminal Emulator or Serial Concentrator, make sure you are emulating a VT100 terminal. Use these port specifications for the connection:

- ✔ 9600 baud
- ✔ 8 data bits
- ✔ No parity
- ✔ 1 stop bit
- ✔ Hardware flow control

2 From the window emulating the VT100 terminal, open the connection to the appliance.

3 Power on the appliance by pressing the power button, located behind the front bezel. Supplying power to the appliance starts the initial boot process. The initial boot process takes approximately three to four minutes. During this time random characters might appear on the screen of your VT100 terminal emulator.

Note See the *Intel NetStructure Cache Appliance Quick Start* for locations of controls and physical features on your system.

4 After your system completes the boot procedure, a console login prompt appears with fields for both a login and password. At the prompt, supply `admin` for both the login and password, and press Enter.

5 After you login, the VT100 terminal emulator screen displays this initial set of menu selections.

```
—setup      Initial Intel Cache Setup
 install    Install Intel Cache
 commit     Commit Setup Changes
```

6 Use the arrow keys to select **setup** and press the Enter key.

Note For information on how to navigate within the CLI, refer to *Navigating the command-line interface, on page 51*.

7 The setup menu appears. This menu allows you to configure network and time parameters as well as view settings you have entered.

```
—network    Configure Network
 timezone   Configure Time Zone
 time       Configure Date and Time
 view       View Settings
```

8 Use the arrow keys to select **network** and press the Enter key. The following network setup fields appear:

```
Enter IP Address        192.168.1.10_______________
Enter Hostname          Intel-NetStructure-Cache__
Enter Netmask           255.255.255.0______________
Enter Nameserver IP     ___________________________
Enter Gateway IP        192.168.1.1________________
Enter Domain            ___________________________
```

9 In each field supply an appropriate value and press the Enter key. Pressing the Enter key moves the cursor to the next field. After you have supplied values for all six fields, press CTRL+X to save your changes and return to the previous menu.

10 The bottom of the screen displays a message that indicates the setup has completed. When the message appears, entries to the screen have been successfully changed and stored. The menu on this screen should appear as follows:

```
 network    Configure Network
-timezone   Configure Time Zone
 time       Configure Date and Time
 view       View Settings
```

11 Use the arrow keys to highlight **timezone** and press the Enter key. Pressing the Enter key causes a scrollable list of available timzones to appear. Here is a partial list:

```
-United States   Eastern
 United States   Central
 United States   Mountain
 United States   Pacific
```

12 Use the arrow keys to scroll through the available zones and highlight the appropriate zone for your area. After highlighting the applicable zone, press the Enter key. Next, press any key to save your selection and return to the previous screen as follows:

```
 network    Configure Network
 timezone   Configure Time Zone
-time       Configure Date and Time
 view       View Settings
```

Note In order for the timezone change to become effective, the appliance must be rebooted. A reboot operation occurs later during the initial setup.

13 Use the arrow keys to highlight **time** and press the Enter key. Pressing the Enter key causes the following fields to appear:

```
Enable(1)/Disable(0) Daylight Savings Time__
Currently Inside (1)/Outside(0) Daylight Savings Time__
Enter Time       [ HH:MM:SS]      __:__:__
Enter Date       [ MM/DD/YYYY]    __/__/__
```

14 Set your Daylight Savings Time options. Then enter the time using a 24-hour format (e.g., for 2:14:56 PM enter 14:14:56). For each part of the format, you must press Enter to accept the value and to move to the next part of the field. For example, after entering the two-digit hour value, pressing Enter causes the value to be accepted and positions the cursor over the minutes part of the time field. Supply the date using the MM/DD/YYYY format. After supplying the date, press the CTRL-X key combination to save your changes and return to the previous menu as follows:

```
 network    Configure Network
 timezone   Configure Time Zone
 time       Configure Date and Time
-view       View Settings
```

15 From this menu you can select **view** to verify the network and time information you have entered. After you are sure all the information you have entered is correct, press the CTRL-X key combination twice to move back to the main menu as follows:

```
 setup      Initial Intel Cache Setup
-install    Install Intel Cache
 commit     Commit Setup Changes
```

16 From the main screen, highlight **install** and press the Enter key. Selecting **install** causes the settings to be written to the boot image. During the

installation, the bottom of the screen keeps you apprised of the installation's progress.

17 After the installation is complete, use the arrow keys to position the cursor on **commit** as follows:

```
 setup      Initial Intel Cache Setup
 install    Install Intel Cache
-commit     Commit Setup Changes
```

18 Pressing the Enter key starts the final phase of the initialization process as well as the cache application. The bottom of the screen indicates that the cache application has started and prompts you to press the Enter key a second time.

19 When the `Initialization Complete!` prompt appears, press the Enter key to reboot the appliance. Rebooting the appliance takes several minutes. During the reboot process, random characters might appear in the window of the VT100 terminal emulator screen.

20 After your system completes the boot procedure, a console login prompt appears with fields for both a login and password. At the prompt, supply `admin` for both the login and password, and press Enter.

21 After the login completes, the initial menu appears with additional selections:

```
 setup      Initial Intel Cache Setup
-main       Main Intel Cache Controls
 config     Intel Cache Configuration
 monitor    View Statistics
 expert     Enter Expert Mode
 save       Save Config to Floppy
 load       Load Config from Floppy
 logoff     Logoff
```

Note The system starts with factory settings. You can further configure or customize the appliance by following the guidelines in *Chapter 4, Configuring the Appliance*.

Once the software is running, you can access the system through a web browser by using the system's IP Address with an appended `:8081` as the URL. For information on accessing the manager UI, refer to *Accessing the Manager UI, on page 12*.

Accessing the Manager UI

The Manager UI is a browser-based interface, consisting of a series of web pages. Use the Manager UI to monitor performance and configure and fine-tune selected nodes in your cluster. You can access any node in the cluster through the same Manager UI.

▼ Accessing the Manager UI

1 Open your web browser.

The Manager UI requires Java and JavaScript; be sure to enable Java and JavaScript.

2 Point your browser at this location, where *nodename* is the IP address you have assigned to the appliance or the qualified DNS name. If the appliance is part of a cluster, you will be logging into that specific node:

```
http://nodename:8081/
```

3 Provide your appliance administrator's ID and password. By default, the administrator ID is `admin` and the password is `admin`. It is recommended that you change the default administrator ID and password. You can change these values by using the **Security** page. For information on how to use the **Security** page, see *Using the Security page, on page 39*.

Note Should you forget your password, contact Customer Service at Intel Corporation for assistance. For information on how to contact Intel Customer Service, see the *Intel NetStructure Cache Appliance Product Support* booklet that came with your system.

Note Changing ID and password values by using the Manager UI changes those values for the node you are logging into only. Furthermore, changing the ID and password for the Manager UI does not change the ID and password for telnet access. You must use the command-line interface (CLI) to change the telnet ID and password for the node.

The Manager UI appears in your browser in the default monitor mode. The **Dashboard** page, as shown *Figure 1*, is the default page. From the **MONITOR** and **CONFIGURE** tabs to the left of the **Dashboard** page, you can reach all other Manager UI pages.

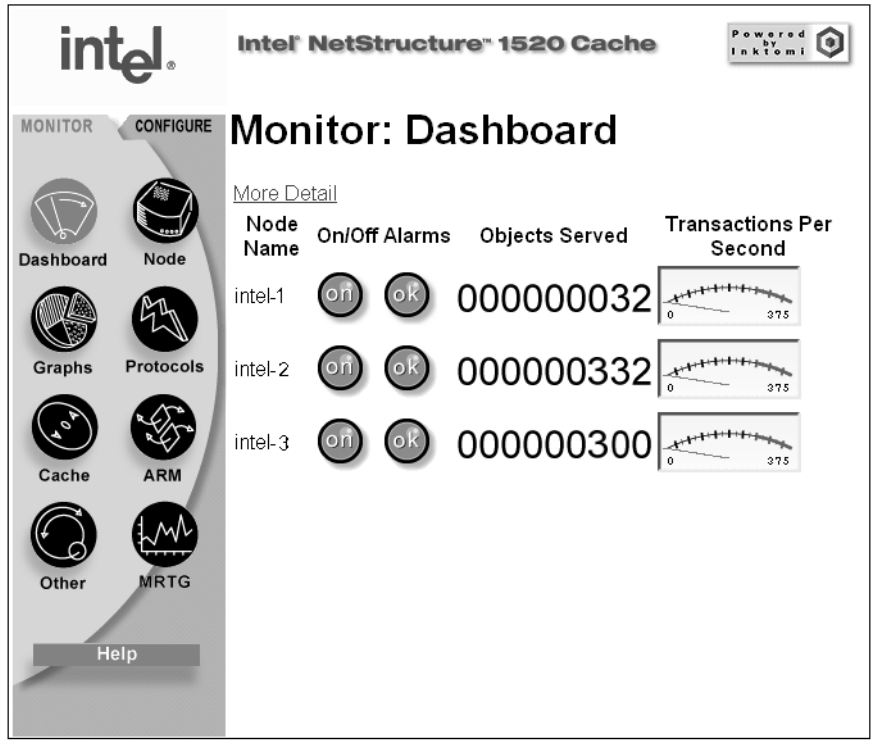

Figure 1 The Dashboard page

Using Monitor and Configure mode

The Manager UI has two modes, Monitor and Configure:

- ✔ In Monitor mode, view performance statistics and graphs. To access Monitor mode, click the top of the **MONITOR** tab.
- ✔ In Configure mode, view and modify the appliance's configuration options. To access Configure mode, click the top of the **CONFIGURE** tab.

Figure 2 shows the control frame buttons for both the Monitor and Configure modes.

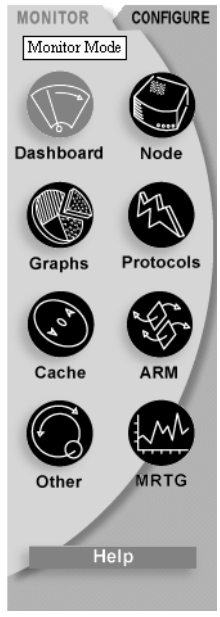

Monitor mode frame

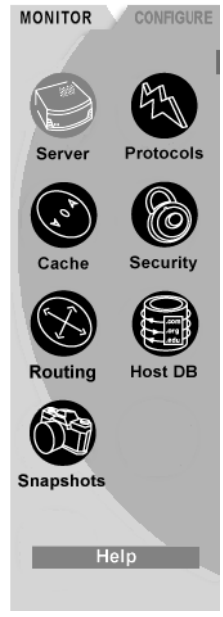

Configure mode frame

Figure 2 The Monitor and Configure Control Frames

When you are in Monitor mode, you can access all the pages that report information about the appliance's performance. With the exception of the information on the **Dashboard** page, information on the **Monitor** pages pertain to the selected node. You can change nodes at any time by returning to the **Dashboard** and clicking the node of your choice. For information about how to use each of the performance screens, see *Accessing monitor pages, on page 18*.

When you are in Configure mode, you can access pages that change system configuration values for the selected node. Each time you click the **Make These Changes** button the selected node's configuration is updated.

Note It is recommended that you save current configuration values before making any changes.

To save and restore an entire set of configuration files, refer to *Using the Snapshots page, on page 47*. For information about all the values you can set in Configuration mode, see *Chapter 4, Configuring the Appliance*.

Using online help

Both the **MONITOR** and **CONFIGURE** tabs have a **Help** page button. When you click the **Help** page button, the online help opens in another browser window. Each of the Manager UI pages has online help available.

Accessing the command-line interface

You can access the command-line interface using one of two methods:

- ✔ Provide a serial connection to the Intel NetStructure Cache Appliance machine. Refer to the *Intel NetStructure Cache Appliance Quick Start Guide* for detailed information.
- ✔ Access the machine through a telnet connection. This method requires you to enter a telnet Administrator ID and password. Refer to *Changing the administrator password for telnet or serial access, on page 60* for information on this ID and password.

For information on using the command-line interface, refer to *Chapter 1, Using the Command-Line Interface*.

Verifying that caching works

After starting the appliance, you should verify that it is up and running. To see if the appliance is processing HTTP requests, do the following:

1. From the **Monitor** tab in the Manager UI, click the **Protocols** button.
2. Make a note of the current **HTTP User Agent Total Document Bytes** statistic.
3. Set your browser to the Intel NetStructure Cache Appliance proxy port.
4. Browse the Internet.
5. Check the **HTTP User Agent Total Document Bytes** value.

   This value should have increased if caching is working.

Changing passwords

Two IDs and passwords exist for each appliance: one to access the Manager UI and one to access the CLI when you are connected to the appliance through a telnet or serial connection. By default, the appliance uses `admin` for both the Administrator's ID and password in each case.

For a given Manager UI session, an ID and password are required the first time you access an appliance or the cluster, or when you attempt to connect to a node through a telnet connection. The Administrator's ID and password are unique for each node in the cluster. It is recommended that you change the default

Administrator's ID and password for both telnet and Manager UI access as soon as possible after installing each node.

To change the password for the Manager UI, see *Using the Security page, on page 39*. To change the password for the telnet or serial connection, see *Changing the administrator password for telnet or serial access, on page 60*.

Chapter 3

Monitoring Appliance Performance

This chapter describes how to use the Manager UI to collect and interpret performance statistics on the Intel NetStructure Cache Appliance.

This chapter contains the following sections:

Accessing monitor pages

The Manager UI uses *monitor pages* to present performance information on the selected appliance and the cluster as a whole. A monitor page is a browser page displayed as a result of "clicking" on a page button in the Manager UI. By default, the Manager UI starts in monitor mode (as opposed to configure mode), which displays Monitor page buttons.

▼ Reaching Monitor pages

1. Open your browser to the Manager UI.
2. Enter the Administrator ID and password. By default, the Administrator ID is `admin` and the password is also `admin`. Intel recommends that the administrator change these values when the appliance is initially installed.

Note Should you forget your password, contact Customer Service at Intel Corporation for assistance. For information on how to contact Intel Customer Service, see the *Intel NetStructure Cache Appliance Product Support* booklet that came with your system.

3. Click on a **MONITOR** tab.

Note Some performance displays rely on Java. To use the Monitor pages or any other pages in the UI, make sure your browser is set to enable Java and JavaScript.

Information displayed on the monitor mode pages fall into two categories: information for the selected node in the cluster, and information for the cluster as a whole. To view information on a given node, you need to access that node as described in *Changing the selected node, on page 20*.

Using the Dashboard page

The **Dashboard** page provides a concise view of the appliance and of the cluster. The page displays all nodes in the cluster by name and tracks essential statistics for each node. In the list of nodes, a single node is currently selected. Its name appears in black text *without* underlining, while the rest of the node names appear appear as hypertext links.

▼ Reaching the **Dashboard** page

1. Be sure you are in monitor mode. If not, click the **MONITOR** tab.
2. Click the **Dashboard** page button.

Note By default, the **Dashboard** page appears after you log onto Manager UI with your Administrator ID and password.

Node-specific information With the exception of the information on the **Dashboard** page and the cluster information on the **Node** page, performance information pertains to a single node.

Use the **Dashboard** page to:

- ✔ Select a node
- ✔ See which nodes are on and which are off
- ✔ See if an alarm condition exists on any node

  If an alarm condition exists, you can click the alarm light to view a description of the alarm and resolve it.
- ✔ See the number (cumulative to date) of objects served to users from each node
- ✔ See the traffic load (as current transactions per second)

  The meter dial shows you how hard a node is working. When the needle is to the left on the dial, the work load is light. When the needle is to the far right (red), the node is overloaded.

Dashboard alert lights

The Dashboard contains two alert lights: an on/off light and an alarm light. Alert lights indicate the following about a node:

Alert light	Condition	Description
on/off light	*Green*	*Caching is active.*
on/off light	*Dark*	*Caching is not active.*
alarm light	*Green*	*No alarms.*
alarm light	*Red with link to alarms*	*Alarms exist for that node. Click the red alarm light for more information.*
alarm light	*Yellow*	*A cluster problem exists.*

Resolving alarms

Alarms alert you to problems or warn you of potential problems. Alarm conditions themselves are built into the appliance—you cannot change them.

If an alarm light is on, you can click it to view a description of the alarm conditions. Click the **Resolve** button to acknowledge that you have been informed of the condition.

Important Clicking the **Resolve** button only dismisses alarm messages; it does not actually resolve the cause of the alarms.

Exposing node detail

Click the **More Detail** link to expose the following information for the listed nodes in the cluster:

- ✔ Cache hit rate
- ✔ Cache hit rate, fresh

- ✔ Cache hit rate, refresh
- ✔ Errors
- ✔ Aborts
- ✔ Active clients/servers
- ✔ Average fresh hit

Note Online help provides descriptions for each of these statistics.

Changing the selected node

As mentioned earlier, information on pages accessed in monitor mode exists for the selected node and for the cluster as a whole. You start the process to change the selected node from the **Dashboard** page by clicking on a node name.

▼ Changing the selected node

1. Click on the node name.
2. Provide the Administrator ID and password, if necessary. It is only necessary to log on to a node once during a given Manager UI session.

Note Should you forget your password, contact Customer Service at Intel Corporation for assistance. For information on how to contact Intel Customer Service, see the *Intel NetStructure Cache Appliance Product Support* booklet that came with your system.

After changing the selected node, that name appears as black text *without* underlining, while the remaining node names appear as hypertext links.

If you need more information about the selected node, click the **Node** page button (described in *Using the Node page, on page 20*).

Note The online help provides descriptions of each of the statistics in the **Dashboard** page.

Using the Node page

The **Node** page provides performance statistics for the currently selected node in your cluster and the cluster as a whole. These statistics include document hit rates, DNS lookups, and client and server transactions.

▼ Reaching the **Node** Page

1. Be sure you are in monitor mode. If not, click the **MONITOR** tab.
2. Click the **Node** page button.

Note Online help provides descriptions for each of the statistics on the **Node** page.

Using the Graphs page

The **Graphs** page provides a list of options for generating performance graphs for cache results, garbage collection, transfer rates, and object size for the currently selected node.

▼ Reaching the **Graphs** page

1. Be sure you are in monitor mode. If not, click the **MONITOR** tab.
2. Click the **Graphs** page button.

Once you reach the **Graphs** page, click a link to generate a graph for viewing.

Using the Protocols page

The **Protocols** page provides cluster-wide statistics for use of the HTTP, FTP, NNTP, ICP, and WCCP protocols for the selected node.

▼ Reaching the **Protocols** page

1. Be sure you are in monitor mode. If not, click the **MONITOR** tab.
2. Click the **Protocols** page button.

Note Online help provides descriptions for each of the statistics in the **Protocols** page.

Using the Cache page

The **Cache** page provides cache statistics for the selected node. Cache statistics report cumulative and current information about connections, transactions, object reads and writes, and document hits and misses.

▼ Reaching the **Cache** page

1. Be sure you are in monitor mode. If not, click the **MONITOR** tab.
2. Click the **Cache** page button.

Note Online help provides descriptions of each of the statistics in the **Cache** page.

Using the ARM page

The **ARM** page provides statistics about the Adaptive Redirection Module used for transparent proxy caching for the selected node. The statistics include information about ARM configuration, WCCP fragments (if you are using a WCCP-enabled router), the Network Address Table (NAT), and security (for example, the number of dropped TCP connections).

▼ Reaching the **ARM** page

1 Be sure you are in monitor mode. If not, click the **MONITOR** tab.

2 Click the **Arm** page button.

Note Online help provides descriptions of each of the statistics in the **ARM** page.

Using the Other page

The **Other** page reports statistics for the various appliance functions, including host database and DNS lookups for the selected node.

▼ Reaching the **Other** page

1 Be sure you are in monitor mode. If not, click the **MONITOR** tab.

2 Click the **Other** page button.

Host database and DNS statistics If you see more lookups on the DNS server than in the host database, you might need to increase the size of your database or adjust database time-out settings. Or, you might need to adjust the time-out and retry settings for DNS look-ups. To make adjustments, see *Using the Host Database page, on page 44*.

Note Online help provides descriptions of each of the statistics in the **Other** page.

Using the MRTG page

Multi Router Traffic Grapher (MRTG) is a graphing tool that enables you to monitor the appliance's performance. The **MRTG** page shows information about virtual memory usage, client connections, document hit rates, hit and miss rates, and so on. MRTG uses five-minute intervals to formulate the statistics and provides useful historical information about your appliance's performance.

▼ Reaching the **MRTG** page

1 Be sure you are in monitor mode. If not, click the **MONITOR** tab.

2 Click the **MRTG** page button.

Once the page is displayed, click on a graph to see daily, weekly, monthly, and yearly statistics for that particular graph.

You can also click on the **daily view** link at the bottom of the **MRTG** page to see daily statistics and on the **weekly view** link to see weekly statistics. Clicking on these links provides a more extensive selection of related graphs.

Note Online help provides descriptions of the graphs.

Chapter 4

Configuring the Appliance

This chapter describes the configuration options that control the Intel NetStructure Cache Appliance behavior and performance, and instructs you on how to set these values in the Manager UI.

This chapter contains the following sections:

Accessing configure pages

The Manager UI uses *configure pages* to display and allow configuration changes to the selected appliance. A configure page is a browser page displayed as a result of "clicking" on a configure page button in the Manager UI.

Note Some performance displays rely on Java. To use the configure pages or any other pages in the UI, make sure your browser is set to enable Java and JavaScript.

▼ Reaching the configure pages

1. Open your browser to the Manager UI.
2. Enter the Administrator ID and password. By default, the Administrator ID is `admin` and the password is also `admin`. It is recommended that you change these default values as soon as possible after the appliance is installed.

Note Should you forget your password, contact Customer Service at Intel Corporation for assistance. For information on how to contact Intel Customer Service, see the *Intel NetStructure Cache Appliance Product Support* booklet that came with your system.

3. Click the **CONFIGURE** tab.

   After you click the **CONFIGURE** tab, the **Server Basics** page appears.

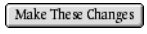

Each configure page allows you to control certain configuration settings for the selected node in a cluster. To update a setting you must provide relevant data or choices and then click the accompanying **Make These Changes** button on the configure page.

The following sections describe each configure page in detail.

Using the Server Basics page

The **Server Basics** page lets you:

- ✔ Turn cache and proxy services on or off
- ✔ Identify the appliance name
- ✔ Restart or reconfigure the caching software
- ✔ Configure the use of virtual IP addresses
- ✔ Auto configure browsers to connect to the appliance
- ✔ Throttle appliance connections
- ✔ Enable SNMP agents

▼ Reaching the **Server Basics** page

- ✔ If you are in monitor mode, click the **CONFIGURE** tab.
- ✔ If you are in configure mode, click **Server** page button.

Setting general options

The following table describes the general configuration settings in the **Intel NetStructure Cache** section.

Option	Description
on/off	*Enables or disables caching. When you disable caching, you shut down all cache and proxy services on a node-by-node basis. That is, you can turn caching on or off only one node at a time.* *You must disable cache services before performing certain maintenance tasks.*
Intel NetStructure Cache Cluster name	*Displays the hostname for the appliance. By default, the name assumes a standalone node and displays the hostname for the appliance as the cluster name. If you are configuring an appliance to be part of an existing management cluster, you must enter the cache cluster name.*
Local Domain Expansion on/off	*Enables or disables local domain expansion.* *If you want the appliance to attempt to resolve unqualified hostnames by expanding to the local domain, enable expansion. For example, if a user makes a request to an unqualified host named* `host_x`*, and if the appliance's local domain is* `y.com`*, the appliance will expand the hostname to* `host_x.y.com`*.*
`.com` *Domain Expansion on/off*	*Enables or disables* `.com` *domain expansion.* *If you want the appliance to attempt to resolve unqualified hostnames by redirecting them to the expanded address prepended with* `www.` *and appended with* `.com`*, enable expansion. For example, if a user makes a request to* `inktomi`*, the appliance redirects the request to* `www.inktomi.com`*.* *If local domain expansion is enabled, the appliance attempts local domain expansion before* `.com` *domain expansion; the appliance tries* `.com` *domain expansion only if local domain expansion fails.*

Setting Web management options

The Web Management section lets you restart the cluster and specify refresh rates as observed in monitor mode. The following table describes these configuration settings.

Option	Description
Restart	*Restarts the entire cluster.* *You must restart the cluster to effect changes you have made to port numbers and virtual IP addresses on the selected node. Restarting the cluster takes about 15 seconds, during which time cache and proxy services are disabled.*
Refresh rate in Monitor mode	*Specifies the refresh rate for the display of the graphs and statistics with which you can monitor the appliance's performance.*

Setting virtual IP addressing options

The **Virtual IP Addressing** section lets you define and maintain the appliance's pool of virtual IP addresses.

The appliance keeps a pool of IP addresses as virtual IP addresses from which to draw and assign IP addresses to nodes as necessary. This practice assures that if a node in the cluster fails, other nodes can assume the failed node's responsibilities.

What are virtual IP addresses?

Virtual IP addresses are really just IP addresses. They are called virtual addresses because they are not tethered to particular machines and can rotate among nodes in a cluster.

It is common for a single machine to represent multiple IP addresses on the same subnet. This machine would have a primary or real IP address bound to its interface card and would also serve many more virtual addresses.

Using virtual IP addressing for node failover

You can set up your user base to use a DNS round-robin pointing at virtual IP addresses, as opposed to using the real IP addresses of the appliance machines in the cluster.

Because virtual IP addresses are not bound to machines, a cluster can steal addresses from inactive nodes and distribute those addresses among the remaining live nodes.

Using a proprietary management protocol, appliance nodes communicate their status with their peers. If a node fails, its peers notice the failure and quickly negotiate which of the remaining nodes will mask the fault by taking over the failed node's virtual interface.

The following table describes the **Virtual IP Addressing** configuration settings.

Option	Description
Virtual IP on/off	*Enables or disables virtual IP addressing.* *If virtual IP addressing is disabled, appliance nodes cannot cover each other's failures.*
Edit virtual IP addresses	*Allows you to edit your list of virtual IP addresses. Changes will not be effective until you click the **Restart** button on the same page.* *Incorrect IP addressing can effectively disable your system. Make sure you understand how virtual IP addresses work before you change them. If you do not assign a range of valid virtual IP addresses to the appliance's manager process, nodes cannot cover each other's failures.*

Adding entries to the Virtual IP address list

You can add or change entries in the Virtual IP address pool by modifying the appliance's Virtual IP address list.

▼ Modifying the Virtual IP address list

1. On the **Server Basics** page, scroll to the **Virtual IP Addressing** section.
2. Click the **Edit virtual IP addresses** link.

   The **Virtual IP** page appears. You can add, remove, or modify Virtual IP addresses by clicking the **Add Entry**, **Delete**, or **Modify** buttons.

▼ Adding a Virtual IP address

1. Click the **Add Entry** button in the **Virtual IP** page.
2. In the **IP Address** field, enter the virtual IP address.
3. In the **Device** field, enter the network interface name (for example, `iprb0`).
4. In the **Subinterface** field, enter the subinterface-ID.

   This is the number between 1-255 that the interface uses for the address.
5. Click the **Add** button.

Note To reset the fields, click the **Reset** button.

Handling multiple interfaces If you have multiple network interfaces, the appliance monitors the state of the interfaces and detects failure. It does this by sending ICMP echo requests, much like the `ping` command.

Setting browser auto configuration options

The **Autoconfiguration of Browsers** section lets you specify an auto configuration file for the selected node. Web browsers use the appliance by specifying a preference to use a proxy server, usually through an auto configuration file.

Note Users must set their browsers to connect to the appliance's auto configuration file. For information on setting your browser to use a proxy, such as the appliance, see your browser documentation. If you are using the transparency option, you do not need auto configuration files.

The following table describes the section's options.

Option	Description
Autoconfiguration file	*Allows you to create or edit an auto configuration file.*

Setting throttling of network connections

The **Throttling of Network Connections** section lets you set a limit on the number of connections the appliance can have. Setting limits on the connections helps to prevent system overload when traffic bottlenecks develop. When network connections reach the limit, new connections are queued until existing connections close.

Note This section is available only if transparency is disabled. If you enable transparency, you do not see this option. See *Configuring load-shedding, on page 28* for information about the transparency load shedding option.

The following table describes the section's options.

Option	Description
Maximum Number of Connections	*Specifies the maximum number of connections that the appliance can have.*

Configuring load-shedding

The **Load Shedding** section lets you configure how the appliance handles overloaded conditions.

When transparency is enabled, the appliance handles overload conditions by forwarding a percentage of new requests to origin servers. You can configure the appliance to automatically shed load if the HTTP-hit transaction times become too long. For example, suppose that the lower limit for HTTP hit-transaction time is 500 milliseconds and the upper limit is 1000 milliseconds. Given these limits, the following is true:

✔ If it takes the appliance more than 500 milliseconds to serve a fresh hit, it begins to shed load.

- ✔ If it takes the appliance more than 750 milliseconds, it begins to shed 50% of its load.
- ✔ If the fresh-hit transaction time exceeds 1000 milliseconds, the appliance begins to shed 100% of its load.

Load shedding is temporary; when hit-transaction times return to acceptable levels, the appliance reverts to handling all incoming requests.

The following table describes the options.

Option	Description
HTTP hit transaction time - low watermark	*The lower limit for HTTP transaction time in milliseconds.* *When the average hit transaction time reaches this value, the appliance forwards a percentage of incoming client requests directly to the origin server.*
HTTP hit transaction time - high watermark	*The upper limit for HTTP transaction time in milliseconds.* *When the average hit transaction time reaches this value, the appliance forwards all incoming client requests directly to the origin server.*

Enabling SNMP agents

The **SNMP** section lets you enable an SNMP agent to monitor information about the appliance and send warning messages, called SNMP traps, to SNMP monitoring stations.

The following table describes the options.

Option	Description
SNMP Agent on/off	*Enables or disables an SNMP agent.* *The appliance SNMP agent supports access to two management information bases (MIBs): MIB-2 (a standard MIB) and the Intel NetStructure Cache Appliance MIB. Enabling the SNMP agent on allows access to both.*

Using the Protocols page

The **Protocols** page lets you view and change the selected appliance's protocol configuration. You can tune HTTP, NNTP, and FTP timeout intervals; and configure the appliance to remove HTTP headers from documents to protect site and user privacy.

▼ Reaching the **Protocols** page

1. Be sure you are in configure mode. If not, click the **CONFIGURE** tab.
2. Click the **Protocols** page button.

The **Protocols** page is divided into four sections for configuring HTTP, NNTP, HTTPS, and FTP.

Configuring HTTP

The HTTP section lets you configure the appliance's handling of HTTP. The following table describes the section's options.

Option	Definition
Keep-Alive Timeout Inbound	*Specifies how long the appliance should keep connections to users open for a subsequent request after a transaction ends.* *If the user does not make another request before the timeout expires, the appliance closes the connection. If the user does make another request, the timeout period starts over.* *The user can close the connection at any time.*
Keep-Alive Timeout Outbound	*Specifies how long the appliance should keep open the connections to Web servers (content servers) for a subsequent transfer of data after a transaction ends.* *If the appliance does not need to make a subsequent request for data before the timeout expires, it closes the connection. Once the connection is closed, the timeout period starts over.* *The Web server can close the connection at any time.*
Inactivity Timeout Inbound	*Specifies how long the appliance should keep connections to users open if a transaction stalls. If the appliance stops receiving data from a user or the user stops reading the data, the appliance closes the connection when this timeout expires.* *The user can close the connection at any time.*
Inactivity Timeout Outbound	*Specifies how long the appliance should keep open connections to Web servers if the transaction stalls. If the appliance stops receiving data from a Web server, the appliance will not close the connection until this timeout has expired.* *The Web server can close the connection at any time.*

Option	Definition (Continued)
Activity timeout Inbound	*Specifies the maximum time the appliance should remain connected to a user. If the user does not finish making a request (reading and writing data) before this timeout expires, the appliance closes the connection.* *The user can close the connection at any time.*
Activity Timeout Outbound	*Specifies the maximum time the appliance should wait for fulfillment of a connection request to a Web server. If the appliance does not establish a connection to a Web server before this timeout expires, the appliance terminates the connection request.* *The Web server can close the connection at any time.*
Remove the following common headers	*Specifies headers for removal. Removing headers can protect the privacy of your site:* ▮ The `From` header. This header identifies the user's e-mail address ▮ The `Referer` header. This header identifies the Web link followed by the user. ▮ The `User-Agent` header. This header identifies the agent—usually a browser—making the request. ▮ The `Cookie` header. This header is often used to identify the user making a request.
Insert Client-ip	*Insert* `Client-ip` *headers to retain client IP addresses.*
Remove Client-ip	*Remove* `Client-ip` *headers for more privacy.*
User Language	*Selects the language in which messages to the user from the appliance are displayed. The default language is English.*

Configuring NNTP

The **NNTP** section lets you configure basic NNTP options. While this section lets you configure basic options, you must use the command-line interface to configure the appliance to cache articles from particular NNTP servers and news groups as well as to set access restrictions and authentication requirements for news readers. See *Configuring NNTP servers, on page 65* for more information.

The following table describes the options.

Option	Definition
NNTP Server on/off	*Enables or disables the appliance to cache and serve news articles.* *After turning NNTP on or off for the selected node, you must restart the cluster to effect the change. Click the **Restart** button on the **Server Basics** page.*
NNTP Server Port	*Specifies the port that the appliance uses for serving NNTP requests. The default port is 119.*
Connect Message (posting allowed)	*Defines the message displayed to news readers when they connect to the appliance with posting allowed.*
Connect Message (posting not allowed)	*Defines the message displayed to news readers when they connect to the appliance with posting not allowed.*
NNTP options	▮ Posting: Allows users to post NNTP articles to parent NNTP servers. ▮ Access Control: Turns access control on or off. To refine access control, use the command-line interface. See *Configuring NNTP access, on page 69* for more information. *If you are using an authentication server, you must enter its name and port (see page 33).* ▮ NNTP V2 Authentication Server: Supports NNTP version 2 authentication. Use this option only if all of your client authentication supports version 2. ▮ Run Local Authentication Server: Runs an authentication program on the selected node. Use the command-line interface to configure which clients must be authenticated. See *Configuring NNTP access, on page 69* for more information. ▮ Allow Feeds: Allows the appliance to accept feeds of news articles from feed or push groups. *Use the command-line interface to designate feed and push groups. The appliance does not cache news articles from feed groups; instead, it receives feeds of news articles as the parent NNTP server receives feeds. Push groups are groups for which the appliance can both retrieve articles on demand and receive news feeds.* *See Configuring NNTP servers, on page 65 for information about designating news groups as push or feed.*

Option	Definition (Continued)
NNTP options (continued)	▮ Background Posting: Causes the appliance to post NNTP articles to parent NNTP servers in the background. ▮ Obey Cancel Control Messages: Sets the appliance to obey `cancel` control messages. *When the appliance gets a cancel control message, it deletes the corresponding article from the cache. You do not need to enable this option if the appliance is caching articles on demand (i.e. no feed groups exist). For all nonfeed news groups, the appliance actively polls parent NNTP servers for cancelled articles. See the Check for Cancelled Articles option, below.* ▮ Obey Newgroups Control Messages: Causes the appliance to obey `newgroup` control messages. *The appliance actively polls parent NNTP servers for new groups; see the Check for New Groups option, below.* ▮ Obey Rmgroups Control Messages: Sets the appliance to obey `rmgroup` (remove group) control messages.
Inactivity Timeout	*Defines the number of seconds that idle connections can remain open. This timeout should be at least three minutes.*
Check for New Groups Every	*Defines the number of seconds that pass before the appliance polls parent NNTP servers for new news groups. The parent group lists change slowly. Consequently, you do not need to check them frequently.* *Use the command-line interface to list the hosts you want the appliance to poll. See Configuring NNTP servers, on page 65 for more information.*
Check for Cancelled Articles Every	*Defines the number of seconds that pass before the appliance polls all nonfeed news groups on the parent NNTP servers for cancelled articles. Checking for new articles should not be done too frequently as it involves communication with the parent NNTP server.*
Check Parent NNTP Server Every	*Defines the number of seconds that pass before the appliance polls the parent NNTP server for new articles.*
Check Cluster Every	*Defines the number of seconds that pass before the appliance checks the nodes on the cluster.*
Check Pull Groups Every	*Defines the number of seconds that pass before the appliance pulls (or caches) news articles from defined pull groups. Use the command-line interface to designate pull groups. See Configuring NNTP servers, on page 65 for more information.*
Authentication Server Host	*The name of the host machine running the authentication server. If the host machine is the appliance, enter "local host".*

Option	Definition (Continued)
Authentication Server Port	*The port on which the locally run authentication server accepts connections. If the authentication server is remote, the appliance connects to it on this port.*
Local Authentication Server Timeout	*The number of milliseconds after which the authentication server aborts an incomplete authorization operation. The client can retry the operation.* *Refer to Configuring NNTP access, on page 69 for information about configuring authentication servers.*
Client Speed Throttle	*The number of bytes per second that clients are limited to during downloading operations. Use a* `0` *(zero) for unlimited downloading.*

Configuring FTP

The **FTP** section lets you configure FTP protocols. The following table describes the options.

Option	Definition
FTP connection mode	▮ PASV/PORT: Specifies the appliance use PASV connection mode. PASV/PORT is the default FTP connection mode. If PASV mode fails, the appliance uses PORT mode to initiate the data connection, and then the appliance accepts it. ▮ PASV only: Specifies that the appliance initiates the data connection to the FTP server, and the FTP server accepts it. This mode is firewall-friendly, however, some FTP servers do not support it. ▮ PORT only: Specifies that the FTP server initiates the data connection, and the appliance accepts it. *FTP transfers require two connections: a control connection to inform the FTP server of a request for data and a data connection to send the data. The appliance always initiates the control connection. FTP mode determines whether the appliance or the FTP server initiates the data connection.*
FTP inactivity timeout (seconds)	*Defines the number of seconds before the appliance waits for a response from the FTP server. If the FTP server does not respond in time, the appliance abandons the user's request.*
Anonymous FTP password	*Specifies an anonymous password for FTP servers that require a password for access.*

Using the Cache page

The **Cache** page allows you to configure the following:

- ✔ Cache activation
- ✔ Object freshness
- ✔ Variable object content

▼ Reaching the **Cache** page

1. Be sure you are in configure mode. If not, click the **CONFIGURE** tab.
2. Click the **Cache** page button.

The following sections describe the sections in the **Cache** page.

Cache activation

The following table describes the cache activation configuration options.

Option	Description
Enable HTTP caching	*Enables caching of objects retrieved through HTTP.*
Enable FTP caching	*Enables caching of objects retrieved through FTP.*
Enable NNTP caching	*Enables caching of objects retrieved through NNTP.*
Ignore user requests to bypass cache	*Instructs the appliance to ignore* `no-cache` *headers. This means the appliance ignores a user's stipulation to ignore their requests served from the cache.*

Storage

The following table describes the storage options.

Option	Description
Maximum HTTP/FTP object size in bytes	*Specifies the maximum size of HTTP or FTP objects the appliance can cache.* *Use a `0` (zero) to indicate no limit.*
Maximum number of alternate versions (HTTP)	*Specifies the maximum number of HTTP alternates that the appliance can cache.* *Use a `0` (zero) to indicate no limit. If a popular URL has thousands of alternates, you might observe increased cache hit latencies (transaction times) as the appliance searches through the alternates for each request. In particular, some URL addresses can have large numbers of alternates due to cookies. If the appliance is set to vary on cookies, you might encounter this problem. See Variable content, on page 38 for more information.*

Freshness

The following table describes the freshness options.

Option	Description
Verify freshness by checking	*Configures the appliance to ask the original content server to verify the freshness of objects according to the following list before serving them.* ▮ when the object has expired ▮ when the object has expired or if the object has no expiration date ▮ always ▮ never
Minimum freshness information for a document to be cached	*Specifies the minimum freshness information required to consider a document able to be cached:* ▮ an explicit lifetime ▮ a last-modified time ▮ nothing
If an object has no expiration date	*Specifies the time limits the appliance will keep an object in the cache:* ▮ minimum time in the cache. You can specify from 15 minutes to two weeks. ▮ maximum time in the cache. You can specify from 15 minutes to two weeks.

Option	Description (Continued)
FTP cached objects expire	*Specifies how long the appliance will keep FTP objects in the cache. You can specify from 15 minutes to two weeks.*
Internet Explorer requests force a check with the origin server	*Configures the appliance to treat Microsoft Internet Explorer requests more conservatively, providing fresher content at the cost of serving fewer objects from the cache according to the following options:* ▮ never: never force a freshness check with the origin server ▮ for IMS revalidation requests: only force a freshness check for IMS (If Modified Since) revalidation requests ▮ always: always force a freshness check with the origin server *Certain versions of Microsoft Internet Explorer do not request cache reloads from reverse proxies and transparent caches when the user presses the browser **Refresh** button. This can prevent content from being loaded directly from the origin servers.*

Variable content

The following table describes the variable configuration options.

Option	Description
Do not cache	*Instructs the appliance to refuse to cache objects served in response to URL addresses that contain:* ▮ ? ▮ ; ▮ `cgi` ▮ end in `.asp`
Enable Alternates	*Instructs the appliance to cache alternate versions of HTTP documents.*
Vary on these HTTP header fields:	*Enables the appliance to serve alternate documents. Selecting the Enable Alternates option allows you to specify values to match for the following fields:* ▮ If the request is for text: The default value is `user-agent` and `cookie`. Some documents can have thousands of alternate cookie versions. If you choose to vary on cookies, it is recommended that you limit the number of alternates cached. See *Storage, on page 36*. ▮ If the request is for images: Images are rarely personalized. ▮ If the request is for anything other than text or images *Using document header information, the appliance can compare cached document specifications against requested specifications to determine if the correct alternate version of the document is in the cache. For example, a document header can specify that the document targets a specific browser, but any browser can request the document from the appliance. If a requested document's fields do not match a cached document's fields, the appliance does not serve the document from its cache, but instead retrieves a fresh copy from the origin server.*
Cache responses to requests containing Cookies for:	*Configures the appliance to cache responses to requests that contain cookies for:* ▮ no content-types ▮ all content-types ▮ only image-content types ▮ content-types that are not text

Using the Security page

The **Security** page lets you configure access to the Manager UI. You can set administrator and guest IDs and passwords (guests have read-only access) for the selected node.

▼ Reaching the **Security** page

1 Be sure you are in configure mode. If not, click the **CONFIGURE** tab.

2 Click the **Security** page button.

The following table describes the Manager access options.

Option	Description
Authentication (basic) on/off	*Enables or disables authentication. Leave authentication on to check the administrator ID and password whenever a user logs on to the Manager UI.*
Administrator's ID	*Specifies the administrator login ID. (The ID is not checked if you turn authentication off.) The administrator has access to both configure and monitor pages in the Manager UI.*
Change administrator's password	*Allows you to change the administrator password. Clicking the link displays the **Change Administrator's Password** page where you can enter and validate a new password. (The password is not checked if you turn authentication off).*
Guest ID	*Specifies the guest login ID. Guests can access only the monitor pages of the Manager UI. The guest login ID is static for all guests.*
Change guest password	*Allows you to change the guest password. Clicking the link displays the **Change Guest's Password** page where you can enter and validate a new password.*

Using the Routing page

The **Routing** page lets you configure the following:

✔ HTTP parent caching

✔ Internet Caching Protocol (ICP) support

✔ Server acceleration (reverse proxy service)

From the **Routing** page, you can also check if transparency and WCCP are enabled.

▼ Reaching the **Routing** page

1. Be sure you are in configure mode. If not, click the **CONFIGURE** tab.
2. Click the **Routing** page button.

Setting HTTP parent caching options

The appliance can participate as a member of an HTTP cache hierarchy. You can point your appliance at a parent network cache—either another Intel NetStructure Cache Appliance or a different caching product—to form a cache hierarchy, wherein a child cache relies upon a parent cache in fulfilling user requests.

parent failover

You can specify more than one parent cache to be queried. If the first parent cache does not respond to the request, the appliance tries the next parent cache.

The appliance supports multiple parent caches and parent failover. Use the command-line interface to configure multiple parent caches and parent failover (which gives appliance a sequence of parent caches to query if the first parent cache misses). See *Controlling parent proxy caching, on page 89*.

The following table describes the options.

Option	Description
Parent Caching on/off	*Enables or disables parent caching. To set parent caching on, you must also name a parent cache.*
Parent Cache	*Identifies the parent cache and port. Using the following format:* `parent_name:port_number`. *The port must be dedicated. If the appliance cannot find a requested object in its own cache, it searches the parent cache before searching the Internet. If you want parent failover, you can specify more than one parent cache; for example,* `parent1:port1; parent2:port2`

Setting ICP options

In the **ICP** section you can establish ICP peers.

The following table describes the ICP options.

Option	Description
ICP Mode	*Enables or disables ICP mode:* ▮ Only Receive Queries ▮ Send/Receive Queries ▮ Disabled
ICP Port	*Specifies the port to use for ICP messages. The default port is 3130.*
ICP Multicast enabled on/off	*Enables or disables multicast. If your appliance has a multicast channel connection to its ICP peers, it can send ICP messages through multicast.*
ICP query timeout	*Specifies the timeout for ICP queries in seconds.*
ICP Peers	*View or modify the appliance's ICP hierarchy.*

Establishing ICP peers

For ICP to work, the appliance must recognize its ICP neighbors (siblings and parents).

▼ Adding an ICP Peer

1. Click the **ICP Peers** link.
2. Click the **Add Entry** button.
3. Enter the information for the ICP peer host. If you want to clear the entire form of information, you can press the **Reset** button.

Field	Description
Hostname	*The hostname for the ICP host. You do not have to enter a hostname if you know the host IP address.* *If you enter a hostname but leave the IP address as 0.0.0.0, the ICP configuration obtains the host IP address via a DNS lookup on the entered hostname. Therefore, if you do not know the IP address, simply leave it as 0.0.0.0.*
Host IP	*The host IP address.* *If you enter an IP address other than 0.0.0.0, the ICP configuration uses the IP address to identify the host. Otherwise, the ICP configuration requires a hostname.*
Type	*ICP host type. Use one of the following options:* ▮ 1 specifies a parent cache ▮ 2 specifies a sibling cache ▮ 3 specifies the local host *Option 3 is reserved for the appliance. In option 3, the hostname must be* `localhost` *and the host IP address must be 0.0.0.0. The ICP configuration enforces this convention.*
Proxy Port	*The appliance's proxy port (usually 8080).*
ICP Port	*The UDP port used for ICP (usually 3130).*
Multicast Member	*Indicates whether the host is on a multicast network with the appliance. Use one of the following options:* ▮ No ▮ Yes
Multicast IP	*The multicast IP address.*
Multicast TTL	*The multicast datagram time to live. Use one of the following options:* ▮ 1: specifies that IP multicast datagrams will not be forwarded beyond a single subnetwork. ▮ 2: allows delivery of IP multicast datagrams to more than one subnet if there are one or more multicast routers attached to the first hop subnet.

4. Click the **Add** button to save your changes.

Setting server accelerator options

The **Server Accelerator** section allows you to configure the appliance as a Server Accelerator (also known as a reverse or server-side proxy). You can enable or disable this function as well as control how the appliance routes document requests to the slower traditional Web servers. For more information about setting up the appliance as a Server Accelerator, see *Setting general controls, on page 62*.

The following table describes Server Accelerator options.

Option	Description
Server Acceleration	*Enables or disables server acceleration.* *If you select on, the appliance is a server accelerator for the Web servers specified in document route rewrite rules defined through the command-line interface.*
Reverse proxy only	*Sets the appliance to operate solely as a server accelerator. If you select Yes, the appliance does not serve requests to unspecified Web servers from the cache. See Understanding server acceleration mapping rules, on page 132 for information on creating document route rewriting rules.* *If you select No, the appliance serves requests from unspecified Web servers as a normal proxy cache.*
Document Route Rewriting Rules	*Allows you to view, modify, or add document route rewrite rules. See Understanding server acceleration mapping rules, on page 132 for information on document route rewrite rules.*
URL to redirect requests without Host header	*Specifies an alternate URL that incoming requests from older clients that do not provide a* `Host:` *header can be directed.* *It is recommended that you set this option to a page that explains the situation to the user and advises a browser upgrade or provides a link directly to the origin server, bypassing the appliance. Alternatively, you can specify a map rule that maps requests without* `Host:` *headers to a particular server.*

▼ Creating a document route rewriting rule

1 In the **Server Accelerator** section, click the **Document Route Rewriting Rules** link.

The **Configure: Routing: URL Rewriting** page appears. This page displays the set of current rules as well as a **Add Entry** button that lets you create new rules.

2 Click the **Add Entry** button.

3 From the **Type** field, select the type of rule you want to set (`map` or `reverse_map`).

4 In the **Target** field, enter the origin or *from* URL for the rule. You can enter up to four components; for example, `<scheme>://<host>:<port>/<path_prefix>`

5 In the **Replacement** field, enter the destination or *to* URL for the rule. You can enter up to four components; for example, `<scheme>://<host>:<port>/<path_prefix>`

6 Click the **Add** button to add the rule.

Note You can abandon the new rule by clicking **Reset**.

Checking transparency

The **Transparency** section indicates whether the appliance is running transparently. If transparency is enabled, you will see the following message:
`The transparency option is installed. Redirected users will be served transparently.`

If transparency is not enabled, you will see the following message:
`The Transparency option is not currently installed.`

For more information about Transparency, see *Transparent proxy caching*, *on page 120*.

Checking WCCP

The **WCCP** section indicates whether WCCP is enabled. If WCCP is enabled, you will see the following message:
`The WCCP option is currently installed.`

If WCCP is not enabled, you will see the following message:
`The WCCP option is not currently installed.`

Using the Host Database page

The **Host Database** page lets you view and change the following:

✔ Host database options

✔ Domain Name Service lookups

▼ Reaching the **Host Database** page

1 Be sure you are in configure mode. If not, click the **CONFIGURE** tab.

2 Click the **Host DB** page button.

Configuring the host database

The appliance host database stores the domain name server (DNS) entries of servers that the appliance contacts to fulfill user requests. You configure the appliance host database by setting options in the **Host Database Management** section. The following table describes the options.

Option	Description
Lookup timeout	*Specifies the DNS lookup timeout in seconds. You can choose from the following:* ▮ 5 seconds ▮ 10 seconds ▮ 15 seconds ▮ 20 seconds ▮ 30 seconds
Foreground timeout	*Specifies how long DNS entries can remain in the database before they are flagged as stale. You can choose from the following:* ▮ 12 hours ▮ 24 hours ▮ 48 hours *For example, if this timeout is 24 hours, and a user requests an entry that has been in the database for 24 hours or longer, the appliance will refresh the entry before serving it.* *You can set the background timeout (see next item) to refresh entries in the background, before objects become stale.* *Be careful not to set the foreground timeout too low. Doing so might slow response time. Additionally, setting the timeout value too high risks accumulation of incorrect information. Setting the foreground timeout to greater than or equal to the background timeout disables background refresh.*

Option	Description (Continued)
Background timeout	*Specifies how long DNS entries can remain in the database before they are flagged as entries to refresh in the background. These entries are still fresh, so they can be refreshed after they are served, rather than before. You can choose from the following:* ▮ 3 hours ▮ 6 hours ▮ 12 hours ▮ 24 hours ▮ 48 hours *For example, the foreground refresh timeout interval is 24 hours and the background timeout is 12 hours. In this situation a user requests an object from* `my.com` *and 16 hours later a user makes a second request for an object from* `my.com`*. The DNS entry for* `my.com` *has not been refreshed in the foreground because the entry is not yet 24 hours old. But since the background timeout has expired, the appliance will first serve the user's request and then refresh the entry in the background.*
Invalid host timeout	*Specifies how long the proxy software should remember that a hostname is invalid. This is often called negative DNS caching. You can choose from the following:* ▮ Immediate ▮ 15 minutes ▮ 30 minutes ▮ 1 hour ▮ 1.5 hours ▮ 2 hours *For example, if a user specifies an invalid hostname, the appliance informs the user that it could not resolve the hostname and the appliance gets another request for the same hostname. If the appliance still remembers the bad hostname, it will not try to look it up again but will simply send another* `invalid hostname` *message to the user.*
Re-DNS on Reload	*Enables or disables the appliance's ability to re-resolve hostnames whenever clients reload pages.*

Configuring DNS

The **DNS Configuration** section lets you configure DNS services. The following table describes the options.

Option	Description
Resolve attempt timeout	*Specifies how long the appliance must wait for the DNS server to respond with an IP address, even if the client request has been cancelled. You can choose from the following:* ▮ 5 seconds ▮ 10 seconds ▮ 15 seconds ▮ 20 seconds ▮ 30 seconds *If the user abandons the request before this timeout expires, the appliance can still obtain the host's IP address in order to cache it. The next time a user makes the same request, the address will be in the cache.*
Number of retries	*Specifies how many times the appliance should allow a lookup to fail before it abandons the lookup and sends an* `invalid hostname` *message to the user. You can choose from the following:* ▮ 1 ▮ 2 ▮ 3 ▮ 4 ▮ 5

Using the Snapshots page

The **Snapshots** page lets you take snapshots of the selected appliance's configurations or lets you restore previously saved configurations. A configuration snapshot consists of a complete set of appliance configuration files.

Note It is a good idea to take a snapshot before doing system maintenance or attempting to tune system performance. Taking a snapshot only takes a few seconds and it can save you hours of correcting configuration mistakes.

▼ Reaching the **Snapshots page**

1 Be sure you are in configure mode. If not, click the **CONFIGURE** tab.

2 Click the **Snapshots** page button.

The following table describes the options.

Option	Description
Name New Snapshot	*Specifies a name for the snapshot. Do not include the forward slash "/" character in the name.*
Take Snapshot	*Takes a snapshot. Taking a snapshop saves a copy of all appliance configuration files. The snapshot is saved under the name specified in the Name New Snapshot field.*
Restore Snapshot	*Restores a snapshot. Clicking the **Restore** button returns the appliance to the configuration previously saved in the snapshot selected from the list.*
Delete Snapshot	*Deletes an existing snapshot. Clicking the **Delete Snapshot** button deletes the previously saved configuration that is selected from the list.*

Note Once you create a snapshot for the appliance, you should remove the floppy diskette from the drive. If you do not remove the diskette from the drive and the system needs to be rebooted remotely, the system will attempt to reboot from the diskette, which does not have a bootable image.

Chapter 5

Using the Command-Line Interface

This chapter describes the command-line utility that you can use to configure the system's network addresses and to control, configure, and monitor the Intel NetStructure Cache Appliance.

This chapter contains the following sections:

Starting the command-line interface

The command-line interface displays automatically on screen when you provide a serial interface connection to the appliance. For information on how to make a serial connection to the appliance, see the *Intel NetStructure Cache Appliance Quick Start Guide*.

Note Make sure your terminal is set to emulate a VT100 terminal when you are communicating with the appliance through a serial interface.

Starting the appliance the first time

The first time you connect to the appliance, the Initial Setup menus display as follows:

```
setup     Initial Intel Cache Setup
install   Install Intel Cache
commit    Commit Setup Changes
```

These menu selections let you do the following:

- ✔ **setup**—Provide the appliance machine with a hostname, IP address, subnet mask address, DNS address, gateway address, domain name, time zone, and date and time.
- ✔ **install**—Install or update the appliance software. This task takes several minutes.
- ✔ **commit**—Save the appliance network configuration after installing the software.

For instructions on how to start the appliance for the first time, see either the *Intel NetStructure Cache Appliance Quick Start Guide* or *Starting the system for the first time, on page 8*.

Note For security reasons, you should change your Administrator ID and password for telnet access as soon as possible after installing and initially configuring your appliance. See *Changing the administrator password for telnet or serial access, on page 60*.

Using the appliance after initial start-up

After initial configuration and when you connect to the appliance through a serial interface, this main selection menu displays on the screen:

```
setup     Initial Intel Cache Setup
main      Main Intel Cache Controls
config    Intel Cache Configuration
monitor   View Statistics
expert    Enter Expert Mode
save      Save Config to Floppy
load      Load Config From Floppy
logoff    Logoff
```

These menu selections let you do the following:

- ✔ **setup**—Change the system's network address configuration and time settings. See *Using the setup menu, on page 52* for more information.
- ✔ **main**—Start or stop the cache and proxy services, check version information, clear statistics, and install and delete software. See *Using the main menu, on page 54* for more information.
- ✔ **config**—Configure the appliance, including server, protocols, security, and routing. See *Using the config menu, on page 61* for more information.
- ✔ **monitor**—Monitor performance by viewing statistics. See *Using the monitor menu, on page 99* for more information.
- ✔ **expert**—Use the appliance's expert feature. See *Using the expert menu, on page 107* for more information.
- ✔ **save**—Save the current configuration to a floppy disk. See *Using the save menu, on page 108* for more information.
- ✔ **load**—Load a saved configuration from a floppy disk. See *Using the load menu, on page 108* for more information.
- ✔ **logoff**—Logoff from the current login.

Navigating the command-line interface

The command-line interface consists of a series of menus that you can access to adjust the system's network configuration and control, and to configure and monitor the appliance.

The following table explains how to navigate the interface:

To do this...	Do this
Move from one menu item to another	*Use the up and down arrow keys*
Select a menu or menu item	*Move to the item and press Enter*
Return to the previous form or menu screen	*Press CTRL-X*
Accept an action confirmation box	*Press CTRL-X*
Accept changes to the form and exit it by returning to the previous form or menu screen	*Press CTRL-X*
Save information you have entered in a form's field and position the cursor at the next field. You must press Enter for each field in the form	*Press Enter*
Cancel all changes to a form and exit it by returning to the previous form or menu screen	*Press ESC*

As you navigate through windows, you see the path of the window displayed in the top menu border, starting with the root menu.

The following steps provide an example of how to view cache performance statistics from the **monitor** menu.

1. From the initial menu, use the down arrow key on your keyboard to navigate to the **monitor** menu item. Doing so highlights that item to show that you have selected it.
2. Press Enter. After pressing Enter, the **monitor** menu appears and the menu border displays `root->monitor`.
3. Press the down arrow key to navigate to the **cache** menu item and press Enter. Doing so displays the cache performance statistics on the screen and the menu border displays `root->monitor->cache`.

Using the setup menu

The **setup** menu lets you do the following:

- ✔ Change the IP address, hostname, and netmask address on the primary network interface controller in the appliance.
- ✔ Change the speed and transmission mode of the primary network interface controller.
- ✔ Change the DNS address and domain name.
- ✔ Change the gateway address.
- ✔ Configure time zone settings.
- ✔ Configure date and time settings.
- ✔ View current network address settings on the primary network interface controller.

Changing network addresses configuration

You can change the network settings of the primary network interface controller (host name, IP address, and netmask address) any time after the initial setup.

Note You must configure the network interface controller the first time you connect to the appliance from a terminal. (See *Starting the command-line interface, on page 50* for more information.)

▼ Changing network address configuration on the NIC

1. Select the **setup** menu and press Enter.
2. Select **ip** and press Enter. Doing so displays the current IP address, hostname, and netmask.
3. In the **New IP Address** field, enter the IP address that you want to assign to the appliance, and press Enter.

4 In the **New Hostname** field, enter the hostname that you want to assign to the appliance, and press Enter.

5 In the **New Netmask** field, enter the netmask address that you want to assign to the appliance system, and press Enter.

6 Press CTRL-X to save your changes and return to the previous menu.

Changing the controller speed and transmission mode

You can change the speed and transmission mode of the primary network interface controller any time after the initial setup.

▼ Changing speed and transmission mode

1 Select the **setup** menu and press Enter.

2 Select **nic** and press Enter.

3 From the list, choose a speed and mode and press Enter. Doing so causes a message to appear indicating the change has been made but will not take effect until the system is rebooted.

Changing the DNS address and domain name

You can change the DNS address and domain name used by the appliance.

▼ Changing the DNS address

1 Select the **setup** menu, and press Enter.

2 Select **dns** and, press Enter. Doing so displays the current DNS address and domain name.

3 In the **New DNS Address** field, enter the DNS address that you want to assign to the appliance, and press Enter.

4 In the **New Domainname** field, enter the domain name that you want to assign to the appliance, and press Enter.

5 Press CTRL-X to save your changes and return to the previous screen.

Changing the gateway address

You can change the gateway address used by the appliance.

▼ Changing the gateway address

1 Select the **setup** menu, and press Enter.

2 Select **gateway,** and press Enter. Doing so displays the current gateway address and a field in which you can enter the new gateway address.

3 In the **New Gateway** field, enter the gateway address that you want to assign to the appliance, and press Enter.

4 Press CTRL-X to save your changes and return to the previous screen.

Configuring time zone settings

You can configure the appliance for the appropriate time zone.

▼ Configuring the time zone setting

1. Select the **setup** menu, and press Enter.
2. Select **timezone,** and press Enter. Doing so displays a list of available time zone settings.
3. Use the up and down arrow keys to scroll through the list and select the appropriate time zone.
4. Once you have selected the item, press Enter.
5. Press any key to continue.
6. Press CTRL-X to return to the previous screen. When you exit the screen, a message appears indicating that the new time zone setting does not take effect until the system is rebooted.

Configuring date and time settings

You can configure the appliance's date and time.

▼ Configuring the date and time settings

1. Select the **setup** menu, and press Enter.
2. Select **time,** and press Enter. Doing so displays time and date fields, each having various fields in which you can enter data.
3. Provide data in each sub-field and use the Enter key to move between sub-fields**.**
   - ✔ Enable or disable Daylight Savings Time
   - ✔ Indicate whether you're inside or outside Daylight Savings Time
   - ✔ Enter time in the format HH:MM:SS
   - ✔ Enter the date in the format MM/DD/YYYY
4. When you have finished, press CTRL-X to confirm your settings and exit the window.

Viewing current network address settings

You can view the current hostname, IP, DNS, and Gateway address settings by selecting **view** from the **setup** menu.

Using the main menu

The **main** menu lets you do the following:

✔ Check the status of the *Server* and *Manager* resident on the appliance.

- ✔ Start the appliance cache and proxy services.
- ✔ Stop the appliance cache and proxy services.
- ✔ View and maintain the version of software installed on the appliance.
- ✔ Clear persistent statistics.
- ✔ Reboot the system.
- ✔ Halt the system.
- ✔ Change Administrator password for telnet and serial access.
- ✔ Reset the appliance to the factory settings.
- ✔ Prepare cache disk.

Checking the status of the Server and Manager

You can check the status of the appliance's Server and Manager applications using the **main** menu.

▼ Checking Server and Manager status

1. Select the **main** menu, and press Enter.
2. Select **status**, and press Enter. Doing so displays a window that indicates whether the Server and Manager are **`UP`** or **`DOWN`**.

Starting the appliance

Starting the caching and proxy services "starts" the appliance.

▼ Starting the appliance

1. Select the **main** menu, and press Enter.
2. Select **start**, and press Enter. Doing so displays a message indicating that the appliance has started successfully.

Stopping the appliance

Shutting down all caching and proxy services "stops" the appliance.

Note You must stop the appliance before doing certain maintenance tasks.

▼ Stopping the appliance

1. Select the **main** menu, and press Enter.
2. Select **stop**, and press Enter. Doing so displays a message indicating the cache has been stopped.

Viewing and maintaining versions of the software

You can have up to two versions of the appliance software installed on the system at the same time. From these versions, you can choose which one is current and executes in the appliance. Installing a new version of the software automatically makes it the current version.

You can use the **versions** menu, which is a submenu of the **main** menu, to do the following:

- ✔ Identify the installed versions.
- ✔ Install new versions.
- ✔ Switch versions.
- ✔ Delete a version.
- ✔ View which version is running.

Identifying which versions of the software are currently installed

▼ Identifying which versions of the appliance software are installed.

1. Select the **main** menu and press Enter.
2. Select **versions** and press Enter.
3. Select **view** and press Enter. Doing so displays a list of version numbers.

Installing a new version of the appliance software

You can update the software on your cache appliance using FTP to download the updated files. When you install a new version of the software, it becomes the current, running version. In addition, the appliance copies the new version to your secondary drive.

▼ Setting up the FTP server

1. Set up the FTP server to provide upgrade files to the appliance. You can use a single FTP server to upgrade multiple appliances.
2. Place the files on an FTP server that's accessible by the appliance, and on a network with sufficient performance for fast transfer of files.
3. Each upgrade must exist in a separate directory. We recommend that the names you choose for your directories indicate the release. This example shows separate directories for application, patch, and OS/application upgrades:

```
<ftp_dir>/app_3.0.9.0
<ftp_dir>/app_3.1.0.0
<ftp_dir>/patch_1
<ftp_dir>/patch_2
<ftp_dir>/os_1
<ftp_dir>/os_2
```

4 Regardless of the type of upgrade, that is, application, patch, or OS/application, each upgrade requires two files, which you must copy into the correct directory on the FTP server:
`upgrade_info`
`<upgrade_name>.tar.gz`

▼ Starting the upgrade from the appliance side

1 Start the command line interface.

2 Go to `root > main > version > install`

3 Enter the following information in the fields provided:

- ✔ IP address or hostname of FTP server
- ✔ Path to upgrade files
- ✔ Username on FTP server
- ✔ User password on FTP server

4 Press **Ctrl-X** to begin upgrading. A message will appear, `Checking FTP Site...` as the appliance connects to the FTP server and retrieves the `upgrade_info` file. Next, the CLI displays the type of upgrade (APP, PATCH, or OS), and a message describing the upgrade. You will see a warning that an OS upgrade later requires you to swap the primary and secondary drives.

5 Press **Ctrl-X** to proceed or **Escape** to abort. If you select Proceed, the upgrade continues, following the procedure for that upgrade type as explained in the corresponding section below.

Application upgrade

After you press Ctrl-X to proceed, the CLI displays this message:
`Ftp'ing Application Upgrade. Please Wait...`

The server transfers the application upgrade file `tar.gz` (approximately 26 MB). When the transfer is complete, the CLI displays this message:
`Upgrade Will Take 4-6 Minutes. Please Wait...`

Once the upgrade is complete, the system automatically reboots. The CLI displays this message:

```
Final Message: Upgrade Complete.
Intel (r) NetStructure (tm) 1520 Cache is rebooting.
Please wait 2-3 minutes for an active console login.
```

After the system has finished rebooting, follow the procedures in *Starting the system for the first time* in chapter 2.

Patch upgrade

After you press Ctrl-X to proceed, the CLI displays this message:
`Installing The Patch. Please Wait...`

The server transfers the application upgrade `tar.gz` (typically less than 10 MB). When the transfer is complete, the CLI displays this message:
`Ftp Fetching Successful`

The appliance starts to install the upgrade. The CLI displays this message:
`Patch Installation In Progress. Please Wait...`

Once the upgrade is installed, the CLI displays this message:
`Patch Installation Successful`

Once the upgrade is complete, the system automatically reboots, then the CLI displays this message:

```
Final Message: Upgrade Complete.
Intel (r) NetStructure (tm) 1520 Cache is rebooting.
Please wait 2-3 minutes for an active console login.
```

Continue to use the appliance as before. If the upgrade requires you to reset the application, you are warned in an upgrade message.

OS/Application upgrade

After you press Ctrl-X to proceed, the CLI displays the message:
`Upgrading To The New OS. Please Wait...`

The server transfers the application image upgrade file tar.gz (typically 310 MB). When the transfer is complete, the CLI displays this message:
`Ftp Fetching Successful`

The appliance begins preparing the secondary disk, and the CLI displays this message:
`Disk Preparation in Progress. Please Wait...`

Once the disk is prepared, the CLI displays this message:
`Disk Preparation Successful`

Next, reboot the system. After the system has finished rebooting, follow the procedures in *Starting the system for the first time* in chapter 2.

Running a different version of the appliance software

You can switch between the two different versions of the software.

▼ Running a different version of the appliance software

1. Select the **main** menu, and press Enter.
2. Select **versions**, and press Enter.
3. Select **switch**, and press Enter. Doing so displays a list of versions. If no other versions exist, a message displays indicating such.
4. Select the version you want to run, and press Enter.

Deleting a version of the appliance software

You can delete a version of the appliance software when you need to add a newer version but you already have two versions installed.

Note You cannot delete the currently running version of the appliance software. To delete that software, you must first switch to the second version and then delete the other version. Also, if you have only one software version installed, you cannot delete it.

▼ Deleting a version of the appliance software

1. Select the **main** menu, and press Enter.
2. Select **versions**, and press Enter.
3. Select **delete**, and press Enter.
4. Select the version you want to delete, and press Enter. Doing so displays a confirmation prompt asking you whether you want to really delete the version.
5. When prompted, press y to confirm or n to cancel.

Viewing which version of the appliance software is currently running

You can check which version of the appliance software is running on your machine.

▼ Viewing the current version of the appliance

1. Select the **main** menu, and press Enter.
2. Select **versions**, and press Enter.
3. Select **current**, and press Enter. Doing so displays a message that indicates the current version number.

Clearing statistics

You can clear statistics that remain through reboot operations (persistent statistics). Clearing statistics from the appliance initializes them to a pre-installation state.

Note Clearing statistics involves stopping and restarting the appliance.

▼ Clearing statistics for the appliance

1. Select the **main** menu, and press Enter.
2. Select **stop**, and press Enter. Doing so stops all caching functions in the appliance and displays a status message indicating such.
3. Select **clear**, and press Enter. Doing so displays a confirmation prompt asking you whether you want to really clear statistics.
4. Be sure that y appears after the confirmation prompt and then press Enter.

5 Press CTRL-X to clear the statistics and return to the previous screen. Choosing to clear the statistics causes a confirmation message to appear.

6 Select **start**, and press Enter. Doing so resumes the caching functions in the appliance.

Rebooting the System

You can reboot the system. Rebooting the system is different than starting or stopping the caching software. A system reboot performs an orderly shutdown of the appliance and restarts the operating system.

▼ Rebooting the system

1 Select the **main** menu, and press Enter.

2 Select **reboot**, and press Enter. Doing so causes the system to reboot. The caching software retains its status (on or off) after the reboot operation.

Halting the System

You can halt the system. Halting the system is different than starting or stopping the caching software or rebooting the system. Halting the system gives little or no warning to users connected to the machine before logging them off. You should halt the appliance only as a last resort to problems.

▼ Halting the system

1 Select the **main** menu, and press Enter.

2 Select **halt**, and press Enter. Doing so causes a message to display that indicates the appliance is halting. Shortly after this message the CLI halts.

Changing the administrator password for telnet or serial access

Connecting to the appliance through telnet or a serial port requires you to enter an administrator ID and password. When you install the appliance, the default ID is `admin` and the password is `admin`. This procedure allows you to change the password. The username remains the same.

Note Should you forget your password, contact Customer Service at Intel Corporation for assistance. For information on how to contact Intel Customer Service, see the *Intel NetStructure Cache Appliance Product Support* booklet that came with your system.

Important For security, it is highly recommended that you change the password.

▼ Changing the password

1 Select the **main** menu, and press Enter.

2 Select **passwd**, and press Enter. Doing so causes a prompt to appear requesting you to type and confirm the new administrator password.

3 Enter and confirm the new password.

4 Press CTRL-X to save your changes and return to the previous screen.

Note Changing the password value using CLI changes only the password for telnet or serial access. It does not change the password for Manager UI access.

Resetting to factory settings

You can reset settings in the appliance to their factory defaults.

Warning Using this command deletes your installation and requires you to reinstall and reconfigure the appliance completely.

▼ Resetting the appliance to default factory settings

1 Select the **main** menu, and press Enter.

2 Select **reset**, and press Enter. Doing so displays a confirmation prompt asking you whether you want to really reset settings.

3 Be sure that y appear after the confirmation prompt and then press Enter.

4 Press CTRL-X to reset the settings and return to the previous screen. Choosing to reset the settings causes the appliance to stop and delete the installation, then returns you to the **setup** menu so you can reinstall the appliance again. See *Using the setup menu, on page 52* for more information.

Preparing a cache disk

You can prepare a cache disk for use in the system. You must prepare a new drive in the system before the caching software can use it. Preparing the drive allows the caching software to recognize the drive as a cache disk.

▼ Preparing a cache disk

1 Select the **main** menu, and press Enter.

2 Select **prep**, and press Enter. Doing so causes the system to examine the cache drives for uninitialized drives and prepare them for use.

Using the config menu

The **config** menu lets you do the following:

✔ Set general controls, such as shut down, bounce, start up, or restart the local appliance, and restart or bounce the cluster.

✔ Configure protocol options.

✔ Configure the cache.

✔ Configure security options.

- ✔ Configure routing options.
- ✔ Configure the Adaptive Redirection Module (ARM) for transparent proxy caching.
- ✔ Configure the host database options.
- ✔ Configure logging options.

Setting general controls

You can stop, start, or restart caching on the local appliance or cluster. You can also bounce the local appliance or the cluster. When you bounce the local appliance, caching is stopped and then quickly restarted on the local appliance. The same is true when you bounce the cluster, caching is stopped and then quickly restarted on each node in the cluster.

▼ Setting general controls

1. Select the **config** menu, and press Enter.
2. Select **server**, and press Enter.
3. Select the configuration option you want to use, and press Enter:
   - ✔ To specify the name of your cluster, select **cache rename**, and press Enter. Doing so displays the current cache name and a field in which you can enter a new name. After entering the new name, press CTRL-X to save your changes and return to the previous screen.
   - ✔ To enter a multicast group address, select **multicast address**, and press Enter. Doing so displays the current multicast address and a field in which you can enter the new multicast address. After entering the new address, press CTRL-X to save your changes and return to the previous screen.
   - ✔ To restart caching on the cluster, select **cluster restart**, and press Enter. See step four for further information.
   - ✔ To restart caching on the local appliance, select **local restart**, and press Enter. See step four for further information.
   - ✔ To shut down caching on the local appliance, select **local shutdown**, and press Enter. See step four for further information.
   - ✔ To start up caching on the local appliance, select **local startup**, and press Enter. See step four for further information.
   - ✔ To bounce the cluster, select **cluster bounce**, and press Enter. See step four for further information.
   - ✔ To bounce the local appliance, select **local bounce**, and press Enter. See step four for further information.
   - ✔ To set up an alarm email address, select **email**, and press Enter. Doing so displays the current alarm email address. You can enter the email

address you want to use in this field and press CTRL-X to save your changes and return to the previous screen.

- ✔ To see whether the appliance is in reverse or forward proxy mode, select **view-mode**, and press Enter. A message displays at the bottom of the screen that indicates reverse or forward proxy enabled.
- ✔ To set the appliance for reverse proxy, select **rev-proxy**, and press Enter.
- ✔ To set the appliance for forward proxy, select **forw-proxy**, and press Enter.

Note To use both forward and reverse proxy, set the appliance to reverse. If you are running in non-transparent mode, the proxy port is 80.

4. In some cases, you are prompted to confirm the action before it is performed. To continue with the action, be sure that `y` appears after the prompt when you press Enter. After pressing Enter, press CTRL-X to return to the previous screen. To cancel the operation, be sure `n` appears after the prompt and press Enter. Or you can press ESC to exit the screen.

Configuring protocol options

You can set HTTP, NNTP, and FTP configuration options. You can also set filter rules and remap rules. Filter rules let you deny or allow particular URL requests and keep or strip header information. Remap rules let you create a set of document routing rewrite rules for reverse proxy caching so that the appliance can handle relative path requests.

Configuring HTTP options

You can view the current configuration settings and remove HTTP headers.

▼ Configuring HHTP options

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **http**, and press Enter.
4. Select the configuration option you want to use, and press Enter:
   - ✔ To view the current HTTP configuration settings, select **view**, and press Enter.
   - ✔ To remove HTTP headers, select **remove**, and press Enter. You can remove the following headers:

     From: identifies the user's email address

     Referer: identifies the Web link followed by the user

User-Agent: identifies the agent making the request, usually a browser

Cookie: identifies the user that made the request

- ✔ To add HTTP headers, select **add**, and press Enter. You can add the following headers:

  From: identifies the user's email address

  Referer: identifies the Web link followed by the user

  User-Agent: identifies the agent making the request, usually a browser

  Cookie: identifies the user that made the request

- ✔ To remove a client IP header or undo the removal, select **remove/undo**, and press Enter. See insert/undo below.
- ✔ To insert a client IP header or undo the insertion, select **insert/undo**, and press Enter. When a client IP header is inserted, it allows the traffic server to track its IP as opposed to other means that common http protocol permits.
- ✔ Language: Messages from the traffic server to users are displayed by default in English.
- ✔ Auth: This is the proxy authorization. Because the proxy authorization header field applies only to the next outbound proxy that demanded authentication using the proxy-authenticate field, this feature is added so that you can force the traffic server to forward the header to the next proxy in the chain. By default, this is disabled. If you are running the traffic server through another proxy (for example, a firewall), you should enable this feature to make http authentication work.

Configuring NNTP options

You can configure enable and disable NNTP caching, view the current NNTP settings, enable and disable NNTP server feeds, enable and disable NNTP access control, configure NNTP servers, configure NNTP access, configure the NNTP port, set timeout values, and remove HTTP headers.

▼ Configuring NNTP options

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **nntp**, and press Enter.
4. Select the configuration option you want to use, and press Enter:
   - ✔ To view the current NNTP configuration settings, select **view**, and press Enter. The configuration settings display on screen.
   - ✔ To enable the appliance to cache and serve news articles select **enable**, and press Enter.

- ✔ To Disable the appliance from caching and serving news articles select **disable**, and press Enter.
- ✔ To allow NNTP server feeds, select the first **feeds** in the menu and press Enter.
- ✔ To inhibit NNTP server feeds select the second **feeds** in the menu and press Enter.
- ✔ To allow NNTP access control, select the first **access** in the menu and press Enter.
- ✔ To inhibit NNTP access control select the second **access** in the menu and press Enter.
- ✔ To configure NNTP servers, select **servers**, and press Enter. Refer to *Configuring NNTP servers* for more information.
- ✔ To configure NNTP access, select **access** and press Enter. Refer to *Configuring NNTP access, on page 69* for more information.

Configuring NNTP servers

You can add, delete, and view NNTP server rules. The appliance uses NNTP server rules to let you specify:

- The parent NNTP servers from which you want the appliance to cache articles.
- The news groups you want the appliance to observe.
- The type of NNTP activity you want the appliance to perform; for example, caching news articles on demand, posting news articles, and receiving news feeds.
- The network interface the appliance uses to contact the parent NNTP server.

▼ Adding NNTP server rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **nntp**, and press Enter.
4. Select **servers**, and press Enter.
5. Select **add rules**, and press Enter.
6. Enter an NNTP server rule, and press Enter.
7. Press CTRL-X to save your changes and return to the previous screen.

Each rule must have the following format:

```
hostname group-wildmat priority interface
```

The *hostname* and *group-wildmat* tags are required; *priority* and *interface* are optional.

The following table describes the tags you can use in a rule:

Tag	Description
hostname	*Choose one of the following:* ▮ host name ▮ host name:port ▮ IP address ▮ IP address:port ▮ .block—Use `.block` to block access to specific news groups.
group-wildmat	*This tag must be a comma-separated list of group names and list files in wildmat format (use * as a wildcard). The list file options are:* `subscriptions`*,* `distributions`*, and* `distrib.pats`*.* *Do not use spaces in the list. Use the prefix "*`!`*" to indicate groups not included in the list. The list is processed in reverse order, so more specific restrictions should be placed later in the list.* *Examples:* ▮ `*,!distrib.pats` *The previous example does not include any* `distrib.pats` *files, but does include all others.* ▮ `*,!alt.*` *The previous example does not include any groups of the form* `alt.*`*, but does include all others.* ▮ `talk.religion.*,!talk.religion.barney,subscriptions` *The previous example includes only subscriptions from all* `talk.religion.*` *groups but excludes* `talk.religion.barney`*.*
priority	*This tag tells the appliance how to treat the specified host and news groups. Use one of the following options:* ▮ <no priority tag> *If you do not use a priority tag, the appliance caches articles from the specified news groups on demand. If you specify multiple groups (such as* `alt.*`*), the appliance maintains a group list and will poll the parent NNTP server regularly to check for changes in the group list.*

Tag (Continued)	Description (Continued)
priority (continued)	▮ feed *The appliance will receive news feeds for the specified groups as the parent NNTP server receives news feeds. The appliance will not cache articles on demand, since it will have them.* ▮ push *The appliance can both receive news feeds and cache articles on demand.* ▮ pull *The appliance actively pulls (caches) all articles from these news groups at a frequency you specify in the appliance Manager UI. The appliance does not wait for user requests.* *A "pull" line must be preceded by a "cache on demand" line. The appliance needs to be aware of the news server and its groups before it can pull articles from a specific group. See the examples following this table.* ▮ pullover *The appliance actively pulls the overview database for the news groups but retrieves news articles on demand.* *A "pullover" line must be preceded by a "cache on demand" line. The appliance needs to be aware of the news server and its groups before it can pull overviews from a specific group. See the examples following this table.* ▮ dynamic *The appliance automatically decides, based on usage patterns, whether a group should be "pull," "pullover," or demand retrieval-based.* ▮ Enter a positive integer *The appliance retrieves articles on demand from the specified server according to the assigned priority. The default priority is 0. Multiple servers assigned the same priority are accessed in a round-robin fashion.* ▮ post *Articles to be posted to the specified news groups are sent to the specified server.*
interface	*Enter the network interface the appliance uses to contact the parent NNTP server.*

Examples

The following rule tells the appliance to block all requests from `rec.*` groups with the exception of `rec.soccer`:

```
.block !rec.soccer,rec.*
```

The following rule is an example of setting the port associated with the hostname:

```
news.webhost.com:999 *
```

The following rule is an example of associating an interface and priority with an IP address:

```
news.webhost.com * 0 10.3.3.2
```

The following rules are examples of establishing priorities for the hostnames:

```
news.webhost.com * 0
news.backup.com * 1
```

The following rules are examples of defining pull and pullover groups.

```
comp.webhost.com *
comp.webhost.com comp.* feed
```

Note Every line designating a pull or pullover group must be preceded by a "cache on demand" line as follows:

```
comp.webhost.com alt.*
comp.webhost.com alt.bicycles pull
```

▼ Deleting NNTP server rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **nntp**, and press Enter.
4. Select **servers**, and press Enter.
5. Select **delete**, and press Enter. Doing so displays a list of rules. If no rules exist, a message displays at the bottom of the screen indicating such.
6. Use the arrow keys to select the rule you want to delete and press Enter.
7. Press CTRL-X to save your change and return to the previous screen.

▼ Viewing NNTP server rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **nntp**, and press Enter.
4. Select **servers**, and press Enter.
5. Select **view**, and press Enter. Doing so displays the file containing the NNTP server rules.

Configuring NNTP access

The appliance uses NNTP access rules to let you control user access to news articles that are cached. Each rule describes the access privileges for a particular group of clients. You can add, delete, and view access rules.

▼ Adding NNTP access rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **nntp**, and press Enter.
4. Select **access**, and press Enter.
5. Select **add rules**, and press Enter.
6. Enter an NNTP access rule, and press Enter.
7. Press CTRL-X to save the rule and return to the previous screen.

Each rule must begin with a specific client group. You can use three ways to specify groups of clients: by IP range, domain, or host name. For example:

```
ip=0.0.0.0-255.255.255.255
ip=127.0.0.1
domain=intel.com
hostname=myhost.mydomain.com
```

Following the client group is an access directive. The access directive is of the form `access=`*`value`*. The allowed access values are `ip_allow`, `ip_deny`, `basic`, `generic`, and `custom`. Depending on the access directive, you can further specify an authenticator program, users, and passwords, as in the following examples:

```
ip=127.0.0.1 access=”generic” authenticator=”homebrew”
user=”joe”
hostname=myhost.com access=”basic” user=”joe” pass=”bob”
```

The following table lists the access directive options:

If access is...	authenticator is...	user is...	pass is...
`ip_allow`	*not required*	*not required*	*not required*
`ip_deny`	*not required*	*not required*	*not required*
`basic`	*not required*	*required*	*optional*
`generic`	*optional*	*not required*	*not required*
`custom`	*required*	*optional; but the only allowed entry is the string "required". (See the following example.)*	*optional; but the only allowed entry is the string "required". (See the following example.)*

The following is an example of *custom* access:

```
ip=127.0.0.1 access=″custom″ authenticator=″hb″ user=required pass=required
```

▼ Deleting NNTP access rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **nntp**, and press Enter.
4. Select **access**, and press Enter.
5. Select **delete**, and press Enter. Doing so displays a list of rules. If no rules exist, a message displays at the bottom of the screen indicating such.
6. Use the arrow keys to select the rule you want to delete and press Enter.
7. Press CTRL-X to save your change and return to the previous screen.

▼ Viewing NNTP access rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **nntp**, and press Enter.
4. Select **access**, and press Enter.
5. Select **view**, and press Enter. Doing so displays file containing the NNTP access rules.

Configuring Secure Socket Layer (SSL) port

You can view and specify the ports to which SSL is restricted.

▼ Viewing SSL ports

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **ssl**, and press Enter.
4. Select **view**, and press Enter. Doing so displays the ports to which SSL is restricted.

▼ Restricting SSL to specific ports

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **ssl**, and press Enter.
4. Select **port**, and press Enter. Doing so displays the current ports to which SSL is restricted and a field in which you can specify additional ports.
5. Supply the ports to which SSL will be restricted, and press Enter. You can enter a maximum of two ports. When entering more than one port, separate them with blank space. Also, you must enter the complete list of ports even if one is already specified in the existing list.
6. Press CTRL-X to save your changes and return to the previous screen.

Configuring FTP options

You can view the current FTP configuration settings, set the connection mode, the inactivity timeout value, and the anonymous password.

▼ Configuring the FTP options

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **ftp**, and press Enter.
4. Select the configuration option you want to use, and press Enter:
   - ✔ To view the current FTP configuration settings, select **view**, and press Enter. The configuration settings display on screen.
   - ✔ To set the connection mode, select **mode**, and press Enter. You can select from three modes: PASV/PORT, PASV only, and PORT only. Pressing Enter makes the selection.

- ✔ To set the inactivity timeout (the length of time the appliance waits for a response from the FTP server before abandoning the user's request for data), select **inactivity**, and press Enter. Doing so causes a field to appear with the current setting displayed. Supply the new value and press Enter. Press CTRL-X to save your changes and return to the previous screen.
- ✔ To set the anonymous password for FTP servers that require a password for access, select **password**, and press Enter. Doing so causes a field to appear with the current password displayed. Supply the new value and press Enter. Press CTRL-X to save your changes and return to the previous screen.

Setting filter rules

The appliance uses filter rules to deny or allow particular URL requests and keep or strip header information. When a URL request is allowed, the appliance will cache and serve the requested document. When a request is denied, the client receives an `access denied` message.

You can add, delete, and view filter rules.

▼ Adding filter rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **filter**, and press Enter.
4. Select **add rules**, and press Enter.
5. Enter a filter rule, and press Enter.
6. Press CTRL-X to save the rule and return to the previous screen.

Each rule must have the following format:

```
primary destination=value  secondary specifier=value  action=value
```

Note You can use more than one secondary specifier in a rule. However, you cannot repeat a secondary specifier.

The following table lists the primary destination tags and their allowed values:

Primary Destination	Allowed Value
`dest_domain`	*Requested domain name*
`dest_host`	*Requested host name*
`dest_ip`	*Requested IP address*
`url_regex`	*Regular expression to be found in a URL*

The secondary specifiers are optional. The following table lists the possible tags and their allowed values:

Secondary Specifier	Allowed Value
`time`	*A time range, such as 08:00-14:00*
`src_ip`	*The IP address of the client*
`prefix`	*A prefix in the path part of a URL*
`suffix`	*A file suffix in the URL*
`port`	*A requested URL port*
`method`	*A request URL method; one of the following:* ▮ get ▮ post ▮ put ▮ trace
`scheme`	*A request URL protocol; one of the following:* ▮ HTTP ▮ FTP

The following table lists the possible action tags and their allowed values:

Action	Value
`action`	▮ ip_allow ▮ ip_deny
`keep_hdr`	*Enter the client request header information that you want to keep:* ▮ date ▮ host ▮ cookie ▮ client_ip
`strip_hdr`	*Enter the client request header information that you want to strip. You have the same options as* `keep_hdr`.

Examples

The following rule tells the appliance to deny FTP document requests to the IP address 112.12.12.12.

```
dest_ip=112.12.12.12 scheme=ftp action=ip_deny
```

The following rule tells the appliance to keep the client IP address header for URL addresses that contain the regular expression `politics` and whose path prefix is `/viewpoint`.

```
url_regex=politics prefix=/viewpoint keep_hdr=client_ip
```

The following rule tells the appliance to strip all cookies to the requested host `www.intel.com`.

```
dest_host=www.intel.com strip_hdr=cookie
```

The following rule tells the appliance not to allow `puts` to the requested host `www.intel.com`.

```
dest_host=www.intel.com method=put action=ip_deny
```

▼ Deleting filter rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **filter**, and press Enter.
4. Select **delete**, and press Enter. Doing so causes a list of the rules to appear. If no rules exist, a message appears at the bottom of the screen indicating such.
5. Use the arrow keys and move to the rule you want to delete, and press Enter.
6. Press CTRL-X to save your changes and return to the previous screen.

▼ Viewing filter rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.
3. Select **filter**, and press Enter.
4. Select **view**, and press Enter. Doing so displays the file containing the filter rules.

Setting remap rules

For reverse proxy caching, the appliance uses remap rules to map an origin server to the appropriate location on the appliance.

Remap rules are also used to modify location headers. Origin servers might respond to a request with a location header that redirects the client to another location. Origin server location headers must be *reverse mapped* so that clients do not bypass the appliance when they make redirected requests.

You can add, delete, and view remap rules.

▼ Adding remap rules

1. Select the **config** menu, and press Enter.
2. Select **protocols**, and press Enter.

3 Select **remap**, and press Enter.

4 Select **add rules**, and press Enter.

5 Enter a remap rule, and press Enter.

6 Press CTRL-X to save your changes and return to the previous screen.

Each rule must consist of three fields: *type target replacement*.

The following table describes the proper format for each field.

Field	Description
type	*Enter either one of the following:* ▮ map—maps an incoming request URL to the appropriate origin server URL. ▮ reverse_map—use for location header modifying rules.
target	*Enter the from URL. You can enter up to four components:* <scheme>://<host>:<port>/<path_prefix>
replacement	*Enter the to URL. You can enter up to four components:* <scheme>://<host>:<port>/<path_prefix>

For more detailed information about remapping rules, refer to *Understanding server acceleration mapping rules*, *on page 132*.

▼ Deleting remap rules

1 Select the **config** menu, and press Enter.

2 Select **protocols**, and press Enter.

3 Select **remap**, and press Enter.

4 Select **delete**, and press Enter. Doing so displays a list of the current remap rules. If no rules exist, a message appears at the bottom of the screen indicating such.

5 Use the arrow keys and position the cursor over the rule you want to delete, and press Enter.

6 Press CTRL-X to save your changes and return to the previous screen.

▼ Viewing remap rules

1 Select the **config** menu, and press Enter.

2 Select **protocols**, and press Enter.

3 Select **remap**, and press Enter.

4 Select **view**, and press Enter. Doing so displays the file containing the remap rules.

Configuring the cache

You can configure cache storage options to do the following:

- ✔ Enable caching of objects for different protocols.
- ✔ Set disk storage options.
- ✔ Set freshness properties.
- ✔ Set caching rules.

Enabling caching for different protocols

You can configure the appliance to cache objects retrieved via the HTTP, NNTP, and FTP protocols. You can also choose to ignore or obey user requests to bypass the cache.

▼ Enabling caching for different protocols

1. Select the **config** menu, and press Enter.
2. Select **cache**, and press Enter.
3. Select **activation**, and press Enter.
4. Select the configuration option you want to change.

Note: You are not prompted for confirmation. Make sure you want to complete the action before you select one of the following options, and press Enter.

- ✔ To enable HTTP caching, select the first **HTTP**, and press Enter.
- ✔ To disable HTTP caching, select the second **HTTP**, and press Enter.
- ✔ To enable NNTP caching, select the first **NNTP**, and press Enter.
- ✔ To disable NNTP caching, select the second **NNTP**, and press Enter.
- ✔ To enable FTP caching, select the first **FTP**, and press Enter.
- ✔ To disable FTP caching, select the second **FTP**, and press Enter.
- ✔ To ignore user requests to bypass the cache (ignore client Cache Control: no-cache headers), select the first **Bypass**, and press Enter.
- ✔ To obey user requests to bypass the cache (obey client Cache Control: no-cache headers), select the second **Bypass**, and press Enter.

After you press Enter, your selection displays at the bottom of the screen.

Setting disk storage options

You can configure the cache to store only objects below a certain size and to store a limited number of alternates.

▼ Setting disk storage options

1. Select the **config** menu, and press Enter.
2. Select **cache**, and press Enter.
3. Select **storage**, and press Enter. Doing so causes the Configure Cache Storage box to appear. This box shows the current settings for maximum object size and maximum number of alternates allowed in the cache.
4. In the **New HTTP/FTP Object Size** field, type the maximum size of the HTTP or FTP objects that you want the appliance to cache, and press Enter.
5. In the **New Maximum number of alternates** field, type the maximum number of alternates that you want the appliance to cache, and press Enter.
6. Press CTRL-X to save your changes and return to the previous screen.

Setting object freshness options

You can configure how fresh you want the appliance to keep your documents in the cache.

▼ Setting freshness properties

1. Select the **config** menu, and press Enter.
2. Select **cache**, and press Enter.
3. Select **freshness**, and press Enter. Doing so displays a list of options. Each of these options has several selections you can choose from. Use the arrow keys to position the cursor over the option you want and press Enter.

The following table shows the options:

Option	Description
Options to Verify freshness	*Choosing this option lets you configure how the appliance asks the original content server to verify the freshness of objects (revalidate them) before serving them.* *Select from one of the following options and press Enter. After pressing Enter press CTRL-X to save your changes and return to the previous screen.* ▮ When The Object Has Expired—The appliance revalidates objects with explicit expiration dates after they expire. Otherwise, it uses heuristic methods to evaluate freshness and revalidates the object should it be stale. ▮ When The Object Has Expired Or Has No Expiry Date—The appliance revalidates objects with explicit expiration dates after they expire. All other documents are revalidated before serving. ▮ Always—The appliance always revalidates objects before serving them. ▮ Never—The appliance never checks object freshness.
Freshness information	*Specifies the minimum freshness information required when considering to cache a document.* *Select from one of the following options and press Enter. After pressing Enter press CTRL-X to save your changes and return to the previous screen.* ▮ An Explicit Lifetime—The appliance only caches objects with `Expires` headers or `Cache-Control: max-age` headers. ▮ A Last Modified Time—The appliance only caches objects with `Expires` headers, or `Cache-Control: max-age headers`, or `Last-Modified` headers. ▮ Nothing—The appliance caches documents regardless of freshness headers.

Option (Continued)	Description (Continued)
Set FTP objects expiry	*FTP objects carry no time stamp or date information. The appliance considers them fresh for the amount of time specified here. This "freshness" time is counted from the time the object arrives in the cache.* *Enter the time in seconds and press Enter. After pressing Enter, press CTRL-X to save your changes and return to the previous screen.*
Internet Explorer options	*Versions of Microsoft Internet Explorer do not request cache reloads from reverse proxies and transparent caches when the user presses the browser Refresh button. This behavior can prevent users from manually reloading content directly from the origin servers. You can configure the appliance to treat Microsoft Internet Explorer requests more conservatively. Doing so provides fresher content at the cost of serving fewer documents from cache.* *Internet Explorer requests force a check with the origin server.* *Select from one of the following options and press Enter. After pressing Enter press CTRL-X to save your changes and return to the previous screen.* ▮ Never ▮ For IMS Revalidation Requests ▮ Always

Configuring caching rules

The appliance uses caching rules to determine how a particular group of URL addresses should be cached. You can add, delete, and view caching rules. Caching rules can specify:

- ✔ Whether to cache objects
- ✔ How long to keep (pin) particular objects in the cache
- ✔ How long to consider cached objects as fresh
- ✔ Whether to ignore no-cache directories from the server

▼ Adding caching rules

1. Select the **config** menu, and press Enter.
2. Select **cache**, and press Enter.

3 Select **rules**, and press Enter.

4 Select **add rules**, and press Enter.

5 Enter a caching rule, and press Enter.

6 Press CTRL-X to save your rule and return to the previous screen.

Each rule must have the following format:

```
primary destination=value   secondary specifier=value   action=<value
```

The following table lists the supported primary destinations and their allowed values:

Primary Destination	Allowed Value
`dest_domain`	*Requested domain name*
`dest_host`	*Requested host name*
`dest_ip`	*Requested IP address*
`url_regex`	*Regular expression to be found in a URL*

The secondary specifiers are optional. The following table lists the possible tags and their allowed values.

Note You can use more than one secondary specifier in a rule. However, you cannot repeat a secondary specifier.

Secondary Specifier	Allowed Value
`time`	*A time range, such as 08:00-14:00*
`src_ip`	*The IP address of the client*
`prefix`	*A prefix in the path part of a URL*
`suffix`	*A file suffix in the URL*
`port`	*A requested URL port*
`method`	*A request URL method; use one of the following:* ▮ get ▮ post ▮ put ▮ trace
`scheme`	*A request URL protocol; use one of the following:* ▮ HTTP ▮ FTP

The following table lists the possible action tags and their allowed values:

Action	Value
`action`	▮ never-cache ▮ ignore-no-cache
`pin-in-cache`	*Enter the amount of time you want to keep the object(s) in the cache. Use the following time formats:* ▮ h for hours, e.g. 10h ▮ m for minutes, e.g. 5m ▮ s for seconds, e.g. 20s ▮ mixed units, e.g. 1h15m20s
`revalidate`	*Enter the amount of time you want to consider the object(s) fresh. Use the same time formats that are shown in* `pin-in-cache`.

Examples

The following rule tells the appliance to never cache FTP documents requested from the IP address 112.12.12.12.

```
dest_ip=112.12.12.12 scheme=ftp action=never-cache
```

The following rule tells the appliance to keep in the cache for 12 hours documents whose URL addresses contain the regular expression `politics` and whose the paths contain the prefix `/viewpoint`.

```
url_regex=politics prefix=/viewpoint pin-in-cache=12h
```

▼ Deleting cache rules

1. Select the **config** menu, and press Enter.
2. Select **cache**, and press Enter.
3. Select **rules**, and press Enter.
4. Select **delete rules**, and press Enter. Doing so displays a list of the current rules. If no rules exits, a message appears at the bottom of the screen indicating such.
5. Use the arrow keys to position the cursor over the rule you want to delete and press Enter.
6. Press CTRL-X to save your changes and return to the previous screen.

▼ Viewing cache rules

1. Select the **config** menu, and press Enter.
2. Select **cache**, and press Enter.
3. Select **rules**, and press Enter.
4. Select **view rules**, and press Enter. Doing so displays the file containing the cache rules.

Configuring security options

You can control client access to the appliance and access to the Manager UI.

Controlling client access to the appliance

The appliance uses *IP Allow* rules to specify ranges of IP addresses that are allowed to use the appliance as a web proxy. If you want to deny access to specific IP addresses, do not include them in an IP Allow rule. You can add, delete, and view IP Allow rules.

▼ Adding IP Allow rules

1. Select the **config** menu, and press Enter.
2. Select **security**, and press Enter.
3. Select **server**, and press Enter.
4. Select **add rules**, and press Enter.
5. Enter an IP allow rule, and press Enter.
6. Press CTRL-X to save your rule and return to the previous screen.

   Each rule must have the following format:

   ```
   src_ip=IPaddress or IPaddress_range action=ip_allow
   ```

   The IP address or range of IP addresses specified in the `src_ip` field are allowed to use the appliance as a web proxy.

   Examples

   The following rule allows all clients to use the appliance as a web proxy:

   ```
   src_ip=0.0.0.0-255.255.255.255 action=ip_allow
   ```

   The following rule allows a specific subnet to use the appliance as a web proxy:

   ```
   src_ip=123.12.3.000-123.12.3.123 action=ip_allow
   ```

▼ Deleting IP Allow rules

1. Select the **config** menu, and press Enter.
2. Select **security**, and press Enter.

3 Select **server**, and press Enter.

4 Select **delete rules**, and press Enter. Doing so displays a list of current rules. If no rules exist, a message displays at the bottom of the screen indicating such.

5 Use the arrow keys to position the cursor over the rule you want to delete, and press Enter.

6 Press CTRL-X to save your changes and return to the previous screen.

▼ Viewing IP Allow rules

1 Select the **config** menu, and press Enter.

2 Select **security**, and press Enter.

3 Select **server**, and press Enter.

4 Select **view rules**, and press Enter. Doing so displays the file containing the IP Allow rules.

Controlling access to the Manager UI

The appliance uses *Manager Allow* rules to specify ranges of IP addresses that are allowed to access the Manager UI. If you want to deny Manager UI access to specific IP addresses, do not include them in a Manager Allow rule. You can add, delete, and view Manager Allow rules.

▼ Adding Manager Allow rules

1 Select the **config** menu, and press Enter.

2 Select **security**, and press Enter.

3 Select **mgmt**, and press Enter.

4 Select **add rules**, and press Enter.

5 Enter a rule, and press Enter.

6 Press CTRL-X to save your rule and return to the previous screen.

Each rule must have the following format:

```
src_ip=IPaddress or IPaddress_range action=ip_allow
```

The IP address or range of IP addresses specified in the `src_ip` field are allowed to access the Manager UI.

Examples

The following rule allows one user to access the Manager UI:

```
src_ip=123.12.3.123 action=ip_allow
```

The following rule allows a range of IP addresses to access the Manager UI:

```
src_ip=123.12.3.000-123.12.3.123 action=ip_allow
```

▼ Deleting Manager Allow rules

1. Select the **config** menu, and press Enter.
2. Select **security**, and press Enter.
3. Select **mgmt**, and press Enter.
4. Select **delete rules**, and press Enter. Doing so displays a list of the current rules. If no rules exist, a message displays at the bottom of the screen indicating such.
5. Use the arrow keys to position the cursor over the rule you want to delete, and press Enter.
6. Press CTRL-X to save your changes and return to the previous screen.

▼ Viewing Manager Allow rules

1. Select the **config** menu, and press Enter.
2. Select **security**, and press Enter.
3. Select **mgmt**, and press Enter.
4. Select **view rules**, and press Enter. Doing so displays the file containing the Manager Allow rules.

Configuring routing options

You can configure ICP peers (parent and sibling caches), control HTTP parent proxy services, and configure Web cache control protocol.

Configuring and maintaining ICP peers

You can do the following when configuring and maintain ICP peers:

- ✔ View and modify ICP rules
- ✔ View current ICP settings
- ✔ Enable ICP
- ✔ Disable ICP
- ✔ Enable multicast
- ✔ Disable multicast
- ✔ Set ICP port numbers
- ✔ Set ICP query timeout

Viewing and modifying ICP rules

The appliance uses ICP rules to define parent and sibling caches. You can add, delete, and view ICP rules.

▼ Adding ICP rules

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **rules**, and press Enter.
5. Select **add rules**, and press Enter.
6. Add an ICP rule, and press Enter.
7. Press CTRL-X to save your rule and return to the previous screen.

Each rule must contain the name and configuration information for a single ICP peer in the following format:

```
host:hostIP:cache_type:proxy_port:icp_port:MC_on:MC_IP:MC_TTL:
```

The following table describes each field:

Field	Description
host	*The host name of the ICP peer. The name* `localhost` *is reserved for the appliance.*
host IP	*The IP address of the ICP peer.*
cache_type	*The cache type. Use the following options:* ▮ 1 to indicate an ICP parent cache ▮ 2 to indicate an ICP sibling cache *Option 3 is reserved for the local host (the appliance itself).*
proxy_port	*The port number of the TCP port used by the ICP peer for proxy communication.*
icp_port	*The port number of the UDP port used by the ICP peer for ICP communication.*
MC_on	*Multicast on/off. Use the following options:* ▮ 0 if multicast is not enabled ▮ 1 if multicast is enabled

Field (Continued)	Description (Continued)
MC_IP	*The multicast IP address.* *If MC_on is disabled, appliance ignores this field.*
MC_TTL	*The multicast time to live. Use the following options:* ▮ 1 if IP multicast datagrams will not be forwarded beyond a single subnetwork ▮ 2 to allow delivery of IP multicast datagrams to more than one subnet (if there are one or more multicast routers attached to the first hop subnet) *If MC_on is disabled, appliance ignores this field.*

Example

The following example configuration is for three nodes: the local host, one parent, and one sibling:

```
localhost:0.0.0.0:3:8080:3130:0:0.0.0.0:0:
host1:123.12.1.23:1:8080:3131:0:0.0.0.0:0:
host2:123.12.1.24:2:8080:3131:0:0.0.0.0:0:
```

▼ Deleting ICP rules

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **rules**, and press Enter.
5. Select **delete rules**, and press Enter. Doing so displays a list of current rules. If no rules exist, a message displays at the bottom of the screen indicating such.
6. Use the arrow keys to position the cursor over the rule you want to delete, and press Enter.
7. Press CTRL-X to save your changes and return to the previous screen.

▼ Viewing ICP rules

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **rules**, and press Enter.
5. Select **view rules**, and press Enter. Doing so causes the file containing the ICP rules to appear.

Viewing current ICP settings

You can find out if the ICP protocol is enabled or disabled, what the ICP port number is, whether ICP multicast is enabled or disabled, and the ICP query timeout by viewing the settings.

▼ Viewing ICP settings

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **view**, and press Enter.

Enabling and disabling ICP

You can enable or disable ICP.

▼ Enabling ICP

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **enable-icp**, and press Enter.

▼ Disabling ICP

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **disable-icp**, and press Enter.

Enabling and disabling multicast in ICP

You can enable or disable multicast in ICP.

▼ Enabling multicast in ICP

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **enable-multicast**, and press Enter.

▼ Disabling multicast in ICP

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **disable-multicast**, and press Enter.

Setting the ICP port number

You can set the ICP port number.

▼ Setting the ICP port number

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **port**, and press Enter. Doing so causes a field to appear that has the current port number displayed.
5. Supply the port number in the data field, and press Enter.
6. Press CTRL-X to save your changes and return to the previous screen.

Setting the ICP query timeout

You can set the ICP query timeout number.

▼ Setting the ICP query timeout number

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **icp**, and press Enter.
4. Select **timeout**, and press Enter. Doing so causes a field to appear that has the current timeout value in seconds displayed.
5. Supply the new timeout value in seconds in the data field, and press Enter.
6. Press CTRL-X to save your changes and return to the previous screen.

Controlling parent proxy caching

The appliance uses parent proxy rules to set up parent proxy hierarchies with multiple parents and parent failover, and to configure selected URL requests to bypass parent proxies.

You can enable and disable parent proxy caching as well as configure parent proxy caching rules.

Note For the parent proxy rules to take effect, HTTP parent proxy services must be enabled in the Manager UI.

▼ Enabling parent proxy caching rules

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **parent**, and press Enter.
4. Select **enable**, and press Enter.

▼ Disabling parent proxy caching rules

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **parent**, and press Enter.
4. Select **disable**, and press Enter.

▼ Adding parent proxy caching rules

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **parent**, and press Enter.
4. Select **rules**, and press Enter.
5. Select **add rules**, and press Enter.
6. Enter a parent proxy rule, and press Enter.

7 Press CTRL-X to save your rule and return to the previous screen.

Each rule must have the following format:

```
primary destination=value secondary specifier=value action=value
```

The following table lists the primary destinations and their allowed values:

Primary Destination	Allowed Value
`dest_domain`	*Requested domain name*
`dest_host`	*Requested host name*
`dest_ip`	*Requested IP address*
`url_regex`	*Regular expression to be found in a URL*

The secondary specifiers are optional. The following table lists the possible tags and their allowed values:

Secondary Specifiers	Allowed Value
`time`	
`src_ip`	*The IP address of the client*
`prefix`	*A prefix in the path part of a URL*
`suffix`	*A file suffix in the URL*
`port`	*A requested URL port*
`method`	*A request URL method; one of the following:* ▮ get ▮ post ▮ put ▮ trace
`scheme`	*A request URL protocol; one of the following:* ▮ HTTP ▮ FTP

The following table lists the allowed action tags and their possible values:

Action Tag	Allowed Value
`parent`	*An ordered list of parent proxies. If the request cannot be handled by the last parent server in the list, it will be routed to the origin server.*

Action Tag	Allowed Value (Continued)
round_robin	▮ true *Enter true if you want the appliance to go through the parent proxy list in a round-robin.* ▮ false
go_direct	▮ true *Enter true if you want requests to bypass parent hierarchies and go directly to the origin server.* ▮ false *Enter false if you do not want requests to bypass parent hierarchies.*

Examples

The following rule sets up a parent proxy hierarchy consisting of the appliance (which is the child) and two parents, `p1` and `p2`. All `get` requests, if they cannot be served by the appliance, are routed to the first parent server, `p1.x.com`. If they are not in the first parent server, they are routed to the second parent server, `p2.y.com`. Because `round_robin=true`, the parent servers are queried in a round-robin fashion.

```
dest_domain=. method=get parent="p1.x.com:8080; p2.y.com:8080" round_robin=true
```

The following rule tells the appliance to route all requests containing the regular expression `politics` and the path /`viewpoint` directly to the origin server (bypassing any parent hierarchies).

```
url_regex=politics prefix=/viewpoint go_direct=true
```

Every rule must contain either a `parent=` or `go_direct=` directive.

▼ Deleting parent proxy caching rules

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **parent**, and press Enter.
4. Select **rules**, and press Enter.
5. Select **delete rules**, and press Enter. Doing so displays a list of current rules. If no rules exist, a message displays at the bottom of the screen indicating such.
6. Use the arrow keys to position the cursor over the rule you want to delete, and press Enter.
7. Press CTRL-X to save your changes and return to the previous screen.

▼ Viewing parent proxy caching rules

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **parent**, and press Enter.
4. Select **delete**, and press Enter.
5. Select **view rules**, and press Enter. Doing so lists the file containing the parent proxy caching rules.

Configuring WCCP options

The appliance supports WCCP 2.0-enabled routers. If you use WCCP, you must specify the IP address of the router.

You can enable, disable, configure, and view WCCP options.

▼ Enabling WCCP

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **wccp**, and press Enter.
4. Select **enable WCCP**, and press Enter.

▼ Disabling WCCP

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **wccp**, and press Enter.
4. Select **disable**, and press Enter.

▼ Configuring WCCP options

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **wccp**, and press Enter.
4. **configure**, and press Enter. Doing so causes a menu to appear that allows you to configure WCCP options.

   ✔ To enable security, select **enable security**, and press Enter. Doing so causes two fields to appear in which you can enter and confirm the password. Supply the password in the top field and press Enter. Supply the password in the second field and press Enter. Finally, press CTRL-X to save your changes and return to the previous screen.

   ✔ To disable security, select **disable security**, and press Enter.

   ✔ To enable multicast communication, select **enable multicast**, and press Enter. Doing so causes two fields to appear. Supply the multicast address in the top field and press Enter. Supply the multicast TTL in the

bottom field and press Enter. Finally, press CTRL-X to save your changes and return to the previous screen.

- ✔ To disable multicast communication, select **disable multicast**, and press Enter.
- ✔ To enable HTTP redirection, select **enable HTTP**, and press Enter.
- ✔ To disable HTTP redirection, select **disable HTTP**, and press Enter.
- ✔ To enable NNTP redirection, select **enable NNTP**, and press Enter.
- ✔ To enable NNTP redirection, select **enable NNTP**, and press Enter.
- ✔ To add a router, select **add router**, and press Enter. Doing so causes a field to appear. Supply the router IP and press Enter. Press CTRL-X to save your changes and return to the previous screen.
- ✔ To delete all routers, select **delete routers**, and press Enter.

▼ Viewing current WCCP options

1. Select the **config** menu, and press Enter.
2. Select **routing**, and press Enter.
3. Select **wccp**, and press Enter.
4. Select **view**, and press Enter.

Configuring the Adaptive Redirection Module (ARM)

You can configure the ARM for transparent proxy caching, set bypass rules, and configure load-shedding options.

Enabling and disabling transparent redirection

You can enable or disable transparent HTTP/NTTP.

▼ Enabling transparent redirection

1. Select the **config** menu, and press Enter.
2. Select **arm**, and press Enter.
3. Select **nat**, and press Enter.
4. Select **enable**, and press Enter.

▼ Disabling transparent redirection

1. Select the **config** menu, and press Enter.
2. Select **arm**, and press Enter.
3. Select **nat**, and press Enter.
4. Select **disable**, and press Enter.

Configuring ARM bypass rules

The appliance uses ARM bypass rules to determine whether to bypass incoming client requests or to attempt to serve them transparently.

You can add, delete, and view ARM bypass rules.

▼ Adding ARM bypass rules

1. Select the **config** menu, and press Enter.
2. Select **arm**, and press Enter.
3. Select **bypass**, and press Enter.
4. Select **rules**, and press Enter.
5. Select **add rules**, and press Enter.
6. Add a bypass rule, and press Enter.
7. Press CTRL-X to save your rule and return to the previous screen.

You can configure three types of bypass rules:

Rule	Description
Source bypass	*Configures the appliance to bypass a particular source IP address or range of IP addresses. For example, use this solution to bypass clients that do not want to use caching.*
Destination bypass	*Configures the appliance to bypass a particular destination IP address or range of IP addresses. For example, these could be destination servers that use IP authentication based on the client's real IP address.* *Destination bypass rules prevent the appliance from caching an entire site. You will experience hit rate impacts if the site you bypass is popular.*
Source/ Destination pair bypass	*Configures the appliance to bypass requests that originate from the specified source to the specified destination. For example, you can route around specific client-server pairs that experience broken IP authentication or out-of-band HTTP traffic problems when cached. Source/destination bypass rules can be preferable to destination rules because they block a destination server only for users that experience problems.*

The bypass rules have the following format:

Rule	Format
source IP bypass	`bypass src src_IP` *Where `src_IP` can be:* ▮ A simple IP address, such as `1.1.1.1` ▮ In Classless Inter-Domain Routing (CIDR) format, such as `1.1.1.0/24` ▮ A range of IP addresses separated by a dash, such as `1.1.1.1-2.2.2.2` ▮ Any combination of the above, separated by commas, such as `1.1.1.0/24, 25.25.25.25, 123.1.23.1 - 123.1.23.123`
destination IP bypass	`bypass dst dst_IP` *Where `dst_IP` can have the same format as `src_IP`*
source/ destination IP bypass	`bypass src IP_address AND dst IP_address` *Where `IP_address` must be a single IP address, such as 1.1.1.1*

Examples

The following examples show source, destination, and source/destination bypass rules:

```
bypass src 1.1.1.0/24, 25.25.25.25, 128.252.11.11 - 128.252.11.255
bypass dst 24.24.24.0/24
bypass src 25.25.25.25 AND dst 24.24.24.0
```

▼ Deleting ARM bypass rules

1. Select the **config** menu, and press Enter.
2. Select **arm**, and press Enter.
3. Select **bypass**, and press Enter.
4. Select **rules**, and press Enter.
5. Select **delete rules**, and press Enter. Doing so displays a list of current rules. If no rules exist, a message displays at the bottom of the screen indicating such.
6. Use the arrow keys to position the cursor over the rule you want to delete, and press Enter.
7. Press CTRL-X to save your changes and return to the previous screen.

▼ Viewing ARM bypass rules

1. Select the **config** menu, and press Enter.
2. Select **arm**, and press Enter.

3 Select **bypass**, and press Enter.

4 Select **rules**, and press Enter.

5 Select **view rules**, and press Enter. Doing so displays the file containing ARM bypass rules.

Configuring load-shedding options

When transparent proxy caching is enabled, the appliance handles overload conditions by forwarding new requests to origin servers. You can configure the appliance to automatically shed load if the cache-hit transaction times become too long.

▼ Configuring load-shedding options

1 Select the **config** menu, and press Enter.

2 Select **arm**, and press Enter.

3 Select **shedding**, and press Enter. Doing so displays a field that has the current value for the maximum number of connections.

4 Supply the maximum number of connections in the field and press Enter.

5 Press CTRL-X to save your changes and return to the previous screen.

Configuring the host database options

The appliance host database stores the domain name server (DNS) entries of servers that are contacted to fulfill user requests. You can configure and view the host database.

▼ Configuring host database options

1 Select the **config** menu, and press Enter.

2 Select **hostdb**, and press Enter.

3 Select **configure**, and press Enter. Doing so displays current values for the options you can set.

4 Supply a value for each field you want to change, and press Enter after filling in each field.

5 Press CTRL-X to save your changes and return to the previous screen.

The following table describes the options:

Option	Description
Lookup Timeout	*Specifies the timeout period in seconds for the IP address lookup operation in the host database.*
Foreground Timeout	*Specifies how long DNS entries can remain in the database before they are flagged as stale. For example, if foreground timeout is 24 hours, and a user requests an entry that has been in the database for 24 hours or longer, the entry is refreshed before being served.* *You can set the background timeout (see next item) to refresh entries in the background, before objects become stale.* *Be careful that you don't set the foreground timeout too low as you might slow response time. Also, setting this time too high risks accumulation of incorrect information. Setting the foreground timeout to greater than or equal to the background timeout disables background refresh.*
Background Timeout	*Specifies how long DNS entries can remain in the database before they are flagged as entries to refresh in the background. These entries are still fresh, so they can be refreshed after they are served, rather than before.* *For example, suppose the foreground timeout is 24 hours and the background timeout is 12 hours. A user requests an object from* `my.com` *and 16 hours later, a user makes a second request for an object from* `my.com`*. The DNS entry for* `my.com` *has not been refreshed in the foreground because the entry is not yet 24 hours old. But since the background timeout has expired, the appliance will first serve the user's request, then refresh the entry in the background.*
Invalid Host Timeout	*Specifies how long the proxy software should remember that a host name is invalid. This is often called negative DNS caching.* *For example, if a user specifies an invalid host name, the appliance informs the user that it could not resolve the name, and the appliance gets another request for the same host name. If the appliance still remembers the bad name, it won't try to look it up again, but will send another "invalid host name" message to the user.*
Re-DNS On Reload	*Re-resolves host names whenever clients reload pages.*

Option	Description
DNS Resolve Timeout	*Specifies how long the appliance should wait for the DNS server to respond with an IP address, even if the client request has been cancelled.* *If the user abandons the request before this timeout expires, the appliance can still obtain the host's IP address in order to cache it. The next time a user makes the same request, the address will be in the cache.*
Number of DNS Retries	*Specifies how many times the appliance should allow a look-up operation to fail before it abandons the operation and sends an "invalid host name" message to the user.*

▼ Viewing host database options

1. Select the **config** menu, and press Enter.
2. Select **hostdb**, and press Enter.
3. Select **view**, and press Enter.

Configuring logging options

You can configure the logging options used in the appliance. The appliance is able to keep system logs of events and statistical information. You can enable, disable, configure, and view the logging options.

▼ Enabling logging options

1. Select the **config** menu, and press Enter.
2. Select **logging**, and press Enter.
3. Select **enable**, and press Enter.

▼ Disabling logging options

1. Select the **config** menu, and press Enter.
2. Select **logging**, and press Enter.
3. Select **disable**, and press Enter.

▼ Configuring logging options

1. Select the **config** menu, and press Enter.
2. Select **logging**, and press Enter.
3. Select **collation**, and press Enter. Doing so displays current values in separate fields. Following are default values:

```
Enter Collation FTP Host:
Collation Interval (hours): 3
Enter Collation FTP User: admin
Enter Collation FTP Password: admin
Enter Collation FTP Directory: ~/logs
```

4 Supply a value for each field you want to change, and press Enter after filling in each field.

5 Press CTRL-X to save your changes and return to the previous screen.

▼ Viewing logging options

1 Select the **config** menu, and press Enter.

2 Select **logging**, and press Enter.

3 Select **view**, and press Enter.

Using the monitor menu

The **monitor** menu lets you view the following:

✔ Node performance statistics

✔ Protocol performance statistics

✔ Cache performance statistics

✔ Other performance statistics, such as host database, DNS, and cluster

Viewing Node statistics

Node statistics report performance information about the appliance system. These statistics include document hit rates, the number of HTTP transactions per second, and the number of open client and server connections.

▼ Viewing node statistics

1 Select the **monitor** menu, and press Enter.

2 Select **node**, and press Enter. Doing so causes statistics to display on the screen. The following table describes the statistics listed. Statistics fall into three categories: cache, in progress, and network.

	Statistic	Description
Cache	*Document Hit Rate*	*The ratio of cache hits to total cache requests, averaged over 10 seconds. This value is refreshed every 10 seconds.*
	Bandwidth Savings	*The ratio of bytes served from the cache to total requested bytes, averaged over 10 seconds. This value is refreshed every 10 seconds.*
	Cache Percent Free	*The ratio of cache free space to total cache space.*
In Progress	*Open Server Connections*	*The number of currently open server connections.*

	Statistic (Continued)	Description (Continued)
	Open Client Connections	*The number of currently open client connections.*
	Cache Transfers in Progress	*The number of cache transfers (cache reads and writes) in progress.*
Network	*Client Throughput (Mbit/sec)*	*The number of bytes per second through node (and cluster).*
	Transactions Per Second	*The number of HTTP transactions per second.*

Viewing Protocol statistics

Protocol statistics report the appliance system's use of the HTTP, NNTP, FTP, and ICP protocols.

▼ Viewing protocol statistics

1. Select the **monitor** menu, and press Enter.
2. Select **protocols**, and press Enter. Doing so causes a list of protocols to appear on screen.
3. Select the protocol you want to view and press Enter.

   The following table describes the statistics for the HTTP-trans protocol.

Statistics	Description
Hits	▮ Fresh—The percentage of hits that are fresh and their average transaction times. ▮ Stale Revalidated—The percentage of hits that are stale and revalidated, turn out to be still fresh and served, and their average transaction times.
Misses	▮ Now Cached—The percentage of requests for documents that were not in the cache (but are now) and their average transaction times. ▮ Server No Cache—The percentage of requests for documents that were not in the cache, but have server no-cache headers (cannot be cached); and their average transaction times. ▮ Stale Reloaded—The percentage of misses that are revalidated, turn out to be changed, reloaded, and served; and their average transaction times. ▮ Client No Cache—The percentage of misses with client no-cache headers and their average transaction times.

Statistics	Description (Continued)
Errors	▮ Connect Failures—The percentage of connect errors and their average transaction times. ▮ Other Errors—The percentage of other errors and their average transaction times.
Aborted Transactions	▮ Client Aborts—The percentage of client-aborted transactions, and their average transaction times. ▮ Questionable Client Aborts—The percentage of possibly client-aborted transactions, and their average transaction times. ▮ Partial Request Hangups—The percentage of early hangups (after partial requests) and their average transaction times. ▮ Pre-Request Hangups—The percentage of pre-request hangups and their average transaction times. ▮ Pre-Connect Hangups—The percentage of pre-connect hangups and their average transaction times.
Other transactions	*Unclassified—The percentage of unclassified transactions and their average transaction times.*

The following table describes the statistics for the HTTP protocol. Statistics exist for both the client and server.

	Statistics	Description
Client	*Total Document Bytes*	*The total amount of HTTP data served to clients since installation.*
	Total Header Bytes	*The total amount of HTTP header data served to clients since installation.*
	Total Connections	*The total number of HTTP client connections since installation.*
	Transactions In Progress	*The total number of HTTP client transactions in progress.*
Server	*Total Document Bytes*	*The total amount of HTTP data received from origin servers since installation.*
	Total Header Bytes	*The total amount of HTTP header data received from origin servers since installation.*

	Total Connections	*The total number of HTTP server connections since installation.*
	Transactions In Progress	*The total number of HTTP server connections in progress.*

The following table describes the protocol for the NNTP protocol. Statistics and descriptions exist for Client, Server, and Operations.

	Statistics	**Description**
Client	*Open Connections*	*The number of open NNTP connections.*
	Bytes Read	*The number of NNTP client request bytes read since installation.*
	Bytes Written	*The number of NNTP client bytes written since installation.*
Server	*Open Connections*	*The number of currently open NNTP server connections.*
	Bytes Read	*The number of bytes read from parent NNTP servers since installation.*
	Bytes Written	*The number of NNTP bytes written to the cache since installation.*
	Statistics	**Description**
Operations	*Article Hits*	*The number of news article hits since installation.*
	Article Misses	*The number of news article misses since installation.*
	Overview Hits	*The number of overview hits since installation.*
	Overview Refreshes	*The number of overview refreshes. An overview refresh occurs when the appliance caches a group overview on demand (as opposed to an overview pull).*
	Group Hits	*The total number of news group hits.*
	Group Refreshes	*The total number of news group refreshes (updates).*
	Posts	*The number of posts through the traffic server.*

Post Bytes	*The number of total bytes posted through the traffic server.*
Poll Bytes	*The number of total bytes polled by the traffic server.*
Feed Bytes	*The number of total bytes fed to the traffic server.*

The following table describes the statistics for the FTP protocol:

Statistics	Description
Open Connections	*The number of open FTP connections.*
PASV Connections Successes	*The number of successful PASV connections since installation.*
PASV Connections Failures	*The number of PASV connection failures since installation.*
PORT Connections Successes	*The number of successful PORT connections since installation.*
PORT Connections Failures	*The number of PORT connection failures since installation.*

The following table describes the statistics for the ICP protocol. Statistics exist for queries originating from the node and for queries originating from ICP peers.

	Statistics	Description
Queries Originating from this Node	*Query Requests*	*The number of HTTP requests that generate ICP query messages.*
	Query Messages Sent	*The total number of ICP query messages sent to ICP peers. This number is larger than the number of ICP Query Requests if there are multiple ICP peers.*
	Peer Hit Messages Received	*The number of ICP peer hit messages received in response to ICP queries from this node.*

	Statistics	Description
Queries Originating from this Node (continued)	*Peer Miss Messages Received*	*The number of ICP peer miss messages received in response to ICP queries from this node.*

	Total Responses Received	*The number of response messages received from ICP peers (siblings and parents).*
	Average ICP Message Response Time	*The average time for an ICP peer to respond to an ICP query message from this node. This is a cumulative average value.*
	Average ICP Request Time	*The average time for an HTTP request (that is sent to ICP) to receive an ICP response. This is a cumulative average value.*
Queries Originating from ICP Peers	*Query Messages Received*	*The number of ICP query messages received from remote ICP peers (siblings and parents).*
	Remote Query Hits	*The number of successful cache lookups in response to queries from ICP peers.*
	Remote Query Misses	*The number of unsuccessful cache lookups in response to queries from ICP peers.*
	Successful Responses Sent to Peers	*The number of successful ICP messages written in response to ICP queries from remote ICP peers.*

Viewing Cache statistics

Cache statistics report information about the cache size, bytes used, object look-up operations, object reads, object writes, update operations, and remove operations.

▼ Viewing Cache statistics

1. Select the **monitor** menu, and press Enter.
2. Select **cache**, and press Enter. Doing so causes the statistics to display on the screen. The following table describes the statistics.

Statistics	Description
Cache Bytes Used	*The number of bytes currently used.*
Cache Size	*The number of bytes devoted to the cache.*
Cache Lookups Completed	*The number of completed cache lookups (for ICP hits) since installation.*
Cache Lookups Failed	*The number of ICP misses since installation.*
Cache Reads Completed	*The number of cache reads completed since installation (NNTP, HTTP, and FTP).*

Statistics (Continued)	Description (Continued)
Cache Reads Failed	*The number of cache read misses since installation (NNTP, HTTP, and FTP).*
Cache Writes Completed	*The number of completed cache writes since installation (NNTP, HTTP, and FTP).*
Cache Writes Failed	*The number of cache write failures since installation (NNTP, HTTP, and FTP).*
Cache Updates Completed	*The number of cache HTTP updates completed since installation.*
Cache Updates Failed	*The number of cache HTTP update failures since installation.*
Cache Removes Completed	*The number of cache removes completed since installation (includes NNTP, HTTP, and FTP removes).*
Cache Removes Failed	*The number of cache remove failures since installation (includes NNTP, HTTP, and FTP removes).*

Viewing Other statistics

Other statistics report information about host database lookups, DNS lookups, cluster connections, and logging.

▼ Viewing host database statistics

1. Select the **monitor** menu, and press Enter.
2. Select **other**, and press Enter.
3. Select **hostdb**, and press Enter. Doing so causes the statistics to display on the screen. The following table describes the statistics.

Statistic	Description
Total Lookups	*The total number of lookups in the appliance host database since installation.*
Total Hits	*The total number of host database lookup hits since installation.*
Average TTL (min)	*The average time-to-live in minutes.*

▼ Viewing DNS statistics

1. Select the **monitor** menu, and press Enter.
2. Select **other**, and press Enter.
3. Select **dns**, and press Enter. Doing so causes the statistics to display on the screen. The following table describes the statistics.

Statistic	Description
Total Lookups	*The total number of DNS lookups (queries to name servers) since installation.*
Successes	*The total number of DNS lookup successes since installation.*
Average Lookup Time (msec)	*The average DNS lookup time.*

▼ Viewing cluster statistics

1. Select the **monitor** menu, and press Enter.
2. Select **other**, and press Enter.
3. Select **cluster**, and press Enter. Doing so causes the statistics to display on the screen. The following table describes the statistics.

Statistic	Description
Bytes Read	*The number of bytes read by this node from other nodes in the cluster since installation.*
Bytes Written	*The number of bytes this node has written to other cluster nodes since installation.*
Connections Open	*The total number of intracluster connections opened since installation.*
Total Operations	*The total number of cluster transactions since installation.*
Network Backups	*The number of times this node encountered intracluster network congestion and reverted to proxy-only mode since installation.*
Clustering Nodes	*The number of clustering nodes.*

▼ Viewing logging statistics

1. Select the **monitor** menu, and press Enter.
2. Select **other**, and press Enter.
3. Select **logging**, and press Enter. Doing so causes the statistics to display on the screen. The following table describes the statistics.

Statistic	Description
Currently Open Log Files	*The number of access log files (formats) that are currently being written.*
Space Used For Log Files	*The current amount of space being used by the logging directory, which contains all of the access and error logs.*
Number of Access Events Logged	*The current number of access events that have been written to log files. This counter represents one entry in one file, so that if multiple formats are being written, a single access will create multiple-access event log entries.*
Number of Access Events Skipped	*The number of skipped access events.*
Number of Error Events Logged	*The current number of events that have been written to the access error log.*

Using the expert menu

The **expert** menu lets you invoke a command shell. From the shell, you can execute the following commands to access features not included in the command-line interface or the Manager UI: `date`, `ifconfig`, `iostat`, `ipnat`, `kill`, `last`, `less`, `ls`, `mpstat`, `netstat`, `ping`, `ps`, `print_bypass`, `pwd`, `snoop`, `tail`, `top`, `traceroute`, `traffic_line`, `vmstat`, and `who`.

Note For information on these UNIX commands, refer to Sun's Product Documentation on the World Wide Web by visiting http://docs.sun.com/.

▼ Entering expert mode

1. Select the **expert** menu, and press Enter. Doing so causes control to switch to the Unix operating system.

Note To return to the CLI, enter exit at the operating system's command-line prompt.

Using the save menu

The **save** menu lets you save the current appliance configuration to a floppy disk.

▼ Saving the current configuration to a floppy disk

1. Select the **save** menu, and press Enter. Doing so causes the system to prompt you to insert a blank floppy disk.
2. Insert a floppy disk into the floppy disk drive, and press Enter. Doing so causes the appliance to copy all the current configuration settings to the floppy disk. After the operation, a message displays on screen indicating the copy was successful and asks you to take the floppy out.

Note Do not leave the floppy inside the drive after saving the configuration. Doing so will cause remote boot and halt operations to malfunction.

Using the load menu

The **load** menu lets you copy a previously saved appliance configuration file from a floppy disk.

▼ Loading a previously saved configuration from a floppy

1. Select the **load** menu and press Enter. Doing so causes the system to prompt you to insert a floppy disk into the drive.
2. Insert the floppy disk containing a previously saved configuration in the floppy disk drive, and press Enter. Doing so causes the appliance to copy the configuration from the floppy disk and load it on the system. After the operation, a message appears on the screen indicating the copy was successful.

Using the logoff menu

The **logoff** menu disconnects you from the appliance and logs you out of the system.

▼ Logging off the system

1. Select the **logoff** menu and press Enter. Doing so causes the system to disconnect you and return control to the VT100 terminal emulator window.

Chapter 6

Troubleshooting Problems

When the system doesn't seem to be operating correctly, you can use the information in this chapter to help you find a solution. If the information in this chapter doesn't solve your problem, refer to the *Intel NetStructure Caching Appliance Product Support* booklet that came with your system.

This chapter provides information on the following topics:

Rebooting your system

Rebooting the Intel NetStructure Cache Appliance causes the underlying operating system to reboot. Rebooting the appliance is not the same as starting and stopping the caching software on your system. For instructions on how to start and stop the caching software by using the command-line interface (CLI), refer to *Starting the appliance, on page 55* and *Stopping the appliance, on page 55*. For information on how to start or stop the caching software by using the Manager UI, refer to *Using the Server Basics page, on page 24*.

You can reboot the appliance in a controlled manner through the CLI. If you find the appliance in a state where you can't reboot it in a controlled manner, you can reboot it by pressing the Reset button located on the front panel. You should use this reboot method as a last resort. (For exact location of the Reset button, refer to the *Intel NetStructure Cache Appliance Quick Start Guide*.)

Note During a reboot operation, the system maintains the state of the caching software. For example, if the caching software is running when the reboot operation is initiated, they will still be running after the reboot. On the other hand, if the caching software is not running at the time of the reboot, it will remain off after the reboot. However, during the system reboot, caching operations cease regardless of whether or not the caching software is running at the time of the reboot.

Rebooting your system from the CLI

You can reboot the appliance from the command-line interface.

▼ Rebooting the appliance from the CLI

1. Select the **expert** menu, and press Enter.
2. Select **reboot**, and press Enter.
3. Wait approximately four to five minutes for the appliance's operating system to properly shut down and then restart.

Rebooting your system from the front panel

You can reboot the appliance from the front panel.

▼ Rebooting the appliance from the front panel

1. Press the system's Reset button. (For exact location of the Reset button, refer to the *Intel NetStructure Cache Appliance Quick Start Guide*.)
2. Wait approximately two to three minutes for the appliance's operating system to shut down and then restart.

Upgrading software

Periodically the caching application that runs on the Intel NetStructure Cache Appliance might need upgrading or might need to have a patch applied. In this case, visit Intel's ISP web site at http://www.intel.com/isp and go to the product page for your appliance. That page contains information on the latest software versions and patches that might apply.

Appendix A

Caching Solutions and Performance

This appendix is an overview of the Web caching capabilities and performance of the Intel NetStructure Cache Appliance.

This chapter covers the following topics.

Web proxy caching

The Intel NetStructure Cache Appliance is a high-performance caching proxy server. It is designed to efficiently handle multiple client connections simultaneously and supports HTTP, FTP, NNTP, ICP, and WCCP 2.0 protocols.

Caching

The idea behind Web caching is to store copies of frequently accessed documents close to users and serve this information to them on demand. Users get their information faster, and Internet bandwidth is freed up for other tasks.

Proxy server

Users direct their requests to Web servers all over the Internet. For a caching server to serve these requests, it must act as a *Web proxy server*. A Web proxy server fields user requests to arbitrary Web servers and either serves the requests, or forwards them on to the *origin server* (the Web server that contains the original copy of the requested information).

Transparent and explicit proxy caching

The proxy supports both *transparent proxy caching*, where the user's client software is unaware that it is communicating with a proxy, and *explicit proxy caching*, where client software (typically a browser) must be expressly pointed at the proxy. Transparent proxy caching is discussed in more detail on *page 120*.

A day in the life of a cache request

Here is an overview of the steps that take place as the appliance acts as a proxy cache and serves a user request.

Step 1

The appliance receives a user request for a document, image, news article, or other Web object.

Step 2

With the object address in hand, the appliance looks up the requested object in its object database (cache).

Step 3

If the object is in the cache, the appliance checks to see if the object is fresh enough to serve. (See *Ensuring cached object freshness, on page 115* for details.) If the object is fresh, the appliance serves it to the user as a *cache hit* (*Figure 1*).

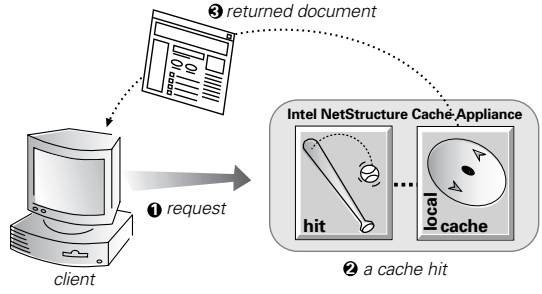

Figure 1 A cache hit

Step 4 If the data in the cache is stale, the appliance connects to the origin server and asks if the document is still fresh. If the document is still fresh, the appliance sends the cached copy to the user immediately.

Step 5 If the object is not in the cache (a *cache miss*) or the server indicates that the cached copy is no longer valid, the appliance gets the document from the Web server, simultaneously streaming it to the user and the cache (*Figure 2*). Subsequent requests for the object will be served faster.

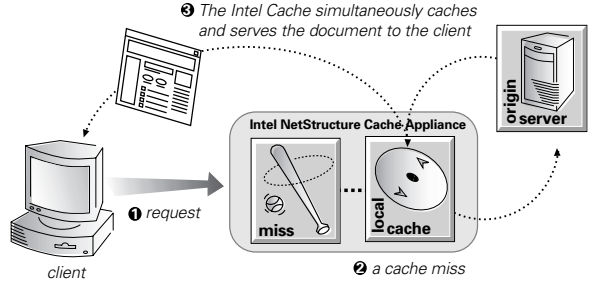

Figure 2 A cache miss

Caching is more complex than the preceding overview suggests. In particular, the overview does not answer these questions:

- ✔ How does the Intel NetStructure Cache Appliance ensure freshness given the different protocols it supports?
- ✔ How does the appliance revalidate stale HTTP objects?
- ✔ How does the appliance test an HTTP object for freshness?
- ✔ How does the appliance decide to serve an HTTP object?
- ✔ How do you configure the appliance's HTTP freshness options?
- ✔ How does the appliance serve correct HTTP alternates?
- ✔ How does the appliance treat requests for objects that cannot or should not be cached?

The following sections discuss these questions.

Ensuring cached object freshness

The Intel NetStructure Cache Appliance handles object freshness differently depending on protocol.

FTP FTP documents stay in the cache for a time period specified by the system administrator. See *Freshness, on page 36*.

NNTP News articles are refreshed each time the appliance polls parent news servers for changes in group lists, article overview lists, and article updates. See *Maintaining the cache: updates and feeds, on page 141*.

HTTP Web documents support optional author-specified expiration dates. The appliance adheres to these expiration dates; otherwise it picks an expiration date based on how frequently the document is changing and on administrator-chosen freshness guidelines. In addition, documents can be revalidated, checking with the server if a document is still fresh.

Revalidating objects

If an HTTP object is stale, the Intel NetStructure Cache Appliance *revalidates* the object. A revalidation is a query to the origin server that asks if the object is unchanged. The result of a revalidation could be:

- ✔ The object is still fresh; the appliance resets its freshness limit and serves the object.
- ✔ A new copy of the object is available; the appliance caches the new object, replacing the stale copy, and serves the object to the user simultaneously.
- ✔ The object no longer exists on the origin server; the appliance does not serve the cached copy.
- ✔ The origin server does not respond to the revalidation query. The appliance serves the stale object along with a `111 Revalidation Failed` warning.

HTTP object freshness tests

Here's how the Intel NetStructure Cache Appliance determines an HTTP document's freshness:

Expires header test:

Some documents come with `Expires` headers or `max-age` headers that explicitly define how long the document can be cached. A simple comparison of the current time with the expiration time determines whether or not the document is fresh.

Last-Modified / Date header test:

If no expiration information exists, the appliance can use the `Last-Modified` and `Date` headers to estimate a freshness limit. The `Last-Modified` header indicates how long ago a document was modified. If a document was last modified two years ago, it is unlikely to suddenly change, so the appliance can cache it safely for a while. But if the document just changed five minutes ago, it might be quite volatile, and the appliance should not cache it very long. The appliance stores an object for some percentage of the time (F) that elapsed since the object last changed. The percentage is 10% by default:

```
freshness limit = F * (Date - Last-Modified)
```

In the above formula, the `Date` header provides the date the object was sent to the appliance and the `Last-Modified` header provides the date the object was last modified on the origin server.

For example, if a document was last modified 32 days ago and was sent to the appliance two days ago, the document is considered fresh in cache for three days after it was sent. (This assumes a factor of 10%.) So for this situation, the document is considered fresh for one more day.

Because this method could result in lengthy freshness times for documents that have not changed for long periods, cache administrators might want to place an upper boundary on the freshness limit. With this boundary in place the freshness limit becomes the smaller of the two values: the boundary or the computed freshness limit. For information on how to configure an upper boundary, refer to the Freshness section of the **Configure: Cache** page of the Manager UI. See *Freshness, on page 36*.

Default test:

For documents that do not have `Expires` headers or do not have both `Last-Modified` and `Date` headers, you can specify an absolute freshness limit in the Freshness section of the **Configure: Cache** page. See *Freshness, on page 36*.

Revalidate rules:

Revalidate rules apply specific freshness limits to specific HTTP or FTP objects. From the command-line interface, you can set freshness limits for objects originating from particular domains or IP addresses, objects with URL addresses that contain specified regular expressions, and objects requested by particular clients. See *Configuring caching rules, on page 79*.

Deciding whether to serve HTTP objects

Even though a document might be fresh in the cache, clients or servers could have constraints that prevent them from retrieving the document from the cache. For example, a client might request that a document not come from a cache, or if it does, the document cannot have been cached for more than 10 minutes.

The Intel NetStructure Cache Appliance bases the servability of a cached document on `Cache-Control` header fields. These headers can appear in both client requests and server responses.

The following cache-control header fields affect whether objects are served:

- ✔ The `no-cache` field, sent by clients, tells the appliance to serve *no* objects directly from the cache; always revalidate. You can configure the appliance to ignore client no-cache fields. See *Cache activation, on page 35*.
- ✔ The `max-age` field, sent by servers, is compared to the document age; if the age is less than the `max-age`, the document is fresh and can be served.

- ✔ The `min-fresh` field, sent by clients, is an *acceptable freshness tolerance*. The client wants the object to be at least this fresh. If a cached document does not remain fresh at least this long in the future, it is revalidated.
- ✔ The `max-stale` field, sent by clients, permits the appliance to serve stale documents provided they are not too old. Some browsers might be willing to take stale documents in exchange for improved performance, especially during periods of poor Internet availability.

The appliance applies `Cache-Control` servability criteria *after* HTTP freshness criteria. For example, a document might be considered fresh, but if its age is greater than its `max-age`, it is not served.

Configuring HTTP freshness options

You can configure the following freshness guidelines for the Intel NetStructure Cache Appliance:

- ✔ How often to revalidate (when to consider objects stale). See *Configuring HTTP revalidation* below.
- ✔ Whether to cache documents without freshness information. See *Configuring HTTP cachability* below.
- ✔ Whether to use an upper boundary to determine if the Last-Modified / Date freshness limit is too long.
- ✔ What absolute freshness lifetime to use when estimating the freshness of documents without `Expires` or `Last-Modified` headers.

See *Freshness, on page 36* for instructions.

Configuring HTTP revalidation

The following HTTP revalidation options are available:

- ✔ Always revalidate (everything is considered stale).
- ✔ Never revalidate (everything is considered fresh).
- ✔ Revalidate all objects without `Expires` headers. Evaluate the freshness of objects with `Expires` headers by first checking the `Expires` header, and then checking `Cache-Control` headers.

✔ Evaluate freshness as follows:

1. Use the `Expires` header test, if applicable, otherwise go to step 2. If the object is stale, revalidate. If it is fresh, check the `Cache-Control` headers.
2. Use the `Last-Modified` / `Date` header test, if applicable, otherwise go to step 3. If the object is fresh according to the `Last-Modified` / `Date` test, check the `Cache-Control` headers for any freshness restrictions.
3. Use the absolute freshness limit specified in the Freshness section of the **Configure: Cache** page. Revalidate if the age is past the freshness limit.

Configuring HTTP cachability

The following HTTP cachability options are available:

✔ Cache only documents that have `Expires` headers

✔ Cache only documents that have `Expires` or `Last-Modified` headers

✔ Do not restrict caching

Caching HTTP alternates

Some Web servers answer requests for the same URL by serving a variety of objects. The content of these objects can vary widely, according to whether a server delivers content for different languages, targets different browsers with different presentation styles, or delivers variable content at different times of the day. Different versions of the same object are termed *alternates*.

Header information identifies alternates. You can configure the Intel NetStructure Cache Appliance to cache all alternates according to a particular header. For example, if you tell the appliance to vary on the `User-Agent` header, the appliance caches all the different user-agent versions of documents it encounters. To configure caching of alternates, see *Variable content, on page 38*.

To cache or not to cache?

Depending on the type of object, you can direct the Intel NetStructure Cache Appliance to cache or not cache an object:

NNTP You can limit article caching to specific news groups. See *Blocking particular groups, on page 140*.

FTP You can configure never-cache rules for specific types of FTP documents by using the command-line interface. See *Configuring caching rules, on page 79*.

HTTP The appliance responds to caching directives from clients and origin servers, as well as configured options in the Manager UI and the command-line interface.

The following table lists the HTTP caching directives that the appliance follows.

Directive source	Caching directives
administration options	▮ Don't cache objects with URL addresses containing `?`, `;`, `/cgi` or end in `.asp`. ▮ Don't cache objects served in response to the Cookie: header. ▮ Set never-cache rules from the command-line interface. Refer to *Configuring caching rules, on page 79*.
client	*Don't cache objects with the following request headers. You can override some of these directives using administration options.* ▮ Cache-Control: no-store header ▮ Cookie: header ▮ Authorization: header
Web server	*Don't cache objects with the following response headers. You can override some of these directives using administration options.* ▮ Cache-Control: no-store ▮ www-Authenticate: header ▮ Set-Cookie: header ▮ Cache-Control: no-cache header ▮ Pragma: no-cache header ▮ Expires: header with value of 0 (zero) or a past date

Transparent proxy caching

In nontransparent proxy caching, client browsers must be configured to send Web requests to the Intel NetStructure Cache Appliance proxy. Many sites have no direct control over user browser settings, making it necessary for site administrators to tell users to configure their browsers to direct requests to the proxy.

Transparency solves this problem. The transparency option enables the appliance to respond to Internet requests without requiring users to reconfigure their browser settings. It does this by redirecting the traffic flow into the cache after it has been intercepted by an L4 switch or router.

This section provides the following:

✔ An overview of how the appliance serves requests transparently. See *Serving requests transparently, on page 121*.

✔ A discussion of interception strategies supported by the Intel NetWorking Cache Appliance. See *Interception strategies, on page 121*.

- ✔ Information on how the ARM changes packet addresses. See *ARM redirection, on page 125*.
- ✔ A description of the appliance's adaptive bypass scheme. See *Appliance adaptive bypass, on page 126*.

Serving requests transparently

Here's how the Intel NetStructure Cache Appliance transparent interception works:

Step 1 The appliance intercepts client requests to origin servers. Several appliance deployment methods exist so that interception can take place. See *Interception strategies, on page 121* for details.

Step 2 The Adaptive Redirection Module (ARM) redirects requests destined for origin servers to the appliance application. See *ARM redirection, on page 125* for details.

Step 3 A very small number of clients and servers do not work correctly through proxies. The appliance identifies these problem clients and servers dynamically, and the ARM adaptively disables interception for these clients and servers, passing their traffic unimpeded to the original server. Also, clients and servers can be manually exempted from caching by configuring the ARM. See *Adaptive interception bypass, on page 126* for more information.

Step 4 The appliance receives and begins processing the intercepted client requests as usual. If a request is a cache hit, the appliance serves the requested document or news article. If a request is a miss, the appliance retrieves the document from the origin server and serves it to the client.

Step 5 On the way back to the client, the ARM changes the source IP address to the origin server IP address and the source port to the origin server port.

Interception strategies

The transparency routing solutions supported by the Intel NetStructure Cache Appliance are:

- ✔ Layer 4-aware switch. See *Using a layer 4-aware switch to filter transparency requests, on page 122*.
- ✔ Cisco IOS-based router using the Web Cache Control Protocol (WCCP). See *Using a WCCP-enabled router for transparency, on page 123*.
- ✔ Policy-based routing. See *Using policy-based routing to filter transparency requests, on page 124*.

How client request traffic reaches the appliance depends on network topology. In a complex network, you must decide which clients are to be served transparently and make sure that the appliance is positioned to intercept their requests. The

appliance, or routers or switches feeding it, is often deployed at a major artery or aggregation pipe to the Internet.

The following sections provide more details about the Intel NetStructure Cache Appliance's transparency routing solutions.

Using a layer 4-aware switch to filter transparency requests

Layer 4-aware switches can rapidly redirect supported protocols to the Intel NetStructure Cache Appliance, while passing all other Internet traffic through directly to its destination. *Figure 3* illustrates this scenario for HTTP.

Layer 4-aware switches offer the following features, depending on the particular switch:

- ✔ A layer 4-aware switch can sense downed hosts on the network and redirect traffic.
- ✔ Single layer 4-aware switches that feed several appliances balance loads among the nodes. Different switches might use different load-balancing methods, such as round-robin or hashing. If a node goes down, the switch automatically redistributes the load. When the node returns to service, some switches automatically return the node to its previous workload, so that the node cache need not be repopulated; this feature is called *cache affinity.* Intel recommends that you do *not* enable the virtual IP failover in this situation, because layer 4-aware switch failover is already in operation.

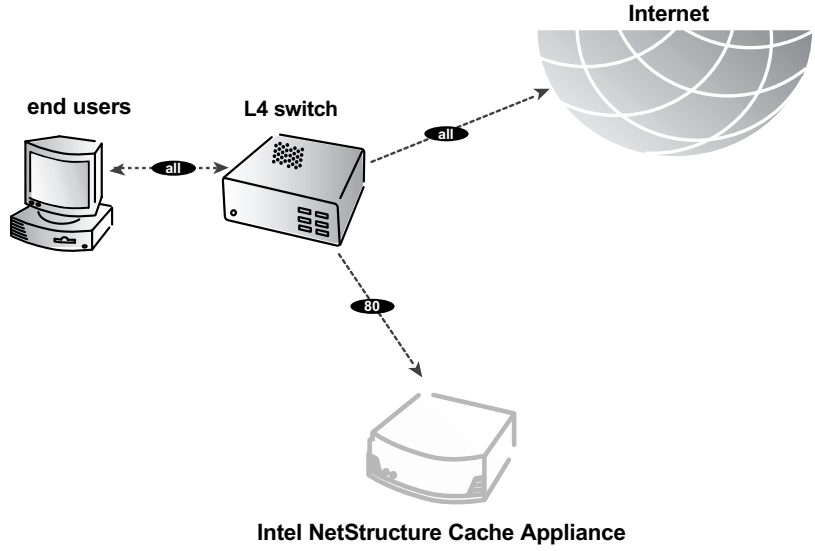

Figure 3 Using a layer 4-aware switch to filter HTTP requests

Using a WCCP-enabled router for transparency

A WCCP 2.0-enabled router can send all port 80 (HTTP) traffic to the Intel NetStructure Cache Appliance, as shown in *Figure 4*. After the WCCP router sends port 80 traffic, the ARM readdresses port 80 to the appliance proxy port (by default, port 8080). Then the appliance processes the request as usual, retrieving the requested document from the cache if it is a hit and sending the response back to the client. Along the way, the ARM readdresses the proxy port in the response header to port 80 (undoing the readdressing it did on the way to the appliance). The user then sees the response exactly as if it were sent directly from the origin server. In addition to port 80 (HTTP) traffic, WCCP 2.0 supports more protocols including NNTP (port 119 traffic).

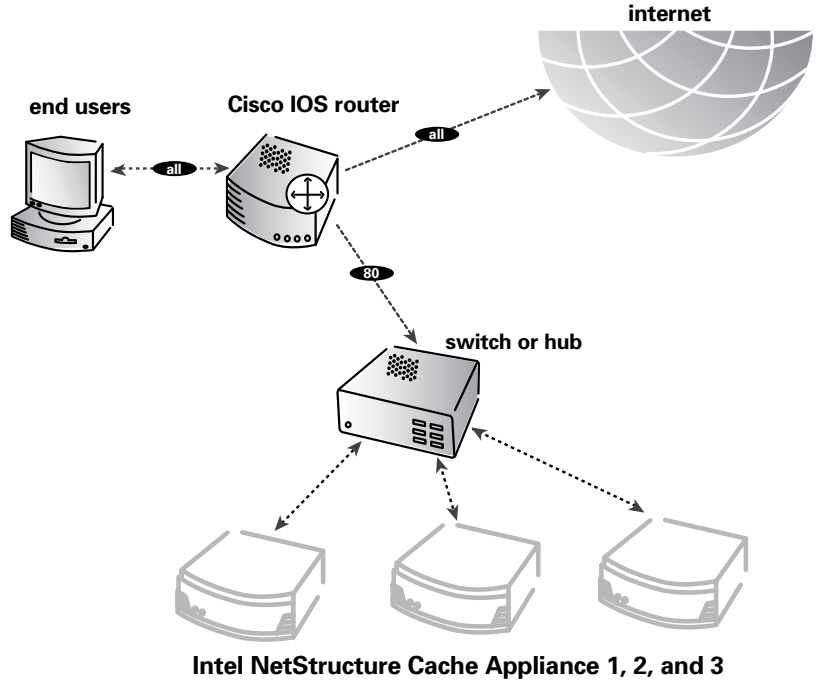

Figure 4 *Using a Cisco IOS router to send port 80 traffic to several Intel NetStructure Cache Appliances*

WCCP provides the following routing benefits:

- ✔ The WCCP-enabled router and the appliance exchange heartbeat messages, letting each other know they are running. The WCCP router automatically reroutes port 80 and port 119 traffic if the appliance goes down.
- ✔ If several appliances receive traffic from a WCCP router, WCCP balances the load among them. The group of appliances is called a *WCCP cache farm*.

- ✔ The appliance handles node failure in WCCP cache farms. If one node goes down, its load is redistributed among the remaining nodes.
- ✔ In WCCP, you can use multiple routers. Traffic flowing through multiple routers can share the same pool of caches.

In *Figure 4*, appliances 1, 2, and 3 form a WCCP cache farm.

If the appliance in the WCCP-enabled routing scheme has an ARM bypass rule, the rule causes the appliance to forward particular client requests directly to the origin server, bypassing the appliance. Bypassed requests are unchanged by the ARM; they retain their client source IP addresses. See *Adaptive interception bypass, on page 126* for details.

In WCCP 2.0, you can exclude certain router interfaces from redirection. The appliance bypass rules can work if you exclude the router interface on which it is connected from using WCCP. To do so, set the router configuration command `ip wccp redirect exclude in` (refer to Cisco's WCCP documentation for information about router configuration).

If a WCCP router serves several nodes, as in *Figure 4*, the router balances their loads. The router sends each node requests aimed at a particular range of IP addresses, so that each node is responsible for caching content residing at particular IP addresses.

The appliance also supports cache affinity. If a node fails and then restarts, the appliance returns the node to its former load distribution. The node's cache need not be repopulated.

The WCCP cache farm acts as a simple form of distributed cache. A WCCP-enabled network device distributes traffic to individual appliances based on the IP address of the destination Web server. Each node caches objects requested from a particular set of Web servers, which belong to that node's assigned range of destination IP addresses.

Virtual IP failover not recommended

If you are running clustered appliances, Intel recommends that you do *not* enable virtual IP failover in WCCP environments. The appliance's WCCP failover mechanism handles node failures and restarts. See *Virtual IP failover, on page 146* for details about virtual IP failover.

Using policy-based routing to filter transparency requests

Instead of the WCCP protocol, you can use the policy-routing capabilities of a router to send traffic to the Intel NetStructure Cache Appliance. WCCP or an L4 switch is generally preferable to policy-based routing because it has a performance impact on the router and does not support load balancing or heartbeat messaging.

Figure 5 illustrates policy-based routing for HTTP objects. This routing scheme has the following characteristics:

- ✔ All client Internet traffic is sent to a router that feeds the appliance.
- ✔ The router sends port 80 (HTTP) traffic to the appliance and sends the remaining traffic to the next hop router.
- ✔ The ARM translates intercepted requests into appliance requests so they can be served.
- ✔ Translated requests are sent to the appliance.
- ✔ Web documents to be served transparently are readdressed by the ARM on the return path to the client, so that the documents appear to have come straight from the origin server.

An appliance cluster with virtual IP failover adds reliability; if one node fails, another node can take up its transparency requests. See *Virtual IP failover, on page 146*.

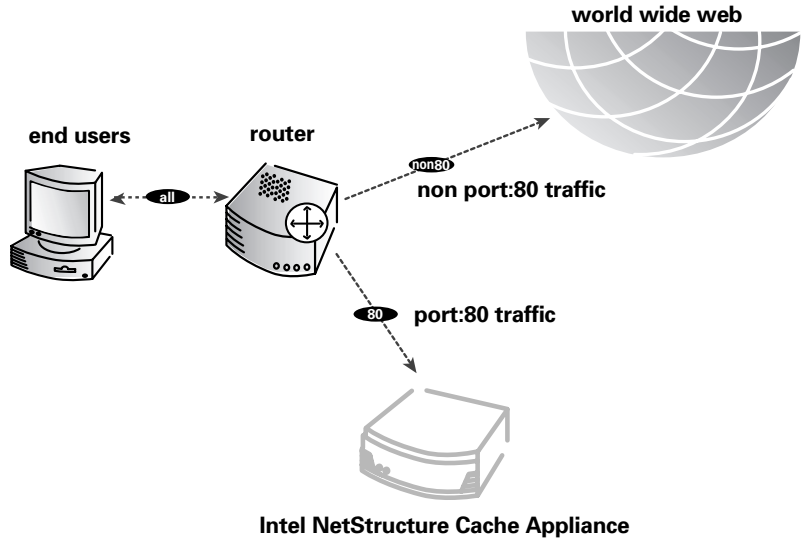

Figure 5 Using a router to filter HTTP requests

ARM redirection

The ARM can make two changes to an incoming packet's address: its destination IP address and its destination port.

- ✔ Typically, HTTP packet destination IPs and ports are readdressed with the IP address of the Intel NetStructure Cache Appliance and the appliance's HTTP proxy port (usually port 8080).

- ✔ NNTP packet destination IPs are readdressed with the IP address of the appliance. If the appliance uses a port other than 119 for NNTP, the destination NNTP port is readdressed as well.

Adaptive interception bypass

The Intel NetStructure Cache Appliance contains an adaptive learning module that recognizes inter operability problems caused by transparent proxy caching and automatically bypasses the traffic around the proxy without operator intervention.

Web proxies are very common in corporate and Internet use, so the frequency of inter operability problems is extremely rare. However, when problems do exist, the reasons usually can be attributed to the following:

- ✔ Client software bugs (homegrown, noncommercial browsers).
- ✔ Server software bugs.
- ✔ Applications that send non-HTTP traffic over HTTP ports as a way of defeating security restrictions.
- ✔ Server IP authentication. In this case the Web server limits access to itself to a few client IP addresses. Since the appliance IP address is different it cannot get access to the server. A server limiting IP addresses is not infrequent. Limitations occur because ISPs dynamically allocate client IP dial-up addresses and use more secure cryptographic protocols.

Appliance adaptive bypass

The appliance watches for certain protocol inter operability errors, and as it detects errors, it configures the ARM to bypass the proxy for the clients and/or servers causing the errors.

In this way, the very small number of clients or servers that do not operate correctly through proxies are auto detected and routed around the proxy, so they can continue to function normally (but without the improvement of caching).

Dynamic rules are temporary

Dynamically generated bypass rules are purged after the appliance restarts.

More about bypass rules

The ARM can bypass the proxy based on the client IP address, the destination server IP address, or both.

You can manually configure bypass rules to direct requests from certain clients or to particular servers. For example, you might want client IP addresses that did not pay for a caching service to be steered around the cache, while paying-clients can obtain the benefits of caching. Or you can remove some servers from caching lists, because the servers don’t want to have their pages cached.

Static and dynamic (adaptive) bypass

Bypass rules can be either static or adaptive. Adaptive bypass rules are dynamically generated if you configure the appliance to bypass in the case of non-HTTP port 80 traffic or HTTP errors.

Static and dynamic rules

Static and dynamic rules look exactly the same. However, the appliance creates dynamic rules when it encounters particular problems, such as non-HTTP port 80 traffic or HTTP errors.

Configuring bypass options

You can bypass requests based on the following criteria:

- ✔ Requests from particular users (identified by source IP addresses); set static source bypass rules from the command-line interface
- ✔ Requests to particular Web sites (identified by destination IP addresses); set static destination bypass rules from the command-line interface
- ✔ Requests from specific sources to specific destinations; set static source/destination bypass rules from the command-line interface

Bypass rules fall into these categories:

Source bypass:
This rule tells the appliance to bypass a particular source IP address or range of IP addresses. For example, you can use this rule to bypass clients that want to opt out of a caching solution. Source bypass rules are not dynamically generated.

Destination bypass:
This rule tells the appliance to bypass a particular destination IP address or range of IP addresses. For example, these could be Web servers that use IP authentication based on the client's real IP address. Destination bypass rules can be dynamically generated.

Hit-rate impact

Destination bypass rules prevent the appliance from caching an entire site. You will experience hit rate impacts if the site you bypass is popular.

Source/destination pair bypass:
This rule tells the appliance to bypass requests that originate from the specified source to the specified destination. For example, you can route around specific client-server pairs that experience broken IP authentication or out-of-band HTTP traffic problems when cached.
Source/destination rules can be dynamically generated.

Reducing hit-rate impact

Source/destination bypass rules might be preferable to destination rules because they block a destination server only for those particular users that experience problems.

Server acceleration

In Web proxy caching, the Intel NetStructure Cache Appliance handles arbitrary Web requests to distant Web servers on behalf of a set of users. *Server acceleration* (also known as reverse proxy caching or virtual Web hosting) is slightly different. In server acceleration, the appliance *is* the Web server to which the user is trying to connect. The Web server host name resolves to the appliance, which is acting as the real Web server.

Having a fast, scalable, fault-tolerant appliance absorb the main Web server request traffic can improve the speed and quality of service of Web serving, reducing load and hot spots on the backup Web servers, while still maintaining the publishing environment available on the backup Web servers.

If the appliance has the desired object in cache, it serves the document quickly. If the document is not in cache, the appliance requests the document from another backup Web server that has all the content. A configuration table specifies which backup Web server has the required content.

A Web host can maintain a scalable appliance serving engine and maintain a set of low-cost, low-performance, less reliable PC Web servers as the backup servers. A single appliance can act as the virtual Web server for multiple backup Web servers, as shown in *Figure 6*.

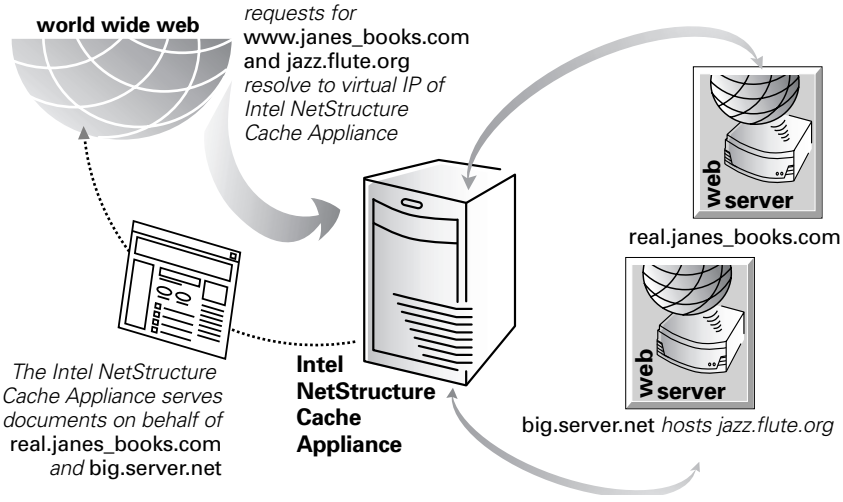

Figure 6 *Intel NetStructure Cache Appliances as server accelerator (reverse proxy) for a pair of Web servers*

Advantages of server acceleration

Server acceleration advantages are similar to Web proxy caching:

- ✔ The appliance is optimized for speed and multiple user connections and can be deployed close to users.
- ✔ Serving cached documents saves network bandwidth.

Server acceleration offers the following server advantages:

- ✔ Web servers can be off-loaded, providing overload insurance. An appliance cluster dynamically mirrors content from heavily loaded Web servers.
- ✔ Web administration is centralized. Administrators maintain the Web server(s) being accelerated, and the appliances do the job of distributing content.

Note Server acceleration described here applies to HTTP requests.

How server acceleration works

When a browser makes a request, it normally sends that request directly to the origin server. When the appliance is in reverse proxy mode, it must intercept the request for that origin server.

Interception occurs by setting up the DNS entry for the origin server (the origin server's *advertised* host name) to resolve to the appliance's IP address. If the appliance is clustered, using a virtual IP address provides added reliability (if a node fails, another node takes on the virtual IP address of the failed node).

When the appliance is set up as the Web server, the browser connects to it rather than to the origin server (see *Figure 6*). The origin server cannot have the same name as the advertised host name, or there would be a DNS conflict.

Retrieving requested documents

Because the appliance is advertised as the origin server, it needs to act as a Web server rather than a proxy server, meaning that it receives server requests, not proxy requests. In this case, the appliance constructs a proxy request from the server request and then satisfies the proxy request.

In HTTP, server requests differ from proxy requests. The main difference is that server requests don't specify the entire URL, just the path. A server request might look like this:

```
GET /index.html HTTP/1.0

HOST: real.janes_books.com
```

Whereas the corresponding proxy request would look like this:

```
GET http://real.janes_books.com/index.html HTTP/1.0

HOST: real.janes_books.com
```

The appliance can construct a proxy request from a server request by using the server information in the host header.

You might have noticed a small problem. The correct proxy request must contain the host name of the origin server, not the advertised host name that names servers associated to the appliance. The advertised host name is what appears in the host header. For example, for the origin server `real.janes_books.com` in *Figure 6*, the server request and host header would be:

```
GET /index.html HTTP/1.0

HOST: www.janes_books.com
```

And the correct proxy request should be:

```
GET http://real.janes_books.com/index.html HTTP/1.0

HOST: real.janes_books.com
```

Document routing rewrite rules

To translate `www.janes_books.com` to `real.janes_books.com`, the appliance needs a set of *document routing rewrite rules* by which it can refer to the full paths on the Web servers it is accelerating. These rules are stored in the `remap.config` file. In the preceding example, the rule to map `www.janes_books.com` to `real.janes_books.com` would be:

```
map www.janes_books.com real.janes_books.com
```

Two types of rules exist: *map* rules and *reverse-map* rules.

- ✔ Map rules specify the location of content that the appliance is accelerating; they enable the appliance to translate a URL requested by a client into one that represents the accelerated content.
- ✔ Reverse-map rules translate origin server redirects to clients. If an origin server sends a redirect response to a client, the appliance translates the redirect so that the client is redirected to the appliance, instead of being redirected around it. For more information on reverse-map rules, see *Web server redirects, on page 131*.

For detailed descriptions of both map rules and reverse-map rules, see *Understanding server acceleration mapping rules, on page 132* and *Examples of rules and translations, on page 133*.

The map rule for the other Web server illustrated in *Figure 6*, `big.server.net`, which hosts `jazz.flute.org`, might look as follows:

```
map jazz.flute.org big.server.net/jazz/
```

This map rule specifies the path `/jazz` for `jazz.flute.org` on the server `big.server.net`.

Generally, you use reverse proxy mode to support more than one origin server. In this case, all of the advertised host names resolve to the IP address or virtual IP address of the appliance (see *Figure 6*). Using host headers, the appliance is able to translate server requests for any number of servers into proxy requests for those servers.

If the appliance receives requests from older browsers that do not support host headers, then it can route these requests directly to a specific server, or send the browser to a URL containing information about the problem. See *Setting server accelerator options, on page 43*.

Web server redirects

Web servers often send redirect responses back to browsers. Redirects tell browsers to go to different pages. Web servers redirect for a variety of reasons. One reason is to balance server load. For instance, if a server is overloaded, it might redirect browsers to a less loaded server. Another reason might be when Web pages have moved to different locations. When the appliance is configured in server acceleration mode, it must readdress redirects from origin servers so that browsers are redirected to the appliance, not to another Web server.

To readdress redirects, the appliance uses reverse-map rules. For example, the reverse-map rule required to convert redirects from `real.janes_books.com` (if the appliance assumes the associated name `www.janes_books.com`) would be:

```
reverse_map real.janes_books.com www.janes_books.com
```

In general, when setting up document rewrite rules, each map rule should have one reverse-map rule, with the source URL and the destination URL of the map rule reversed in the reverse-map rule.

You create and modify document reverse-map rules from the Server Acceleration section of the **Routing** page. See *Setting server accelerator options, on page 43* for more information on how to create reverse-map rules. For more information about how reverse-map rules work, see the following section, *Examples of rules and translations*.

Understanding server acceleration mapping rules

Rewrite rules each consist of three space-delimited fields: `type`, `target`, and `replacement`.

- ✔ `Type` indicates the type of rule.
- ✔ `Target` specifies the URL from which the request originates.
- ✔ `Replacement` specifies the URL the appliance uses in place of the target URL.

Using map rules

When the appliance receives a request as a server accelerator, it first constructs a complete request URL from the relative URL and its headers. The appliance then compares the complete request URL with its list of target (*from*) URL addresses, looking for a match. For the request URL to match a target URL, the following conditions must be true:

- ✔ The scheme of both URL addresses must be the same.
- ✔ The host in both URL addresses must be the same.

  If the request URL contains an unqualified hostname, it will never match a target URL with a fully qualified host name.
- ✔ The ports in both URL addresses must be the same.

  If no port is specified in a URL, the default port for the scheme of the URL is used.
- ✔ The path portion of the target URL must match a prefix of the request URL.

If the appliance finds a match, it translates the request URL into the replacement URL in the rule. It sets the host and path of the request URL to match the replacement URL. The appliance removes the prefix of the path that matched the target URL and substitutes for it the path from the replacement URL.

Note Cross-scheme mappings are not permitted. For example, you cannot map HTTP requests to FTP replacements.

Using reverse-map rules

Reverse mappings rewrite the location headers in origin server responses instead of the headers in the user agent requests. Origin servers use location headers to redirect clients to another location.

For example if there is a directory `/pub` on an origin server at `www.molasses.com`, and a user agent sends the request to that server for `/pub`, the server will probably reply with a redirect to `http://www.test.com/pub/` to let the client know that it was a directory it had requested, instead of a document. (A common use of redirects is to normalize URL addresses so that clients can bookmark documents properly.)

The appliance uses reverse mappings to prevent redirects from origin servers from causing clients to bypass the appliance in favor of direct access to the origin servers.

In a typical Server Accelerator configuration, there should be a reverse-map rule for every map rule, with the origin URL and replacement URL of the map rule reversed.

Examples of rules and translations

The following examples illustrate several important cases of rewrite rules.

Example 1 This map rule does not specify a path prefix in the target or replacement:

```
map http://www.x.com/ http:/server.hoster.com/x/
```

This rule results in the following translations:

User Request	Translated Request
http://www.x.com/Widgets/index.html	*http://server.hoster.com/x/Widgets/index.html*
http://www.x.com/cgi/form/submit.sh?arg=true	http://server.hoster.com/x/cgi/form.submit.sh?arg=true

Example 2 Map rules with path prefixes specified in the target:

```
map http://www.y.com/marketing http://marketing.y.com/

map http://www.y.com/sales http://sales.y.com

map http://www.y.com/engineering http://engineering.y.com/

map http://www.y.com/ http://info.y.com/
```

These rules result in the following translations:

User Request	Translated Request
http://www.y.com/marketing/projects/manhattan/specs.html	*http://marketing.y.com/projects/manhattan/specs.html*
http://www.y.com/marketing/projects/boston/specs.html	*http://info.y.com/marketing/projects/boston/specs.html*
http://www.y.com/engineering/marketing/requirements.html	*http://engineering.y.com/marketing/requirements.html*

Example 3 The order of the rules matters:

```
map http://www.g.com/ http://external.g.com/

map http://www.g.com/stuff http://stuff.g.com
```

These rules result in the following translation:

User Request	Translated Request
http://www.g.com/stuff/a.gif	*http://external.g.com/stuff/a.gif*

In these examples, the second rule is never applied because all URL addresses that match the second rule also match the first rule. The first rule takes precedence because it appears earlier in the `remap.config` file.

Example 4 A mapping with a path prefix specified in the target and replacement:

```
map http://www.h.com/a/b http://server.h.com/customers/x/y
```

This rule results in the following translation:

User Request	Translated Request
http://www.h.com/a/b/c/d/doc.html	*http://server.h.com/customers/x/y/c/d/doc.html*
http://www.h.com/a/index.html	*Translation fails*

Example 5 Reverse mapping:

```
map http://www.x.com/ http://server.hoster.com/x/
```

```
reverse_map http://server.hoster.com/x/ http://www.x.com/
```

These rules result in the following translations:

User Request	Translated Request
http://www.x.com/Widgets	*http://server.hoster.com/x/Widgets*

User Request	Origin Server Header	Translated Header
http://www.x.com/Widgets	*http://server.hoster.com/x/Widgets/*	*http://www.x.com/Widgets/*

For browsers that do not support host headers

When accelerating multiple servers, the appliance is unable to route to URL addresses from older browsers that do not send the `Host:` header. The best solution is to direct the user to a page that explains the situation and advises a browser upgrade or provides a link directly to the origin server, bypassing the appliance. For information on how to do this, see *Setting server accelerator options, on page 43*.

Understanding cache hierarchies

Cache hierarchies consist of levels of caches that communicate with each other. Hierarchical caching can give you information about the local access requirements of your users; this information might not appear in a large central cache. The Intel NetStructure Cache Appliance supports several types of cache hierarchies, but all cache hierarchies recognize the concepts of parent and child caches.

In a cache hierarchy a parent cache is a cache higher up, to which the appliance can forward requests. A child cache is a cache lower down for which the appliance is a parent.

In the event of a cache miss, instead of forwarding the request to a distant origin server, it might be faster to try another nearby cache in the hierarchy. If a forwarded request is a miss on the parent cache, the parent cache forwards the request to the origin server. See *Figure 7, on page 136* for an illustration. The appliance supports multiple parent caches; if a request misses on all parents, the appliance chooses a specific parent to forward the request to the origin server.

The Intel NetStructure Cache Appliance can function as a member of the following cache hierarchies:

- ✔ HTTP cache hierarchy
- ✔ ICP (Internet Cache Protocol) hierarchy
- ✔ NNTP hierarchy

The following sections describe these cache hierarchies.

HTTP cache hierarchies

The Intel NetStructure Cache Appliance supports HTTP cache hierarchies, using other Intel NetStructure Cache Appliances or even other caching products as parents or children in a chain of interdependent caches.You can create small, regional caches (for an organizational department or for users in a defined geographic area), and link them to larger parent caches, defining larger areas.

If a regional cache does not have a requested document (a cache miss) and HTTP parent caching is enabled, the appliance forwards the HTTP request to a parent cache in the hierarchy rather than contacting the origin server. If the parent cache (or caches) cannot serve the object they can forward the request to other caches further up in the hierarchy.

The appliance supports multiple HTTP parent caches and parent failover. This feature gives the appliance a sequence of parent caches to query if the first parent cache misses.

For information on how to enable parent caching from the Manager UI, see the parent caching section on the **Configure: Routing** page (see *Setting HTTP parent caching options, on page 40*). For information on how to enable parent failover using the command-line interface, see *Controlling parent proxy caching, on page 89*.

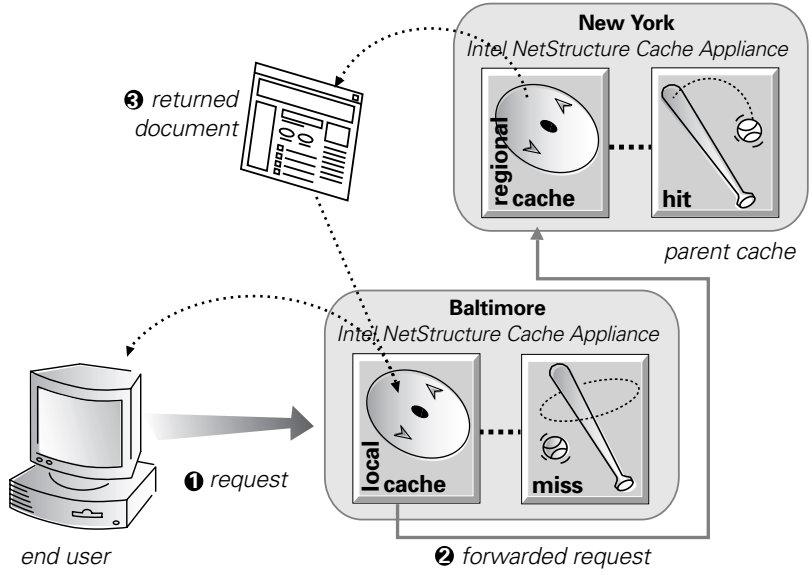

Figure 7 A cache hierarchy in action

ICP cache hierarchies

Internet Cache Protocol (ICP) is a protocol for proxy caches to exchange information about their content. ICP query messages ask other caches if they are storing a particular URL. ICP response messages reply with a hit or miss answer.

Peer, sibling, and parent caches

ICP hierarchies employ sibling caches as well as parent caches. Sibling caches exist at the same hierarchical level, while parent caches exist one level up in the hierarchy. A cache exchanges ICP messages only with specific ICP peers. An ICP peer can be a sibling cache or a parent cache.

If the Intel NetStructure Cache Appliance has ICP enabled, it sends out ICP queries to its sibling caches in the event of a cache miss on an HTTP request. If there are no hits on siblings, the appliance sends ICP queries to ICP parents. If there are no hits on ICP parents, the appliance forwards the request to its HTTP parents. If there are no HTTP parent caches established, the appliance forwards the request to a selected ICP parent cache (which resolves the request by communicating with the origin server).

How an ICP hit can be a miss

If the appliance receives a hit message from an ICP peer, then it sends the HTTP request to that peer. It might turn out to be an actual miss, because the original HTTP request contains header information that is not communicated by the ICP query. For example, the *hit* might not be the requested alternate. If an ICP hit turns out to be a miss, the appliance forwards the request to either its HTTP parent caches or to the origin server.

For information on now to enable and configure ICP options using the Manager UI, see the ICP section of the **Configure: Routing** page (see *Setting ICP options, on page 41*). For information on how to configure ICP options using the command-line interface see *Configuring and maintaining ICP peers, on page 84*.

NNTP cache hierarchies

Using an Intel NetStructure Cache Appliance as parent to another group of appliances can reduce load on a parent news server and take advantage of the large number of concurrent connections that server supports.

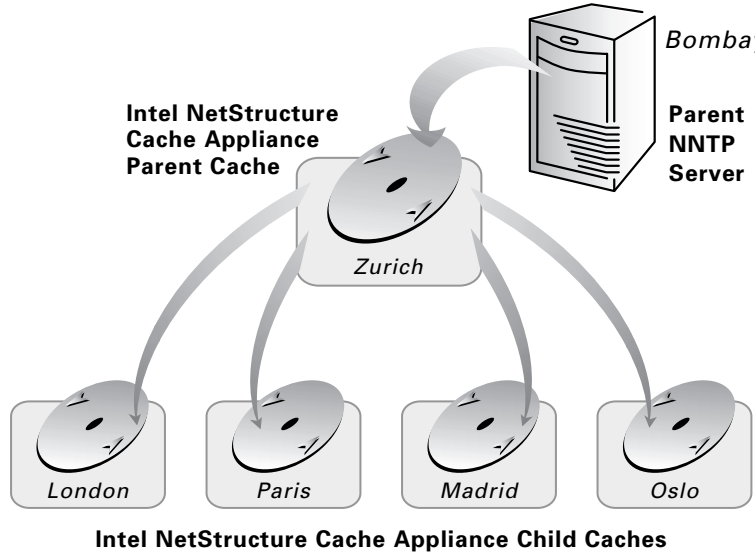

Figure 8 Hierarchy of news caching servers

In *Figure 8* above, the *parent* news server for each of the child appliances is the parent appliance. The parent appliance is a child cache to the distant parent news server.

News article caching

The Intel NetStructure Cache Appliance can function as a news server or a caching news server. News, also known as USENET and *discussions*, is a system of online discussion groups. NNTP is the protocol used to retrieve and distribute these discussion groups. The appliance supports NNTP as specified in RFC 977 and many common and proposed extensions.

To read news articles, users need a news reader, such as Netscape Communicator or Microsoft Internet Explorer, and access to a news server. The appliance is a caching news server. It can be configured to sit transparently between users and a *parent* or *backing* news server, increasing responsiveness for the user and decreasing network bandwidth use and the load on the parent news server.

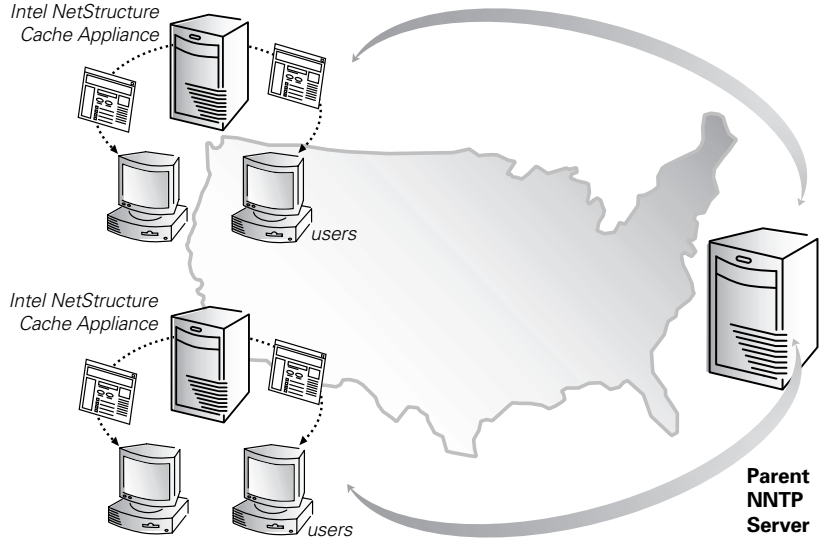

Figure 9 Intel NetStructure Cache Appliances caching news articles for a distant NNTP server

The appliance provides many options that you can configure for supporting parent NNTP servers. The rest of this section describes the appliance's NNTP features.

The appliance as a news server

As a news server, the Intel NetStructure Cache Appliance does the following:

- ✔ Maintains lists of supported news groups
- ✔ Accepts news feeds for each supported news group
- ✔ Serves requested articles to users
- ✔ Accepts and numbers user postings to the supported news groups

The appliance as a caching proxy news server

As a caching proxy news server for a particular news server, the Intel NetStructure Cache Appliance does the following:

- ✔ Maintains lists of the news groups on its parent NNTP servers. You can configure the frequency that the appliance updates its copies of group lists.
- ✔ Caches and serves article overview lists on demand. You can also tell the appliance to pull article overview lists from the parent news server periodically.
- ✔ Caches and serves articles on demand. The appliance can also accept news feeds, like any news server.
- ✔ Caches and serves miscellaneous *LIST* files, such as subscription files.
- ✔ Sends user postings to the parent news server.

When clients issue news requests, the appliance intercepts these requests and serves them from its cache, reducing traffic to parent news servers. If a particular overview or article is not in the cache, the appliance forwards requests to the parent server.

Supporting several parent news servers

The Intel NetStructure Cache Appliance can cache articles for several news servers. You specify the parent news servers for the appliance from the command-line interface (see *Configuring NNTP servers, on page 65*). For each parent news server, the appliance can cache some or all of that server's news groups.

Some of the possible parent configurations that the appliance supports are as listed below:

Several news servers supplying the same groups:

Several news servers can be configured to redundantly serve the same groups, providing enhanced reliability. The appliance provides the following features for managing these configurations:

- ✔ Priorities

  If the appliance has to contact a parent news server for information about a group supplied by several news servers, then it contacts the news server with the highest priority.

- ✔ Round-robin

  If several parent news servers supplying the same group have the same priority, the appliance selects a parent news server in round-robin fashion.

- ✔ Failover

  If a request to a parent server fails, the appliance tries the next server in the round robin (of the same priority) and then servers of lower priority.

- ✔ Background retries

  Failed servers are retried in the background and are used (restored to their specified priority) when they become available.

Several servers supplying different groups:

Several news servers can be configured with news servers supplying different (disjoint) groups. You can use this feature to spread the load based on group.

Nonstandard ports and network interfaces

You can configure the interface from which to connect to a parent news server port. You can also configure the port on the parent server to which the appliance connects.

Blocking particular groups

You can block particular groups on specified news servers. Clients do not see blocked groups in news server group lists. For information on how to list all block groups by using the from the command-line interface, see *Configuring NNTP servers, on page 65*.

Clustering

You can configure large clusters of Intel NetStructure Cache Appliances to act as a single large virtual cache that has all the storage and serving power of the aggregate. The high-performance object store maintains all articles, overview lists, group lists, and LIST files across the cluster. This information is updated at configured intervals so that users and child caches see a consistent view of news.

Two types of clusters are supported: soft clusters and management-only clusters. A soft cluster consists of multiple appliances that use an external clustering device such as an L4 Switch or router to handle load balancing and routing responsibilities. A management-only cluster also consists of multiple appliances whose functions are managed through a proprietary communications protocol accessible through the Manager UI. A management-only cluster does not use an external clustering device.

For more information about clustering, see *Clustering, on page 144*.

Transparency

The Intel NetStructure Cache Appliance can transparently intercept NNTP traffic bound for a well known NNTP server. By transparently intercepting, caching, and serving the NNTP data from a centralized parent news server, the appliance simplifies migration and administration while both increasing responsiveness and decreasing network use.

Posting

The Intel NetStructure Cache Appliance sends user article postings to the parent news server. You can specify the parent news server that receives postings for a particular group or set of groups from the command-line interface. For procedural information, see *Configuring NNTP servers, on page 65*. When acting as the news server (accepting article feeds), the appliance accepts postings.

With background posting, the appliance queues posted articles until the posting news server can accept the posted article.

Maintaining the cache: updates and feeds

The Intel NetStructure Cache Appliance can maintain the freshness of its cache by:

- ✔ Updating its cache on demand
- ✔ Actively retrieving (pulling) updates at configured intervals
- ✔ Accepting news feeds

You can configure the following options from both the Manager UI and the command-line interface:

Pull the overview information for specified groups:

For all groups designated as *pullover*, the server will retrieve the overview database information (using the `OVER`/`XOVER` commands) automatically and periodically. Pulling overview information can be useful for high volume groups that are frequently read but from which only a subset of the articles are accessed.

Pull the articles for specified groups:

For all groups designated as *pull*, the appliance will retrieve the articles automatically and periodically. Pulling groups is useful when you do not want to or cannot set up a full or partial feed.

Dynamically subscribe to specified groups:

The appliance can monitor the usage pattern for groups, and those for which the overview database is very frequently accessed can be treated as pullover groups. Likewise, those for which the articles are very frequently accessed can be treated as pull groups.

Take a partial feed (push) for specified groups:

For all groups designated as *push*, the appliance verifies that it has any requested articles and retrieves them from the parent server if they are not available locally. Partial feeds are useful for groups where *some* articles are always accessed, or for shifting article transport to a time of day when bandwidth is cheaper or more plentiful.

Take a full feed for some or all groups:

For all groups designated as *feed*, the appliance does not connect to the parent news server, and instead acts like a conventional news server. In particular, if a cache miss occurs, the appliance does not forward the request to a parent news server.

You can use full feeds for very high volume groups in which most or all the articles are accessed. You can also use them for shifting article transport to a time when bandwidth is cheaper or more plentiful.

Caution Taking a full feed is not recommended as the server will have no way to retrieve an article if it is lost for any reason (*e.g.* such as lack of space or hardware failure).

For information on how to configure update frequencies by using the Manager UI, see the **Configure: Protocols** page (*Using the Protocols page, on page 30*). For information on how to control the appliance's caching behavior for specific news groups from the command-line interface, see *Configuring NNTP servers, on page 65*.

Configuring Access control

You can configure different types of user authentication based on source domain, host name, or IP range from the command-line interface. See *Configuring NNTP access, on page 69* for more information.

Obeying NNTP control messages

By default, the Intel NetStructure Cache Appliance periodically checks the parent server for new groups, cancelled articles, and new articles for nonfeed news groups. If you have enabled these periodic checks in the **Configure: Protocols** page, you do not need to configure the appliance to obey NNTP control messages. See *Configuring NNTP, on page 31* for more information.

However, you can configure the appliance to obey NNTP control messages. In particular, you can enable the appliance to obey `cancel`, `addgroup`, and `rmgroup` messages in the **Configure: Protocols** page of the Manager UI. For example, if you select **Obey cancel control messages**, the appliance pulls `cancel` messages automatically to obey them.

Client bandwidth throttling

You can limit the amount of bandwidth allotted to clients for downloading articles. Clients that attempt to exceed the bandwidth limit will have each operation slowed to keep their bandwidth consumption to the limit. See *Configuring NNTP, on page 31* for more information.

Carrier-class architecture

The Intel NetStructure Cache Appliance is designed for carrier class operation. It offers the following:

- ✔ High performance
- ✔ High availability
- ✔ Node fault tolerance
- ✔ Expansion capabilities
- ✔ Centralized management

Performance

By combining Intel NetStructure Cache Appliance nodes into clusters, you can multiply individual performance. The following sections describe the appliance's performance features.

Self-tuning DataFlow core

A streaming DataFlow I/O core transfers data to and from disk and network connections. This core adapts to both TCP network dynamics and disk performance dynamics to result in fast and continuous data flow through many thousands of simultaneous connections.

Fine-grained parallelism

The appliance uses a highly parallel application that can manage hundreds of thousands of concurrent activities by combining kernel multithreading with an internal scheduling system called Nanothreading.

Raw-disk object store

The appliance stores all cached documents in a custom, high-speed database called the *object store*. The object store is a streaming database that supports storing alternate versions of the same object, varying on spoken language or browser type.

Alias-free caching

The object store uses content-fingerprinting technology to recognize when two URL addresses refer to the same content, and shares the content copy in cache. Thus, the appliance caches identical content only once. The alias-free cache frees cache space and provides a higher aggregate hit rate for the same cache size.

Fast space reclamation

A space-reclamation algorithm ensures the appliance collects and removes stale data. This garbage collection runs continuously and in real time.

RAM caching

To serve popular objects fast and reduce load on disks, the appliance maintains a small RAM memory cache of extremely popular objects.

Clustering

The appliance uses soft clustering and managed clustering to meet the performance needs of today and to scale to the needs of tomorrow. You can increase the Intel NetStructure Cache Appliance performance incrementally by adding new nodes to the cluster. For more information about clustering, see *Clustering, on page 140*.

Fast DNS resolver

The appliance includes a fast, asynchronous DNS resolver to streamline conversion of host names to IP addresses.

Host database

The appliance maintains a database of information about

- ✔ Internet hosts
- ✔ DNS information
- ✔ HTTP versions of hosts
- ✔ Host reliability and availability information

For information about how you can configure the host database, see *Using the Host Database page, on page 44*.

Advanced protocol features

The appliance supports

- ✔ Performance features of the emerging HTTP 1.1 protocol, such as persistent connections, request pipelining, and cache-control features.
- ✔ A rich set of commands to optimize the performance of NNTP browsing, including support for RFC 977, such as the OVER, PAT, XREPLIC and NEXT/PREV commands, and all common extensions.
- ✔ Caching of all NNTP data types and reception of news article feeds. See *News article caching, on page 138*.
- ✔ FTP caching, Internet Cache Protocol (ICP) messaging, and the SNMP protocol for network management.

Fast kernel packet engine

The appliance contains a high-speed core TCP/IP network packet engine called the Adaptive Redirection Module (ARM). This packet engine supports high-speed interception of traffic for transparency, supports automatic bypass of sites that do not function properly with proxy caches, and efficiently streams data to the Intel NetStructure Cache Appliance. See *Transparent proxy caching, on page 120* for more information.

High-availability

The Intel NetStructure Cache Appliance contains high-availability features that work together to increase reliability, minimizing the impact of hardware or software failures. The following sections describe these features.

Alarms

The appliance signals an *alarm* for any detected failure condition. Alarms are presented on the Manager UI, and can be configured to send email or page support personnel.

Pending alarms are indicated on the Dashboard of the Manager UI as a red lamp, as shown in *Figure 10*.

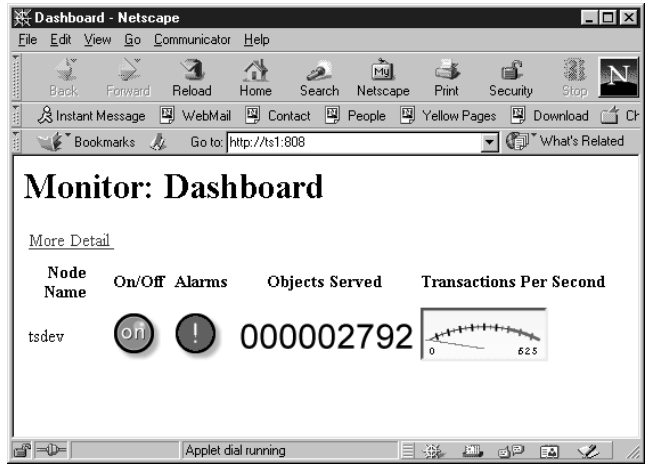

Figure 10 The Monitor Dashboard

The appliance also supports email notification for alarms. You set the email address to which alarms are sent from the command-line interface.

Virtual IP failover

The virtual IP failover option is available to clustered Intel NetStructure Cache Appliances. When virtual IP failover is enabled, the appliance maintains a pool of virtual IP addresses that it assigns to the nodes in the cluster. These addresses are virtual only in the sense that they are not tied to a specific machine; the appliance has the flexibility to assign them to any node in the cluster. To the outside world, these virtual IP addresses are the addresses of the appliance cluster.

The appliance handles virtual IP failover in the following ways:

- ✔ By maintaining cluster communication. Nodes automatically exchange statistics and configuration information through multicast communication. If multicast heartbeats are not received from one of the cluster nodes, the other nodes recognize it as unavailable.
- ✔ By reassigning IP addresses of failed nodes to operational nodes within approximately 30 seconds. This feature allows service to continue without interruption.
- ✔ By using the *ARP rebinding* process to handle IP reassignment. With this process, the IP addresses are assigned to new network interfaces, and the new assignment is broadcast to the local network.

You assign virtual IP addresses through the Manager UI as described in *Setting virtual IP addressing options, on page 26*. Note that virtual IP addresses must be pre-reserved like all IP addresses, before they can be assigned to an appliance.

Load shedding

Overload conditions, such as network outages, misconfigured routers, or security attacks, can slow down the Intel NetStructure Cache Appliance's response time. In transparent configurations, the appliance can use its ARM bypass functionality to forward overload requests directly to origin servers, bypassing the cache. When the overload condition dissipates, the appliance automatically returns to full caching mode.

You can configure the appliance to shed load if HTTP hit transaction times become large. See *Configuring load-shedding, on page 28* for instructions on configuring load shedding options.

Node fault tolerance

The appliance tolerates failures on any of the cache disks. The two classes of failures that the appliance handles are partial disk failures and total disk failures:

- ✔ A partial failure is one in which a small portion of a disk becomes unusable. If this occurs the appliance marks that portion of disk as corrupt and continues to use the rest of the disk while avoiding the corrupt portion.
- ✔ A total disk failure is one in which the hard drive becomes unusable. In this case, the appliance marks the entire disk as corrupt and continues using the remaining disks.

For either failure, an alarm is sent to the Manager UI indicating which disk failed so that an administrator can replace it. The appliance maintains two boot images should the primary drive completely fail.

Note If all of the cache disks fail, the appliance operates in proxy-only mode.

Expansion capabilities

The Intel NetStructure Cache Appliance automatically detects the addition or removal of nodes.

If you connect an additional node to a cluster, you need only install the appliance software on the new node, making sure that the cluster name and port assignments match those of the existing cluster. The new node is recognized automatically.

If a node fails or is shut down and removed, and if virtual IP failover is enabled, then requests destined for the missing node are handled by another node in the cluster.

Centralized administration

The Intel NetStructure Cache Appliance incorporates many native command and control features for carrier-class system management and administration. The following list provides an overview of these features:

Single system image
The appliance maintains a single system image for every node configured into the appliance cluster.

Multicast management protocol
The appliance uses a multicast management protocol to manage the cluster's single system image. Information about cluster membership, configuration, and exceptions is shared across all managers in the cluster, and the appliance automatically propagates configuration changes to all nodes in the cluster.

Node discovery
The appliance automatically detects new appliance nodes on your network and adds them to the cluster, propagating the latest configuration information to the newcomer. This feature provides a convenient way to bootstrap new machines.

Browser-based management interface
The Manager UI is a Web-based interface that you can access through a browser. You can configure management access to the Manager UI through password authentication. See *Chapter 4, Configuring the Appliance* for information about using the Manager UI.

Command-line interface
In addition to the Manager UI, the appliance supports a command-line interface (CLI). The CLI provides a text-based interface that lets you configure the system's network addresses and control, configure, and monitor the appliance. See *Chapter 1, Using the Command-Line Interface* for more information.

SNMP
The appliance supports SNMP access for reading statistics and sending traps (SNMP alarms). The appliance integrates into existing SNMP-managed networks, appearing as additional managed device.

SNMP is a standard way of managing everything that is a part of your network environment. SNMP-compliant devices or agents store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP managers. SNMP managers probe devices for status and SNMP agents report whether a device is functioning properly.

Note The Intel NetStructure Cache Appliance supports two MIBs: MIB-2 (a standard MIB) and the Intel NetStructure Cache Appliance MIB. You can

enable SNMP access to either one or both of these MIBS on your Intel NetStructure Cache Appliance. See *Enabling SNMP agents, on page 29*.

If a device fails, it can send a warning message or an SNMP trap to the SNMP monitoring station. All SNMP agents require you to configure the trap destination IP address before they can send traps. This configuration varies among agent implementations. It can also depend on the MIB.

Client ACL

In addition to supporting SSL security, the appliance also supports client access control lists (ACLs). The appliance serves only requests from clients whose IP addresses are on the ACL. You can edit the ACL from the command-line interface (see *Chapter 1, Using the Command-Line Interface*).

Appendix B

Error Messages

This appendix contains the following sections:

- *HTML messages sent to clients, on page 152* describes the HTML error messages that the Intel NetStructure Cache Appliance sends to browser clients (not to be confused with standard HTTP response codes)
- *Standard HTTP response messages, on page 154* describes the standard HTTP response codes that web servers send to browser clients

HTML messages sent to clients

The appliance returns detailed error messages to browser clients when there are problems with the HTTP transactions requested by the browser. These response messages correspond to standard HTTP response codes, but provide more information. A list of the more frequently encountered HTTP response codes is provided on *page 154*.

The following table lists the appliance's hard-coded HTTP messages and their corresponding HTTP response codes.

Title	HTTP code	Description
Access Denied	*403*	*You are not allowed to access the document at location <URL>.*
Bad HTTP request for FTP Object	*400*	*Bad HTTP request for FTP object.*
Cache Read Error	*500*	*Error reading from cache. Please retry request.*
Connection Timed Out	*504*	*Server has not sent any data for too long a time.*
Content Length Required	*400*	*Could not process this request because no Content-Length was specified.*
Cycle Detected	*400*	*Your request is prohibited because it would cause an HTTP proxy cycle.*
Forbidden	*403*	*<port_number> is not an allowed port for SSL connections.* *(You have made a request for a secure SSL connection to a forbidden port number.)*
FTP Authentication Required	*401*	*You need to specify a correct username and password to access the requested FTP document <URL>.*
FTP Connection Failed	*502*	*Could not connect to the server <server name>.*
FTP Error	*502*	*The FTP server <server name> returned an error. The request for document <URL> failed.*

Title	HTTP code	Description
Host Header Required	*400*	*An attempt was made to transparently proxy your request, but this attempt failed because your browser did not send an HTTP "Host" header. Please manually configure your browser to use* `http://<proxy_name>:<proxy port>` *as an HTTP proxy. Please refer to your browser's documentation for details.* *Alternatively, end users can upgrade to a browser that supports the HTTP "Host" header field.*
Host Header Required	*400*	*Your browser did not send a Host HTTP header field and therefore the virtual host being requested could not be determined. To access this web site correctly, you will need to upgrade to a browser that supports the HTTP Host header field.*
HTTP Version Not Supported	*505*	*The web server <`server name`> is using an unsupported version of the HTTP protocol.*
Invalid HTTP Request	*400*	*Could not process this <`client request HTTP method`> request for <`URL`>.*
Invalid HTTP Response	*502*	*The host <`server name`> did not return the document <`URL`> correctly.*
Malformed Server Response	*502*	*The host <`server name`> did not return the document <`URL`> correctly.*
Malformed Server Response Status	*502*	*The host <`server name`> did not return the document <`URL`> correctly.*
Maximum Transaction Time exceeded	*504*	*Too much time has passed transmitting document <`URL`>.*
No Response Header From Server	*502*	*The host <`server name`> did not return the document <`URL`> correctly.*
Not Cached	*504*	*This document was not available in the cache, and you (the client) only accept cached copies.*
Not Found on Accelerator	*404*	*The request for <`URL`> on host <`server name`> was not found. Check the location and try again.*
NULL	*502*	*The host <`host name`> did not return the document <`URL`> correctly.*

Title	HTTP code	Description
Proxy Authentication Required	*407*	*Please login with username and password.*
Server Hangup	*502*	*The server <host name> closed the connection before the transaction was completed.*
Temporarily Moved	*302*	*The document you requested, <URL>, has moved to a new location. The new location is <new URL>.*
Transcoding Not Available	*406*	*Unable to provide the document <URL> in the format requested by your browser.*
Tunnel Connection Failed	*502*	*Could not connect to the server <host name>.*
Unknown Error	*502*	*The host <host name> did not return the document <URL> correctly.*
Unknown Host	*500*	*Unable to locate the server named <host name> the server does not have a DNS entry. Perhaps there is a misspelling in the server name, or the server no longer exists. Double-check the name and try again.*
Unsupported URL Scheme	*400*	*Cannot perform your request for the document <URL> because the protocol scheme is unknown.*

Standard HTTP response messages

The following standard HTTP response messages are provided for your information. For a more complete list and descriptions, see the Hypertext Transfer Protocol — HTTP/1.1 Specification.

Message	Description
200	*OK*
202	*Accepted*
204	*No Content*
206	*Partial Content*
300	*Multiple Choices*
301	*Moved Permanently*
302	*Found*
303	*See Other*
304	*Not Modified*

Message	Description
400	*Bad Request*
401	*Unauthorized; retry*
403	*Forbidden*
404	*Not Found*
405	*Method Not Allowed*
406	*Not acceptable*
408	*Request Timeout*
500	*Internal server error*
501	*Not Implemented*
502	*Bad Gateway*
504	*Gateway Timeout*

Glossary

Alternates

Different versions of the same web object. Some web servers answer requests to the same URL with a variety of objects. The content of these objects can vary widely, depending on whether a server delivers content for different languages, targets different browsers with different presentation styles, or delivers variable content at different times of the day.

ARM

Adaptive Redirection Module. Used in transparent proxy caching, ARM is an Intel NetStructure Cache Appliance component that redirects intercepted client traffic destined for an origin server to the Intel NetStructure Cache Appliance application. Before the traffic is redirected by the ARM, it is intercepted by an *L4 switch* or router.

Cache

Stores copies of frequently accessed objects close to users and serves them to users when requested. See also *Object store*.

Cache hierarchy

Levels of caches that communicate with each other. All cache hierarchies recognize the concepts of *Parent cache* and *Child cache*.

Cache hit

An object in the cache that can be served directly to the client.

Cache miss

An object that is *not* in the cache or that is in the cache but no longer valid. In both cases, the Intel NetStructure Cache Appliance must get the object from the *Origin server*.

Caching web proxy server

A web proxy server with local cache storage that allows the proxy to fulfill client requests locally, using a cached copy of the origin server's previous response.

CGI

Common Gateway Interface. A set of rules that describe how a web server and another piece of software (a *CGI program*) located on the same machine communicate.

cgi-bin

The most common directory name on a web server in which *CGI* programs are stored.

Child cache

A cache lower in a *Cache hierarchy* for which the Intel NetStructure Cache Appliance is a parent. See also *Parent cache*.

Cluster

A group of the Intel NetStructure Cache Appliance nodes that are configured to act as a single large virtual cache. For information on the supported cluster schemes, see *Management-only clustering* and *Soft Cluster*.

Configure mode

One of two modes in the *Intel NetStructure Cache Appliance Manager*. Configure mode lets you configure the Intel NetStructure Cache Appliance from a web browser. See also *Monitor mode*.

Configure page

A web-based page that appears on the Manager UI when you click on an active button while in Configure mode. See also *Monitor page*.

Cookie

A piece of information sent by a web server to a web browser. The browser software saves the information and sends it back to the server whenever the browser makes additional requests from the server. Cookies enable web servers to keep track of users.

DNS

Domain Name Service. The Intel NetStructure Cache Appliance includes a fast, asynchronous DNS resolver to streamline conversion of host names to IP addresses.

Explicit proxy caching

A configuration option where client software (typically a browser) must be specifically configured to send web requests to the Intel NetStructure Cache Appliance proxy.

FTP

File Transfer Protocol. A protocol based on TCP/IP for reliable file transfer.

Full clustering

An Intel NetStructure Cache Appliance cluster distributes its cache across its nodes into a single, virtual object store, rather than replicating the cache, node by node. See also *Management-only clustering*.

HTTP

Hyper Text Transfer Protocol. The client-server protocol upon which the World Wide Web is based.

ICP

Internet Cache Protocol. A protocol for proxy caches to exchange information about their content.

Intel NetStructure Cache Appliance Manager

The Intel NetStructure Cache Appliance's browser-based interface consisting of a series of web pages that enable you to monitor performance and change configuration settings.

IP

Internet Protocol. The lowest-layer protocol under TCP/IP responsible for end-to-end forwarding and long packet fragmentation control.

IP Allow rule

Specifies ranges of IP addresses allowed to use the appliance as a web proxy.

ISP

Internet Service Provider. An organization that provides access to the Internet.

JavaScript

A network-oriented programming language specifically designed for writing programs that can be safely downloaded to your computer through the Internet.

L4 switch

An ethernet switch that can control network traffic flow using Level 4 rules. The switch can intercept desired client protocol packets and direct them to a proxy for transparent operation.

Manager

A functional software module resident in the appliance that acts as an interface between the status of the appliance and the Manager UI.

Manager Allow rule

Specifies ranges of IP addresses allowed to access the Manager UI.

Management-only clustering

An Intel NetStructure Cache Appliance option where all nodes in a cluster automatically share configuration information through a proprietary communications protocol. See also *Full clustering*.

MIB

Management Information Base. The set of parameters that an SNMP management station can query in the SNMP agent of a network device (for example, a router). The Intel NetStructure Cache Appliance supports two MIBs: MIB2 (a well-known standard MIB) and the proprietary Intel NetStructure Cache Appliance MIB, which provides more specific node and cluster information.

Monitor mode

One of two modes in the *Intel NetStructure Cache Appliance Manager*. Monitor mode lets you monitor the Intel NetStructure Cache Appliance's performance from a web browser. See also *Configure mode*.

Monitor page

A web-based page that appears on the Manager UI when you click on an active button while in Monitor mode. See also *Configure page*.

MRTG

Multi Router Traffic Grapher. A graphing tool provided with the Intel NetStructure Cache Appliance that enables you to monitor the Intel NetStructure Cache Appliance's performance.

News server

A web server you can access to read and post to usenet news groups.

NNTP

Network News Transfer Protocol. A protocol used to distribute, inquire, retrieve, and post news articles.

Object store

A custom high-speed database where the Intel NetStructure Cache Appliance stores all cached objects.

Origin server

The web server that contains the original copy of the requested information.

Parent cache

A cache higher up in a *Cache hierarchy*, to which the Intel NetStructure Cache Appliance can send requests.

POP

1. Point of Presence. Usually a city or location to which a network can be connected, often with dial up phone lines.

2. Post Office Protocol. The basic protocols for addressing e-mail.

Proxy server

See *Web proxy server*.

Reverse proxy

A option that allows the Intel NetStructure Cache Appliance to be configured as a web server for convenient geographical distribution of server content. Reverse proxy also off loads static content service from servers building dynamic content and provides a peak load buffer or *surge protector* for web servers. Sometimes referred to as *Server acceleration*.

Router

A device that handles the connection between 2 or more networks. Routers look at destination addresses of the packets passing through them and decide which route to send them on.

Server

The software engine resident in the appliance that enables the appliance to cache objects.

Server acceleration

See *Reverse proxy*.

SNMP

Simple Network Management Protocol. A set of standards used for communication with devices connected to a TCP/IP network. SNMP-compliant devices (agents) store information about themselves in *MIB*s and provide this information to SNMP Managers.

Soft Cluster

Multiple appliances that use an external clustering device such as an L4 Switch or router to handle load balancing and routing responsibilities.

SSL

Secure Sockets Layer. A protocol that enables encrypted, authenticated communications across the Internet. Used mostly in communications between web servers and web browsers.

TCP

Transmission Control Protocol. An Internet Standard transport layer protocol. TCP provides reliable end-to-end communication by using sequenced data sent by IP.

Transparent proxy caching

A configuration option that enables the Intel NetStructure Cache Appliance to intercept and respond to Internet requests without requiring users to reconfigure their browser settings. It does this by intercepting traffic destined for an origin server and redirecting that traffic through the cache.

URL

Uniform Resource Locator. The address that defines the route to a file on the web or other Internet facility.

Virtual IP failover

An option available to clustered Intel NetStructure Cache Appliances, where the appliance maintains a pool of virtual IP addresses that it assigns to the nodes of a cluster. If a node fails, the remaining nodes mask the fault and take over the failed node's virtual interface.

WCCP

Web Cache Control Protocol. A protocol used by Cisco IOS-based routers to redirect traffic during transparent proxy caching.

Web proxy server

Forwards client requests to *Origin server*s. The proxy server may deny requests according to filter rules or security limitations.

Web server

A computer that provides World Wide Web services on the Internet. See also *Origin server*.

WPAD

Web Proxy Auto-Discovery. A protocol that allows clients to automatically locate a web proxy, providing the benefits of a proxy without the need for explicit client configuration.

Index