

HP Domain ServiceControl 1.0 Concepts and Operation Guide

Edition 3



**Manufacturing Part Number: J1592-90009
E0699**

U.S.A.

© Copyright 1999, Hewlett-Packard Company.

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Warranty. A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

Restricted Rights Legend. Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

HEWLETT-PACKARD COMPANY 3000 Hanover Street Palo Alto,
California 94304 U.S.A.

Use of this manual and flexible disk(s) or tape cartridge(s) supplied for this pack is restricted to this product only. Additional copies of the programs may be made for security and back-up purposes only. Resale of the programs in their present form or with alterations, is expressly prohibited.

Copyright Notices. ©copyright 1983-98 Hewlett-Packard Company, all rights reserved.

Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under the copyright laws.

©copyright 1979, 1980, 1983, 1985-93 Regents of the University of California

This software is based in part on the Fourth Berkeley Software Distribution under license from the Regents of the University of California.

©copyright 1980, 1984, 1986 Novell, Inc.

©copyright 1986-1992 Sun Microsystems, Inc.

©copyright 1985-86, 1988 Massachusetts Institute of Technology.

©copyright 1989-93 The Open Software Foundation, Inc.

©copyright 1986 Digital Equipment Corporation.

©copyright 1990 Motorola, Inc.

©copyright 1990, 1991, 1992 Cornell University

©copyright 1989-1991 The University of Maryland

©copyright 1988 Carnegie Mellon University

Trademark Notices. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

X Window System is a trademark of the Massachusetts Institute of Technology.

MS-DOS and Microsoft are U.S. registered trademarks of Microsoft Corporation.

OSF/Motif is a trademark of the Open Software Foundation, Inc. in the U.S. and other countries.

Netscape, Netscape Communications, the Netscape Communications logo, Netscape Communicator, Netscape Enterprise Server, Netscape FastTrack Server, and Netscape Navigator are trademarks of Netscape Communications Corporation.

Cisco is a registered trademark of Cisco Systems, Inc.

Contents

1. HP Domain ServiceControl

Installation Requirements	11
Hardware	11
Software	11
Disk Space (Estimated)	11
Installing HP Domain ServiceControl	11
Configuring HP Domain ServiceControl	12

2. HP AdmissionControl

What is HP AdmissionControl?	14
What is a Session?	14
When Does a Web Server Become Overloaded?	14
How Does HP AdmissionControl Work?	14
Installing HP AdmissionControl	16
Installation Requirements	16
Hardware	16
Software	16
Disk Space (Estimated)	16
Installing the Software	16
Configuring HP AdmissionControl	17
Running the setup Script	17
Setting the Configurable Parameters	18
Admit Sessions	20
End Sessions	21
Respond to Rejected Sessions	21
Creating an Allocation Page for Rejected Sessions	23
Starting HP AdmissionControl	24
Stopping HP AdmissionControl	25
Removing HP AdmissionControl	26

Contents

Configuration Parameters	27
Defaults Already Configured in config.ac	31
Logging HP AdmissionControl Information	33
Saving Information to the Server Access Log	33
Logging Statistics to a File	34
Other Configuration Considerations	36
Tuning Your Web Server's Cache	36
Allowing Equal Access to a System's Resources	36
Detecting Failed Servers in a Cluster	37
3. HP LocalDirector Controller	
What is HP LocalDirector Controller?	40
Installing HP LocalDirector Controller	41
Installation Requirements	41
Hardware Requirements	41
Software Requirements	41
Other Requirements	41
Installing the Software	41
Configuring Cisco LocalDirector	42
Configuring HP LocalDirector Controller	43
Starting HP DESMS	43
Configuring the HP LDControl Server	44
Configuring the HP LDControl Daemon	45
Starting and Stopping the HP LDControl Daemon	47
Configuring a System with More than One LAN Card	48
Troubleshooting the HP LD Controller	49
General Troubleshooting	49
Problems Accessing the HP LD Controller from HP DESMS	49
HP LDControl Daemon	49

Printing History

The manual printing date and part number indicate its current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The manual part number will change when extensive changes are made.

Manual updates may be issued between editions to correct errors or document product changes. To ensure that you receive the updated or new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

First Edition: June 1998 (HP-UX Release 10.20 and 11.0)

Second Edition: August 1998 (HP-UX Release 10.20 and 11.0)

Third Edition: June 1999 (HP-UX Release 10.20 and 11.0)

1 HP Domain ServiceControl

HP Domain ServiceControl (HP DSC) is a web-based solution that provides the quality of service needed to maintain your web applications.

HP DSC is made up of the following components:

HP Domain ServiceControl

- **HP AC - HP AdmissionControl** controls the amount and length of sessions admitted to a web server such that the system does not become overloaded.
- **HP DESMS - HP Domain Enterprise Server Management System** configures, administers, operates, and monitors your system including HP AC and HP LDC.
- **HP LDC - HP LocalDirector Controller** automatically generates the information needed by Cisco's LocalDirector to distribute load and manage TCP/IP connections.

Installation Requirements

The following is the list of requirements to install all the components of HP Domain ServiceControl (HP DSC). If you are installing a subset of the components, go to the appropriate chapter(s) in this manual for more information. Information about HP Domain Enterprise Server Management System can be found in the *Getting Started with HP Domain Software* manual.

Hardware

- HP 9000 computer system
- 32 MB of memory
- CD-ROM drive to load the software

Software

- HP-UX 10.20 or 11.0
- HP DESMS
- Netscape Enterprise or FastTrack Server version 3.0 or later
- Cisco LocalDirector version 2.1 or later

Disk Space (Estimated)

- 80 MB

Installing HP Domain ServiceControl

To install the HP DSC software, run `swinstall` (as root) and install the bundle J1592AA.

Configuring HP Domain ServiceControl

Each component of HP DSC needs to be configured individually. Go to the following areas for more information about configuration:

- HP AC - See Chapter 2 , “HP AdmissionControl,” on page 13.
- HP DESMS - See the manual *Getting Started with HP Domain Software*.
- HP LDC - See Chapter 3 , “HP LocalDirector Controller,” on page 39.

Also refer to the *HP Domain ServiceControl Release Note* for other important information.

2 HP AdmissionControl

What is HP AdmissionControl?

HP AdmissionControl (HP AC) prevents a web server from becoming overloaded by controlling the amount and length of sessions admitted.

Using the HP Domain Enterprise Server Management System (HP DESMS, HP's web-based Domain software administrator), you can configure when to admit new sessions (based on the system's load), when to end a session (based on the time interval between requests and/or the length of a session), and how to respond to a deferred session (redirect the request or return a rejection message). You can also display the status of all the systems in a web server's cluster and display performance statistics of HP AC.

What is a Session?

A session is an HTTP request or sequence of HTTP requests made to a web server by a single user. A simple session might consist of a request for a server's home page. If that page contains images, the session would also consist of requests for each image.

A more complicated and longer session might consist of a request for a server's home page, following a link to a catalog, browsing the catalog and adding items to a shopping cart, and finally supplying payment information to complete a purchase.

When Does a Web Server Become Overloaded?

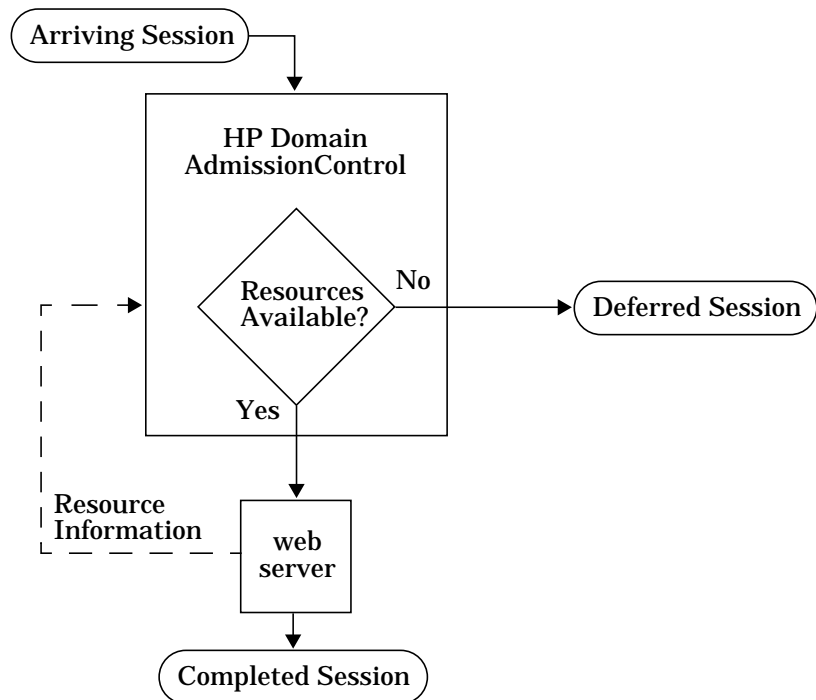
A web server becomes overloaded when there are too many requests for the server to fulfill; the server runs out of resources. Requests may be delayed or dropped. For example, a user involved in a long session may experience poor performance or may have to reload a page more than once for it to display.

How Does HP AdmissionControl Work?

HP AC monitors the resources and the number of arriving sessions on the web server. Based on available resources and how HP AC is configured, it will accept or defer a session. HP AC makes sure the web server does not become overloaded by deferring any new sessions for which it does not have resources. It also allows an admitted session to be completed.

When the web server's resources are available, HP AC allows the session to begin and will fulfill all requests made during this session. HP AC monitors the time between requests as well as the length of the session. If one of these variables exceeds the thresholds configured in HP AC, the session is terminated. Otherwise, all requests are fulfilled until the session is completed.

When the web server's resources are not available, HP AC will defer the session. The web administrator can choose to redirect the request to another server or serve a page with an error message or some helpful information.



Installing HP AdmissionControl

HP AdmissionControl (HP AC) is part of the HP Domain ServiceControl (HP DSC) product. You do not need to load the entire HP DSC product if you only want to use HP AC.

Installation Requirements

Hardware

- HP 9000 computer system
- 32 MB of memory
- CD-ROM drive to load the software

Software

- HP-UX 10.20 or 11.0
- Netscape FastTrack Server version 3.0 or later OR
Netscape Enterprise Server version 3.0 or later

Disk Space (Estimated)

- 38 MB

Installing the Software

To install HP AC, run `swinstall` (as root), double-click on the bundle J1592AA to view the software, and select the HP AC software.

Configuring HP AdmissionControl

If Netscape Enterprise Server version 3.0 or later was already installed on your system in the `/opt/ns-enterprise3` directory before you installed HP Domain ServiceControl, HP AC is automatically configured and enabled on your system.

If Netscape FastTrack Server version 3.0 or later was already installed on your system in the `/opt/ns-ftrack` directory before you installed HP Domain ServiceControl, HP AC is automatically configured and enabled on your system.

If both servers were already installed on your system before you installed HP Domain ServiceControl, HP AC is automatically configured and enabled for your Netscape Enterprise Server.

If neither server was installed before you installed HP Domain ServiceControl, you must first install the Netscape Enterprise Server version 3.0 or later or Netscape FastTrack Server version 3.0 or later and then run the setup script (see the next section for information about running the setup script).

If you installed either server after installing HP Domain ServiceControl, you must run the setup script (see next section for information about running the setup script).

If you installed either server in another server root directory other than `/opt/ns-enterprise3` (for the Netscape Enterprise Server) or `/opt/ns-ftrack` (for the Netscape FastTrack Server), you must run the setup script (see the next section for information about running the setup script).

Running the setup Script

You should only run the setup script if you installed the Netscape Enterprise Server or Netscape FastTrack Server after installing HP Domain ServiceControl or if the server root directory is something other than `/opt/ns-enterprise3` (for the Netscape Enterprise Server) or `/opt/ns-ftrack` (for the Netscape FastTrack Server).

To run the setup script, type:

```
/opt/hpac/setup
```

Configuring HP AdmissionControl

At the Root pathname of Netscape server prompt, enter the Netscape Server root directory.

The setup script will automatically configure and enable HP AC.

Default values set for all applicable configuration parameters are listed at the end of this chapter.

Setting the Configurable Parameters


You can edit the configuration parameters in three different ways:

- You can access the “HP AdmissionControl Settings” administration web page at the URL `http://system_name/hpac`. Your Netscape Server must be running to access this page. Use the system’s root user and password to access this page.
- You can edit the configuration file
`/opt/ns-server_name/server_id/config/hpac/config.ac`
- You can use the HP DESMS interface to access the “HP AdmissionControl Settings” administration web page (HP AC is located under Server Control). Your Netscape Server and HP administration server must be running.

See “Starting HP DESMS” on page 43 for more information on how to start the HP administration server.

To access HP AC, you must enter the system’s root user and password.

The following is an example of the “HP AdmissionControl Settings” page:



The following sections describe the parameters that can be set using the “HP AdmissionControl Settings” administration web page. All configurable parameters in the configuration file are listed at the end of this chapter.

HP AdmissionControl
Configuring HP AdmissionControl

Admit Sessions

For the default value, the load threshold is set to 2.

Table 2-1

HP AC Configurable Parameters: Admit Sessions

Parameter	Description
Shutdown	After the specified number of seconds, do not admit any new sessions and reject all remaining sessions.
Always	Disable HP AC and admit all sessions to the web server. If you are permanently disabling HP AC, you should remove HP AC from you system (see “Removing HP AdmissionControl” on page 26 for more information).
Load	If the load is less than this specified amount, admit the session. The load threshold works best when a web application is processor-intensive.

To calculate the load threshold, do the following:

1. Run the `top` command. Note the first listed “Load averages” and the percentage listed under “IDLE” of “Cpu states”.

Example 2-1

Sample Output from the top Command

```
System: web_server          Wed Apr 22 10:41:59 1998
Load averages: 1.45, 1.46, 1.78
149 processes: 145 sleeping, 3 running, 1 stopped
Cpu states:
LOAD  USER  NICE  SYS  IDLE  BLOCK  SWAIT  INTR  SSYS
1.45 100.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0%
```

2. If IDLE is greater than zero, then set the threshold to $(100 - \text{IDLE})/100$.

If IDLE equals zero, then set the threshold to the first listed load average.

Using the above sample output, you would set the load threshold to 1.45.

Run `top` when performance is poor and acceptable to determine an acceptable load level.

End Sessions

For the default, the interval between requests is set to 600 seconds (10 minutes) and the session duration is disabled.

Table 2-2

HP AC Configurable Parameters: End Sessions

Parameter	Description
Interval between requests	The amount of time, in seconds, between requests within the session. If the time between requests exceeds the specified threshold, the session is ended.
Session duration	The amount of time of the session. If the session exceeds the specified threshold, the session is ended. Set this parameter to zero to disable this threshold.

Respond to Rejected Sessions

Rejected sessions can either be redirected or served a web page containing a standard error message or other information.

If redirection is disabled or the redirect request is not accepted, a web page containing a standard error message or other information is returned.

The default is to disable redirection and serve an error page to a rejected session.

Table 2-3

HP AC Configurable Parameters: Redirect Rejected Sessions

Parameter	Description
Disable	Do not redirect the rejected session. Instead, return a web page with an error message or other information.
Address	Redirect the rejected session to a load balancing system (such as a cluster manager), to another web server, or to another system set up to specifically handle rejected sessions. Specify the full domain name or IP address.
Cluster	Redirect the rejected session to a randomly selected member of a cluster.

Table 2-3 HP AC Configurable Parameters: Redirect Rejected Sessions

Parameter	Description
Domain	<p>Redirect the rejected session to the domain to which the systems in the cluster belong. For example, if the cluster consists of the two systems <code>ww1.sales.acme.com</code> and <code>ww2.sales.acme.com</code>, specify the domain as <code>.sales.acme.com</code>. Note the leading dot.</p> <p>The domain specifies the scope over which the client-side cookies can be shared (HP AC uses client-side cookies to encode session information).</p>
Group	The cluster group to which this system belongs.
TTL	<p>Time to Live. The number of redirections allowed before rejecting the session. Add one to the TTL parameter to get the number of redirects. For example, setting the TTL to 0 allows one redirection. Setting the TTL to 1 allows two redirections.</p> <p>It is recommended that the TTL be set to a value no larger than 2.</p>

When redirection is disabled, the following types of web pages can be served:

Table 2-4 HP AC Configurable Parameters: Respond to Rejected Sessions

Parameter	Description
Error	The “Server Temporarily Unavailable” error message is served to rejected requests. This is the easiest but least informative way to respond to a rejected request. Extra load on the system may be generated by users trying to immediately re-establish a connection.
Page	The specified web page is served to rejected requests. You can provide information to the customer such as why the session was rejected and when would be a more opportune time to return to the site. A user is less likely to try immediately re-establishing a connection and more likely to return to the site.

Table 2-4 HP AC Configurable Parameters: Respond to Rejected Sessions

Parameter	Description
Allocate sessions	<p>A “countdown” web page is served to rejected requests. The user is asked to wait a certain amount of time before getting access to the site. This type of session is given priority over new sessions.</p> <p>Enter the number of sessions you wish to have admitted each minute. For example, if you allocate five sessions per minute, one session will be admitted every twelve seconds. If seven sessions are rejected, five will be admitted the next minute and two will be admitted the minute after. Make sure the allocation rate is less than the system’s capacity.</p> <p>You may want to compensate the allocation for no-shows.</p> <p>Monitoring the “% Priority Requests” statistic will show the system’s load due to priority sessions. Reduce the allocation rate if the number approaches or is over 50%.</p> <p>Also monitor the “Last Allocation Delay” statistic. It shows how long a deferred user waited before his/her session was admitted.</p>
External	<p>Use an external function or application. For example, the user could be offered a discount or other incentive to return to the site; or, the user could be offered to enroll in a service plan that guarantees access to the site. Currently, no supported functions or applications exist. If selected, the Error page is returned.</p>
Exclude	<p>After the initial rejection, the user is prevented from returning to the site after the specified amount of time (in seconds). This discourages a user from trying to immediately re-establish a connection.</p> <p>If a session is allocated, this is the minimum delay that can be scheduled.</p>

Creating an Allocation Page for Rejected Sessions

The following is the default allocation page served to the user:

HP AdmissionControl

Configuring HP AdmissionControl

Example 2-2

Server Busy

We are sorry our server is busy, and in order to ensure adequate levels of service to our current visitors, we cannot serve you at this time. We have allocated you a position in line and will serve you as soon as we can.

You will automatically enter the site in 1:23
(minutes:seconds).

To change the content of this page, you must edit the file
`/opt/ns-server_name/server_id/config/hpac/config.ac`

To change the “Server Busy” title, edit the `deferral.title` parameter. To change the content of the message, edit the `deferral.message` parameter. To edit the countdown line, edit the `deferral.enterMsg` parameter. Note that each parameter is ended by a single period on its own line.

The default allocation page would include the following entries in the `config.ac` file:

```
deferral.title
Server Busy
.
deferral.message
We are sorry our server is busy, and in order to ensure
adequate levels of service to our current visitors, we
cannot serve you at this time. We have allocated you a
position in line and will serve you as soon as we can.
.
deferral.enterMsg
You will automatically enter the site in
.
```

Starting HP AdmissionControl

HP AC is automatically enabled after you have installed HP ServiceControl (provided you have met certain conditions; see “Configuring HP AdmissionControl” on page 17 for more information) or run the setup script (see “Running the setup Script” on page 17 for more information).

If you have stopped HP AC, then you can start it again by configuring and specifying the load from the “HP AdmissionControl Settings” administration web page.

Stopping HP AdmissionControl

To stop or disable HP AC, set the Admit Session to “Always” from the “HP AdmissionControl Settings” administration web page.

If you are permanently stopping HP AC, you should remove HP AC from your system (see “Removing HP AdmissionControl” on page 26 for more information). Because HP AC does not cache HTML, performance may be slow for static pages.

Removing HP AdmissionControl

To remove HP AC, do the following:

1. Type `/opt/hpac/setup -r`
2. At the Enter root pathname of Netscape server prompt, enter the Netscape Server root directory.

Configuration Parameters

Table 2-5 HP AC Configurable Parameters

Parameter	Range/Type	Default	Description
admit.always	N/A	N/A	Configures HP AC to always accept sessions.
admit.load	N/A	N/A	Sets the “Admit Sessions” to use the load threshold.
admit.load.threshold	$0.0 \leq x$	2.0	Sets the load threshold for admitting a session.
admit.priority.damping	$0.0 \leq x \leq 1.0$	0.001	The damping factor used to compute the priority percentage. Larger values make the controller more responsive to recent activity.
admit.priority.threshold	$0.0 \leq x \leq 1.0$	0.9	The value of the priority fraction when priority requests are rejected (0.9 = 90% priority percentage).
admit.shutdown	N/A	N/A	Sets the “Admit Sessions” to shutdown.
admit.shutdown.delay	$0 \leq x$ (integer)	600	Sets how much time before HP AC shuts down and rejects all (new and remaining) sessions.
cluster.disable	N/A	N/A	Disables clustering. The system will not try to open the <code>hpac-cluster</code> port.
cluster.domain	string		The domain used to limit the scope of the cluster over which the client-side cookies can be shared.

HP AdmissionControl
Configuration Parameters

Table 2-5 HP AC Configurable Parameters

Parameter	Range/Type	Default	Description
cluster.group	string	default	Sets the cluster group.
cluster.keepalive.interval	1 <= x (integer)	60	The number of seconds between status broadcasts.
cluster.keepalive.misses	1 <= x (integer)	3	The number of missed status broadcasts before a server is identified as down.
cluster.TTL	0 <= x (integer)	0	Sets the cluster TTL which is the number of redirections allowed before rejecting the session.
cookie.password	string	N/A	Sets the password used in creating the MD5 digest.
deferral.enterMsg	string	You will automatically enter the site in .	Part of the allocation page message served when a session is rejected.
deferral.message	string	We are sorry our server is busy, and in order to ensure adequate levels of service to our current visitors, we cannot serve you at this time. We have allocated you a position in line and will serve you as soon as we can. .	Part of the allocation page message served when a session is rejected.
deferral.title	string	Server Busy .	Part of the allocation page message served when a session is rejected.

Table 2-5 HP AC Configurable Parameters

Parameter	Range/Type	Default	Description
endsession.duration	0 <= x (integer)	0	Sets, in seconds, the maximum session duration
endsession.requestInterval	0 <= x (integer)	300	Sets, in seconds, the maximum interval between requests
isolate.damping	0.0 <= x <= 1.0	0.001	The damping factor used to compute the share of resources being used by each virtual server. Larger values make the controller more responsive to recent activity.
isolate.pattern	string	N/A	The pattern used to identify virtual servers. If provide, HP AC will isolate traffic to each virtual server.
load.damping	0.0 <= x <= 1.0	0.1	The damping factor used to compute the web server load. Larger values make the controller more responsive to recent activity.
log.file	string	N/A	The full path name of the HP AC log file. The statistics from the “HP AdmissionControl Statistics” web page are logged to this file.
log.interval	1 <= x (integer)	600	How often, in seconds, the statistics are logged to the log.file.
mime.cache	N/A	See “Defaults Already Configured in config.ac” on page 31.	The web server should cache the specified mime types.

HP AdmissionControl
Configuration Parameters

Table 2-5 HP AC Configurable Parameters

Parameter	Range/Type	Default	Description
mime.nocache	N/A	N/A	The web server should cache all but the specified mime types.
mime.type	string	N/A	Specifies the mime type used by the caching policy.
reject.error	N/A	Server Temporarily Unavailable	Sets the response to a rejected session to the “Server Temporarily Unavailable” error message. This message is not configurable.
reject.excludeTime	0 <= x	300	Sets how long, in seconds, the user is prevented from returning to the site.
reject.external	N/A	Server Temporarily Unavailable	Sets the response to a rejected session to use an external function or application. Currently, it will only serve the “Server Temporarily Unavailable” error message.
reject.future	N/A	N/A	Sets the response to a rejected session to a “countdown” web page.
reject.future.rate	1 <= x (integer)	10	Set the number of sessions that can be admitted each minute.
reject.page	N/A	N/A	Sets the response to a rejected session to a customizable web page.
reject.page.path	string	N/A	The full path and file name of the customizable web page.
shlib.initialize	string	N/A	The name of the function that initializes the library.

Table 2-5 HP AC Configurable Parameters

Parameter	Range/Type	Default	Description
shlib.library	string	N/A	The file name of the shared library.
shlib.loadMetric	string	N/A	The name of the function used as the load metric.
shlib.release	string	N/A	The name of the function that releases resources allocated by the library.
shlib.reject	string	N/A	The name of the function used when the reject policy is set to External.
shlib.sessionState	string	N/A	The name of the function that modifies the session state. This function may change the priority of sessions or end sessions.

Defaults Already Configured in config.ac

The following are defaults that have been configured in the `/opt/ns-server_name/server_id/config/hpac/config.ac` file:

```
# Warning, manual changes to this file will be
# overwritten by changes made
# via the Admission Control Policy screen.
mime.cache # a list of likely cache candidates follows
mime.type audio/basic
mime.type audio/midi
mime.type audio/x-wav
mime.type audio/x-liveaudio
mime.type audio/x-pn-realaudio
mime.type image/fif
mime.type image/gif
mime.type image/ief
mime.type image/ifs
mime.type image/jpeg
```

HP AdmissionControl

Configuration Parameters

```
mime.type image/tiff
mime.type image/wavelet
mime.type image/vnd
mime.type image/x-photo-cd
mime.type video/mpeg
mime.type video/quicktime
mime.type video/x-mpeg2
mime.type video/x-msvideo
```

Logging HP AdmissionControl Information

Saving Information to the Server Access Log

HP AC admission decisions can be saved to the server access log. To include them, do the following:

1. Start the Netscape Administration Server.
2. Select the server to configure.
3. Click on **Server Status**.
4. Select **Log Preferences**.
5. Enter the name of the new access logfile to create (Netscape will not let you change the format of a logfile that is in use).
6. Select **Custom format** near the bottom of the page.
7. Append `%Req->vars.hpacSessionState%` at the end of the Custom format field. Be sure to include a space at the beginning of this entry to separate it from the other entries.
8. Click **OK** to make the change.
9. Click on **Save and Apply** to save your changes.

One of the following values will be logged to the access logfile:

Table 2-6

Admission Decision Log Values

Value	Description
-	A management request was served from the cache.
2	The request was rejected.
3	The request was processed as part of a basic priority session.
4	The request was processed as part of a high priority session.
5	The request was deferred to a future time slot.
6	The request was redirected to another server.

Logging Statistics to a File

The statistics displayed on the “HP AdmissionControl Statistics” web page can be logged to a file. To log the statistics to a file add the following lines to `/opt/ns-server_name/server_id/config/hpac/config.ac`:

```
log.file filename  
log.interval seconds
```

Where *filename* is the name of the file to which to save the statistics and *seconds* is how often the statistics are saved (in seconds). The `log.interval` default is 60 seconds.

Restart the web server for the changes to take effect.

The following is an example of a log entry:

```
10/Mar/1998:20:39:00,15,0,0,0,0,0,0,0,181,0,0,0.5,0.00,92
```

where each field is delimited by a comma. The fields are:

Table 2-7

HP AC Log File Statistics

Statistic Name	Example Value
Time (GMT)	10/Mar/1998:20:39:00
Basic Sessions	15
Priority Sessions	0
Redirected Sessions	0
Sessions Rejected	0
Sessions Rejected Again	0
Sessions Timed Out	0
Sessions Too Long	0
Bad MD5 Requests	0
Basic Requests	181
Priority Requests	0
Management Requests	0
Load	0.5

Table 2-7

HP AC Log File Statistics

Statistic Name	Example Value
Percentage Priority Requests	0.00
Last Allocation Delay	92

Other Configuration Considerations

This section covers tasks beyond basic configuration.

Tuning Your Web Server's Cache

To control session admissions, HP AC tracks the session state by including cookies in HTTP responses. For HP AC to operate effectively, the server cannot generate all of its responses from the cache. This may cause the performance of the web server to be slow.

However, you can tune the cache by configuring the `/opt/ns-server_name/server_id/config/hpac/config.ac` file to allow specified mime types to be cached or not cached.

For example, if you serve HTTP pages with embedded images, you can cache the images by adding the following lines to the configuration file:

```
mime.cache
mime.type image/gif
mime.type image/jpeg
```

Or, if your pages are generated by an application such as an NSAPI plug-in or CGI script, you can allow all mime types to be cached except for the mime types associated with the application. For example, to cache all mime types except those generated by CGI scripts, add the following lines to the configuration file:

```
mime.nocache
mime.type magnus-internal/cgi
```

You must restart the web server if you make changes to the configuration file.

If a page is not cachable, such as a dynamically created page or other pages marked by NSAPI as uncachable, you cannot cache it even if you specify it in the configuration file as cachable.

Allowing Equal Access to a System's Resources

If a web server hosts multiple virtual web sites, HP AC can be configured to balance traffic between each site, allowing each site to share the system's resources equally.

The description or pattern of the web servers' document roots for each virtual server must be configured in the `/opt/ns-server_name/server_id/config/hpac/config.ac` configuration file. HP AC uses the `isolate.pattern` parameter to monitor each site's traffic. You must restart the web server after editing the configuration file.

For example, a system is hosting two web sites for Company A and Company B. Company A's home page is located at `/web_pages/company_a/index.html` and Company B's home page is located at `/web_pages/company_b/index.html`. To monitor each site's traffic, add the following to the configuration file:

```
isolate.pattern /web_pages/%[^\s/]
```

Refer to documentation on the `sscanf` command for more information on constructing the pattern.

After editing the configuration file, restart the web server.

Detecting Failed Servers in a Cluster

When a web server is running, it periodically broadcasts a message that it is up. When a web server is shut down, it broadcasts a message that is unavailable. When a web server is not running, it does not broadcast a message.

Each web server in a cluster listens for these broadcast messages. Based on the messages received or not received, it will mark each system as up or down.

The web server marks a system as up when it receives the appropriate broadcast message from that system. The web server marks a system as down when it receives the appropriate message or it fails to receive a message after a specified number of broadcast intervals.

In the HP AC

`/opt/ns-server_name/server_id/config/hpac/config.ac` configuration file, you can configure how often messages are sent (broadcast interval) and how many messages can be missed before the system is marked as down. The parameters are:

```
cluster.keepalive.interval 60  
cluster.keepalive.misses 3
```

where the interval is measured in seconds. The values given are the default values.

Other Configuration Considerations

If you configure these parameters, the web server must be restarted. Also, all systems in the cluster should use the same values for these parameters.

Decreasing the broadcast interval decreases the time it takes to detect a failure. However, it also increases broadcast traffic.

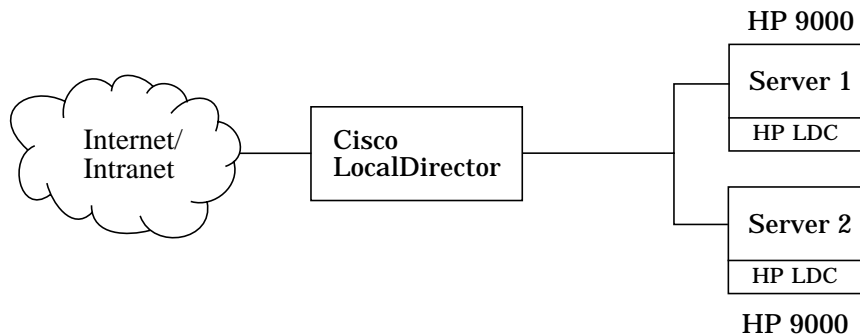
Sometimes, not all broadcast messages are received. Therefore, the number of misses should be set to a value greater than one. However, in a large cluster, removing a machine falsely does not severely impact the cluster's capacity and improves the responsiveness to failures.

The maximum time to detect a failure can be determined by multiplying the broadcast interval by the number of misses. For the default values, the maximum failure detection time is three minutes.

3 HP LocalDirector Controller

What is HP LocalDirector Controller?

HP LocalDirector Controller (HP LDC) works with Cisco LocalDirector by automatically generating the information needed by Cisco LocalDirector to load balance TCP/IP traffic across multiple servers. HP LDC, specifically the HP LocalDirector Control daemon (`ldc_agt`), must be configured and running on each system that is managed by Cisco LocalDirector. Information regarding Cisco LocalDirector can be found at the URL <http://www.cisco.com/warp/public/751/loDIR/>.



The information generated by HP LDC is referred to as “weights” in this document. The weights help Cisco LocalDirector determine which system has the resources to process an incoming request.

Initial weights are generated by the Webstone benchmark tool. Webstone generates traffic to a system and calculates the throughput based on getting different-sized static pages in three minutes. The throughput is used to calculate the initial weights.

HP LDC adjusts the weights depending on CPU usage and system load. Both the CPU usage and system load (measured by a one minute average job queue) must reach their configured thresholds before the weights are lowered. The weights will range between one and the initial weights.

Installing HP LocalDirector Controller

HP LocalDirector Controller (HP LDC) is part of the HP Domain ServiceControl (HP DSC) product. You do not need to load the entire HP DSC product if you only want to use HP LDC.

Installation Requirements

Hardware Requirements

- HP 9000 computer system
- 32 MB of memory
- CD-ROM drive to load the software

Software Requirements

- HP-UX 10.20 or 11.0
- HP DESMS
- An HTTP server such as Netscape Enterprise or FastTrack Server version 3.0 or later
- A web browser that supports JDK 1.1, such as Netscape Navigator 4.05 or Microsoft Internet Explorer 4.0

Other Requirements

- A system already running Cisco LocalDirector version 2.1 or later

Installing the Software

To install HP LDC, run `swinstall` (as root), double-click on the bundle J1592AA to view the software, and select the HP LDC software.

Configuring Cisco LocalDirector

Because HP LocalDirector Controller sends information to Cisco LocalDirector using SNMP traps, the *snmp-server host* parameter must be configured in the Cisco LocalDirector.

On a system managed by Cisco LocalDirector, run `netstat -rn` and check the routing table information. Use the IP address associated with the Cisco LocalDirector gateway to configure the *snmp-server host* parameter

Refer to Cisco LocalDirector's manuals for information on setting this parameter and other configuration information.

Configuring HP LocalDirector Controller

There are two components you need to configure before using HP LDC: the HP LDControl server and the HP LDControl daemon.

You must configure HP LDC using the HP Domain Enterprise Management System (HP DESMS). Instructions on how to start HP DESMS are included in this manual. For more information on HP DESMS, please refer to the *Getting Started with HP Domain Software* manual.

Starting HP DESMS

If this is the first time you are running HP DESMS, you must run the `/opt/hpwebsuite/hpsetup startup` script first.

To start HP DESMS, do the following:

1. Start the HP administration server. Type the following:
`/opt/hpwebsuite/start-admin`
2. Start a browser that supports frames, is Java-enabled, and supports JDK 1.1.
3. Go to the URL

`http://hostname:8181`

where *hostname* is the name of the system on which you are running the HP administration server.

Reload this page if you have visited this page prior to this installation.

4. Enter the administration user ID (`hpadmin`) and password (which you configured using the `/opt/hpwebsuite/hpsetup startup` script).

If you are starting HP DESMS for the first time, the page to configure the HP LDControl daemon is displayed.

Configuring the HP LDControl Server

After starting HP DESMS, click on “Service Control” from the side menu and then “Server Configuration.” The following page displays:

The screenshot shows a configuration window titled "Server Configuration". It contains three input fields: "Server Directory" with the value "/opt/ns-enterprise3", "Port Number" with the value "80", and "Document Directory" with the value "/opt/ns-enterprise3/docs-https-default". Below these fields are "OK" and "Cancel" buttons. At the bottom of the window is a large, empty rectangular box.

Set the following parameters to configure the HP LDControl server:

Table 3-1

HP LDControl Server Configurable Parameters

Parameter	Description
Server Directory	The home directory of the web server. For example, /opt/ns-enterprise3.
Port Number	The port number used by the web server. This must be a non-SSL port number. For example, 80.
Document Directory	The document root directory of the web server. For example, /opt/ns-enterprise3/docs-https-default.

The HP LDControl daemon must be started/restarted for these values to take effect. See “Starting and Stopping the HP LDControl Daemon” on page 47 for more information on starting HP LDC.

NOTE

The web server configured in the `/etc/hpldc/webstone/conf/server-list` file must be running before the HP LDCControl daemon is started.

Configuring the HP LDControl Daemon

After starting HP DESMS, select “Service Control” from the side menu and then “Configuration.” The following page displays:

The screenshot shows a configuration window for the HP LDControl Daemon. It contains the following fields and controls:

- Local Director Host Name:** A text input field containing the value `localhost`.
- Dynamic Weight:** A dropdown menu with the value `on` selected.
- CPU Threshold:** A text input field containing the value `90`, with a range of `(50..100)` indicated to the right.
- Max Queue Threshold:** A text input field containing the value `2`, with a range of `(1..10..0)` indicated to the right.
- Time Interval:** A text input field containing the value `5`, with a range of `(5..300) secs` indicated to the right.
- Buttons:** `OK` and `Cancel` buttons are located at the bottom of the configuration area.
- Footer:** A long, empty rectangular box is located at the bottom of the page.

Configuring HP LocalDirector Controller

Set the following parameters to configure the HP LDControl daemon:

Table 3-2**HP LDControl Daemon Configurable Parameters**

Parameter	Description
LocalDirector Hostname	The hostname of the system on which Cisco LocalDirector is running.
Dynamic Weight	Starts/Stops the HP LDControl daemon. When activated, the daemon will monitor the system and adjust the weights, if necessary. When deactivated, the HP LDControl daemon sends the initial weights to Cisco LocalDirector and exits.
CPU Threshold	The highest percentage the CPU usage should reach before the weight of the server is automatically decreased. Note that the run queue threshold (average load) must also be reached before the weights are adjusted. Default: 80 Range: 50 - 100% (integer only)
Run Queue Threshold	The highest amount the average load should reach before the weight of the server is automatically decreased. Note that the CPU usage threshold must also be reached before the weights are adjusted. Default: 2 Range: 1.0 - 10.0 (real number)
Time Interval	How often, in seconds, the thresholds are checked and adjustments are made, if needed, to the weights. Default: 5 seconds Range: 5 - 300 seconds (integer only)

Click on “OK” to save your changes. The values are saved in the file `/etc/opt/hpldc/conf/ldc.conf`.

NOTE

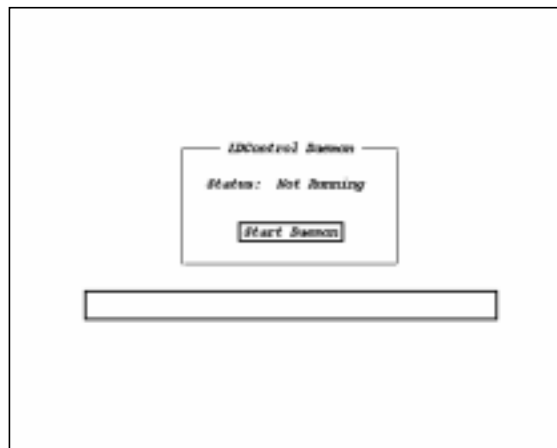
The web server configured in the `/etc/hpldc/webstone/conf/server-list` file must be running before the HP LDCControl daemon is started.

Starting and Stopping the HP LDCControl Daemon

The HP LDCControl daemon can be started and stopped from HP DESMS. Select “Service Control” from the side menu and then “Start/Stop Daemon.” A page similar to the following displays (what is displayed depends on the status of the HP LDCControl daemon: running or not running):

NOTE

The web server configured in the `/etc/hpldc/webstone/conf/server-list` file must be running before the HP LDCControl daemon is started.



To start or stop the HP LDCControl daemon from the command line, type:
`/opt/hpldc/bin/ldc_agt`

Configuring a System with More than One LAN Card

If the system on which you are running the web server has more than one LAN card, you must set the *snmp-server host* and *real* parameters in Cisco LocalDirector. Do the following:

1. On the system with multiple LAN cards, run `netstat -rn` and check the routing table information. Use the IP address associated with the Cisco LocalDirector gateway to configure the *snmp-server host* parameter of Cisco LocalDirector.
2. On the system with multiple LAN cards, run `nslookup local_hostname` and use this IP address to configure the *real* parameter of Cisco LocalDirector.

Refer to Cisco LocalDirector's manuals for information on setting these parameters.

Troubleshooting the HP LD Controller

General Troubleshooting

If an error occurs, check the following two log files for error messages, causes, and recommended actions:

1. `/var/opt/hpldc/ldc.log`
2. `/var/opt/hpldc/webstone.log`

Problems Accessing the HP LD Controller from HP DESMS

Table 3-3

Browser Error Message and Action

Error Message	Action
Applet ConfApplet can't start.	The browser that you are using to administer the system must support JDK 1.1 in order to run HPLDC. Two browsers that support JDK 1.1 are Netscape Navigator 4.05 and Microsoft Internet Explorer 4.0.

HP LDControl Daemon

If the HP LDControl daemon dies, check for the following error messages in the `/var/opt/hpldc/ldc.log` file. Complete the listed actions for the specified messages.

Table 3-4

HP LDControl Daemon Error Log Messages and Actions

Error Message	Action
Cannot open file <i>filename</i> .	Check that the file <i>filename</i> exists and that the HP LDControl daemon has write permission to its directory.

Table 3-4 HP LDControl Daemon Error Log Messages and Actions

Error Message	Action
Server might be down. Initial weight set to 0. Exiting.	<ol style="list-style-type: none">1. Make sure the web server is running on the system. Restart/Start the web server and then restart the HP LDControl daemon.2. Check the parameters in the file <code>/etc/hpldc/webstone/conf/server-list</code>.3. In <code>/etc/passwd</code>, check for the following entry: <code>webstone:trKI3jiXvewTw:33:1:::/tmp:/sbin/sh</code> If this entry does not exist, add it to the file.4. If the web server is running, make sure it is using a non-SSL port. The HP LDControl daemon must use a non-SSL port to generate the initial weight.
Cannot send weight to LocalDirector.	<ol style="list-style-type: none">1. The <code>snmp-server host</code> parameter must be set on the system running Cisco LocalDirector. Refer to Cisco LocalDirector's manual for more information on setting this parameter.2. Check that <code>LDHostname</code> is set to the system running Cisco LocalDirector in the <code>/etc/opt/hpldc/conf/ldc.conf</code> file.