

I n s i d e W e b O b j e c t s

---

# Deploying WebObjects Applications



January 2002

Apple Computer, Inc.  
© 2001–2002 Apple Computer, Inc.  
All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, mechanical, electronic, photocopying, recording, or otherwise, without prior written permission of Apple Computer, Inc., with the following exceptions: Any person is hereby authorized to store documentation on a single computer for personal use only and to print copies of documentation for personal use provided that the documentation contains Apple's copyright notice.

The Apple logo is a trademark of Apple Computer, Inc. Use of the "keyboard" Apple logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

No licenses, express or implied, are granted with respect to any of the technology described in this book. Apple retains all intellectual property rights associated with the technology described in this book. This book is intended to assist application developers to develop applications only for Apple-labeled or Apple-licensed computers.

Every effort has been made to ensure that the information in this document is accurate. Apple is not responsible for typographical errors.

Apple Computer, Inc.  
1 Infinite Loop  
Cupertino, CA 95014  
408-996-1010

Apple, the Apple logo, Mac, Macintosh, and WebObjects are trademarks of Apple Computer, Inc., registered in the United States and other countries.

Enterprise Objects and Enterprise Objects Framework are trademarks of NeXT Software, Inc., registered in the United States and other countries.

Java is a registered trademark of Sun Microsystems, Inc. in the United States and other countries.

Simultaneously published in the United States and Canada

**Even though Apple has reviewed this manual, APPLE MAKES NO WARRANTY OR REPRESENTATION, EITHER EXPRESS OR IMPLIED, WITH RESPECT TO THIS MANUAL, ITS QUALITY, ACCURACY, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. AS A RESULT, THIS MANUAL IS SOLD "AS IS," AND YOU, THE PURCHASER, ARE ASSUMING THE ENTIRE RISK AS TO ITS QUALITY AND ACCURACY.**

**IN NO EVENT WILL APPLE BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES RESULTING FROM ANY DEFECT OR INACCURACY IN THIS MANUAL, even if advised of the possibility of such damages.**

**THE WARRANTY AND REMEDIES SET FORTH ABOVE ARE EXCLUSIVE AND IN LIEU OF ALL OTHERS, ORAL OR WRITTEN, EXPRESS OR IMPLIED. No Apple dealer, agent, or employee is authorized to make any modification, extension, or addition to this warranty.**

**Some states do not allow the exclusion or limitation of implied warranties or liability for incidental or consequential damages, so the above limitation or exclusion may not apply to you. This warranty gives you specific legal rights, and you may also have other rights which vary from state to state.**

# Contents

Figures, Listings, and Tables 9

## **Chapter 1** About This Book 13

---

Why Read This Book 13  
Getting More Information 15  
What You Should Know 15  
Operating Systems Supported 15

## **Chapter 2** Introduction to WebObjects Deployment 17

---

The WebObjects Deployment Model 17  
The WebObjects Deployment Environment 21  
    Communication Paths 21  
    Deployment Tools 23  
Keeping Your Site Secure 26

## **Chapter 3** Installing the Software 29

---

Choosing What to Install 29  
    HTTP Adaptor Files 30  
    Deployment Files 30  
    Types of WebObjects Deployment Installations 31  
Environment Variables in UNIX Platforms 31  
Confirming That wotaskd Is Active 31  
Default Adaptor Installation 33  
Building Adaptors From Source 34

**Chapter 4 HTTP Adaptors 35**

---

Adaptors, Applications, and Hosts	36
Types of Adaptors	38
CGI Adaptors	38
API-Based Adaptors	39
State Discovery	39
Using a Multicast Request	41
Multicast Request	41
Host Polling	42
Using a Defined Host List	44
Using a Configuration File	44
The HTTP Adaptor Configuration File	44
Creating the HTTP Adaptor Configuration File	47
The WebObjects Adaptor Information Page	50
Customizing the HTTP Adaptor	52
Setting the Multicast Address and Port	52
Setting the Host List	53
Setting the Name of the HTTP Adaptor Configuration File	54
Setting Access to the WebObjects Adaptor Information Page	54
Setting an Alias for cgi-bin in the WebObjects URL	55
Setting the Document Root Path of the HTTP server	56

**Chapter 5 Managing Application Instances 57**

---

Configuration Files	57
Lifebeats	60
wotaskd Processes	61
Starting WebObjects Services	62
Starting WebObjects Services Automatically	62
Starting Monitor Manually	62
Controlling WebObjects Services in Solaris	64
Windows 2000	65

**Chapter 6 Deployment Tasks 67**

---

Setting Up Hosts	68
------------------	----

## C O N T E N T S

Adding a Host	68
Configuring a Host	70
Viewing a Host's Configuration	71
Installing Applications	74
Setting Up Applications	75
Adding an Application	76
Configuring an Application	76
New Instance Defaults	77
Application Settings	79
Scheduling	80
Email Notifications	81
Load Balancing and Adaptor Settings	82
Adding Application Instances	83
Configuring Instances	88
Instance Settings	90
Adaptor Settings	90
Setting a Password for the Instance Statistics Page	90
Configuring Sites	91
Setting Monitor Preferences	94
Monitor Password	95
Detail-View Refresh Settings	96
Load Balancing	96
Deploying Multiple Sites	97

### **Chapter 7** Application Administration 99

---

Monitoring Activity	99
Monitoring Application Performance	99
The Application Detail Page	101
The Instance Statistics Page	102
Logging and Analyzing Application Activity	104
Logging and Analyzing Adaptor Activity	105
Creating the Adaptor Log File	105
Analyzing the Adaptor Log File's Contents	106
Improving Performance	107

**Chapter 8**    **Deployment Settings Reference**    109

---

Application Configuration Settings	109
Application Settings	110
Adaptor	110
Adaptor Threads	110
Listen Queue Size	110
Maximum Adaptor Threads	110
Minimum Adaptor Threads	111
Name	111
Phased Startup	111
Starting Port	111
Time Allowed for Startup	111
Instance Settings	111
Additional Arguments	112
Auto-Open In Browser	112
Auto Recover	112
Caching Enabled	112
Debugging Enabled	112
ID	112
Lifebeat Interval	112
Minimum Active Sessions	113
Output Path	113
Path	113
Port	113
Email Notification Settings	114
Load Balancing and Adaptor Settings	114
Connect Timeout	114
Connection Pool Size	114
Dormant	114
Load-Balancing Scheme	114
Receive Buffer Size	115
Receive Timeout	115
Redirection URL	115
Retries	115
Send Buffer Size	115
Send Timeout	115
URL Version	115

# C O N T E N T S

Scheduling Settings	116
Is Scheduled	116
Graceful Scheduling	116
Types of Schedule	116
Command-Line Arguments	116
General Command-Line Arguments	117
WOApplicationName	117
WODirectConnectEnabled	117
WOHost	117
WOLifebeatDestinationPort	118
WOLifebeatEnabled	118
WOLifebeatInterval	118
WONoPause	118
WOOutputPath	118
WOPort	119
WOREcordingPath	119
WOSessionTimeOut	119
WOStatisticsPassword	119
Monitor and wotaskd Command-Line Arguments	120
WODeploymentConfigurationDirectory	120
wotaskd Command-Line Arguments	120
WOAssumeApplicationIsDeadMultiplier	120
WOMulticastAddress	120
WOREspondsToMulticastQuery	121
WOSavesAdaptorConfiguration	121

## **Appendix A Special Deployment Issues** 123

---

Deployment Issues With Java Client Applications	123
Deploying WebObjects 4.5.1 and WebObjects 5 Applications	124

## Glossary 127

---

## Index 131

---

C O N T E N T S



# Figures, Listings, and Tables

## Chapter 2 Introduction to WebObjects Deployment 17

---

Figure 2-1	WebObjects deployment model	19
Figure 2-2	WebObjects deployment model—multiple instances of an application	20
Figure 2-3	Deployment using two computers	21
Figure 2-4	The data path of a WebObjects deployment	22
Figure 2-5	The control path of a WebObjects deployment	23
Figure 2-6	The symbols used to represent the data path and the control path	23
Figure 2-7	Two sites deployed on one computer	24
Figure 2-8	Two sites deployed on two computers	25

## Chapter 3 Installing the Software 29

---

Figure 3-1	The configuration page of a wotaskd process	32
Table 3-1	The WebObjects deployment and administration tools	30
Table 3-2	The adaptors installed in each platform	33

## Chapter 4 HTTP Adaptors 35

---

Figure 4-1	Deployment on one computer, using one adaptor	36
Figure 4-2	Deployment on one computer using two adaptors	37
Figure 4-3	Deployment using three computers using one adaptor	37
Figure 4-4	Dynamic site configuration using multicast request and polling	43
Figure 4-5	Copying the information that makes up the HTTP adaptor configuration file	48
Figure 4-6	Creating and saving the HTTP adaptor configuration file	49
Figure 4-7	The WebObjects Adaptor Information page	51
Listing 4-1	A WebObjects adaptor configuration file	45
Listing 4-2	Structure of the HTTP adaptor configuration file	45

Table 4-1	API-based adaptors and supported platforms	39
Table 4-2	The properties of the HTTP adaptor configuration file	46

---

**Chapter 5**   **Managing Application Instances**   57

---

Figure 5-1	WebObjects configuration-file distribution	59
Figure 5-2	Monitor—empty Applications page	64
Listing 5-1	Starting Monitor	63

---

**Chapter 6**   **Deployment Tasks**   67

---

Figure 6-1	The Hosts page	69
Figure 6-2	Newly added host in Monitor	70
Figure 6-3	Host configuration page	71
Figure 6-4	Host configuration information page	72
Figure 6-5	Adding an application using Monitor’s Applications page	76
Figure 6-6	The New Instance Defaults section of the application configuration page	78
Figure 6-7	The Application Settings section of the application configuration page	80
Figure 6-8	The Scheduling section of the application configuration page	81
Figure 6-9	The Email Notifications section of the Application Configuration page	82
Figure 6-10	The Load Balancing and Adaptor Settings section of the application configuration page	83
Figure 6-11	The Applications page with one application	84
Figure 6-12	The application detail page	85
Figure 6-13	The application detail page after an instance has been added	86
Figure 6-14	The application detail page with two instances added	87
Figure 6-15	Instance configuration page	89
Figure 6-16	Setting a password for an instance’s statistics page	91
Figure 6-17	The site configuration page	93
Figure 6-18	The preferences page	94
Figure 6-19	Login page displayed by Monitor on a password-protected site	95

# F I G U R E S   A N D   T A B L E S

Figure 6-20	Page returned by wotaskd when the site is password-protected	96
Figure 6-21	Multiple application environments on one computer	98

## **Chapter 7**   **Application Administration**   99

---

Figure 7-1	The Applications page	100
Figure 7-2	The application detail page	101
Figure 7-3	The instance statistics page—part 1 of 2	103
Figure 7-4	The instance statistics page—part 2 of 2	104

F I G U R E S   A N D   T A B L E S

# About This Book

---

This book describes the tools and techniques that system administrators and website managers perform to deploy WebObjects applications. The WebObjects Deployment package allows you to deploy applications developed with the WebObjects Development package, so that they can be accessed through an HTTP server. You need a WebObjects deployment license to deploy WebObjects applications.

The book is intended primarily for system administrators. Application developers can also benefit from the information it provides but it's not required reading for them. This edition of the book is intended for WebObjects 5.1.

## Why Read This Book

---

To deploy WebObjects applications and to administer a deployment, you need to become acquainted with the deployment model of WebObjects. This book shows you how your HTTP server interacts with the elements of a WebObjects deployment. It also explains what measures you should take to increase your site's performance.

WebObjects Deployment provides tools for most of the tasks you need to accomplish on a regular basis to maintain your site. If you prefer doing things manually, you can use the command line to start individual application instances or the deployment tools themselves.

This book has the following chapters:

## About This Book

- [Chapter 2, “Introduction to WebObjects Deployment”](#) (page 17), gives you an overview of the deployment approach taken with WebObjects 5. In addition, it lists the ways in which WebObjects Deployment helps you to maintain a secure site.
- [Chapter 3, “Installing the Software”](#) (page 29), explains which WebObjects Deployment components need to be installed on a computer, taking into account the computer’s purpose in your site.
- [Chapter 4, “HTTP Adaptors”](#) (page 35), describes the function of the HTTP adaptor in your site. It also describes how you customize the adaptors included in WebObjects Deployment if the default configuration does not suit your needs. State discovery is how the HTTP adaptor keeps track of the application instances of your site. The chapter describes the different ways that the adaptor can obtain that information and how you configure the adaptor to use one of those methods.
- [Chapter 5, “Managing Application Instances”](#) (page 57), introduces you to the deployment tools you use to configure and maintain your site. It also describes the mechanism used in WebObjects to ensure that application instances are always running, helping you maximize your site’s up time.
- [Chapter 6, “Deployment Tasks”](#) (page 67), explains how to perform configuration and maintenance tasks on your site. It also shows how to maintain multiple sites using the same hardware.
- [Chapter 7, “Application Administration”](#) (page 99), shows you how to monitor and improve your site’s performance.
- [Chapter 8, “Deployment Settings Reference”](#) (page 109), lists the deployment properties you can use to customize your site.
- [Appendix A, “Special Deployment Issues”](#) (page 123), lists issues to keep in mind when deploying Java Client applications. It also tells you what to do if you want to deploy WebObjects 4.5.1 applications together with WebObjects 5 applications.
- [Appendix B, “Document Revision History”](#) (page 125), lists changes made from previous editions of the book.

## Getting More Information

---

To get an overview of the WebObjects platform, you should read the book *Inside WebObjects: WebObjects Overview*. You can find general information about WebObjects at

<http://developer.apple.com/webobjects>.

These are other WebObjects documents that provide helpful information:

- *Installation Guide* includes system requirements, compatibility information, and installation instructions. (The *Installation Guide* is included in your WebObjects box.)
- *Post-Installation Instructions* lists tasks that you must perform after installing WebObjects Deployment to confirm that the installation was successful and to tailor it for your specific needs.

## What You Should Know

---

This book assumes you have a background in system administration. You must be familiar with the operation of your platform, especially how to use its command shell editor to issue commands. You must also be acquainted with the operation of your HTTP server software and TCP/IP networking. Knowledge of WebObjects application development is helpful, but not required.

## Operating Systems Supported

---

The following are the supported deployment platforms for WebObjects 5:

- **Mac OS X Server:** Mac OS X Server 10.1.1.

# C H A P T E R 1

## About This Book

- **Solaris:** Sun Solaris 2.8.
- **Windows 2000:** Microsoft Windows 2000 Pro with J2SE 1.3.1.

WebObjects applications can be deployed on other platforms running J2SE 1.3 or later.



# Introduction to WebObjects Deployment

---

This chapter introduces the essential concepts and tools you use when you deploy WebObjects applications.

The chapter addresses the following topics:

- “The WebObjects Deployment Model” (page 17) introduces you to the WebObjects way of deploying applications. It explains how the users of your applications send requests to application instances running on your site and how responses (Web pages) are generated and sent back to users.
- “The WebObjects Deployment Environment” (page 21) describes the functions of several elements (both in your platform and in WebObjects Deployment) in a site.
- “Keeping Your Site Secure” (page 26) lists the security-minded features available in WebObjects Deployment.

## The WebObjects Deployment Model

---

A WebObjects deployment has six major parts:

- **Client:** a user’s Web browser or the client-side of a Java Client application.
- **HTTP server:** application that receives HTTP requests from clients and sends responses back to them.
- **HTTP adaptor:** application that serves as an interface between your HTTP server and your application instances. The HTTP adaptor routes requests from the HTTP server to the appropriate instance and sends the responses generated

## Introduction to WebObjects Deployment

back to the HTTP server. The adaptor does this while performing **load balancing** to distribute an application's users among its active instances. Load balancing helps to spread the user load of your site evenly across your application hosts.

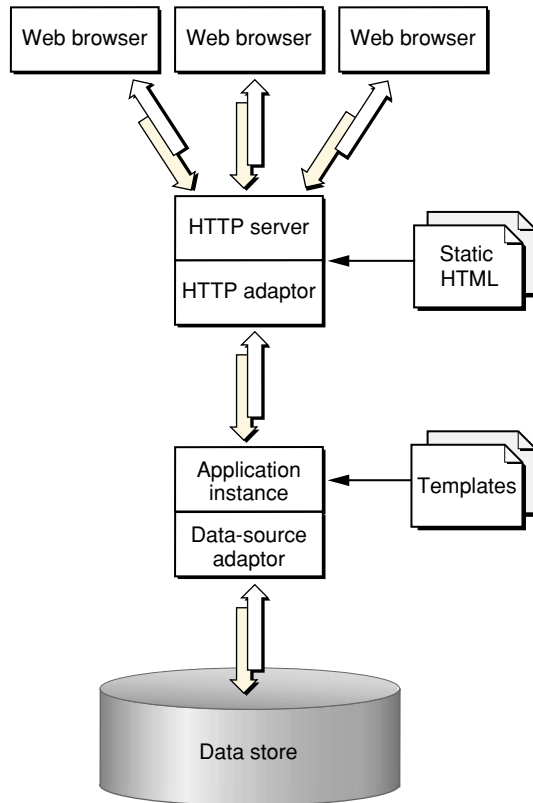
- **Application instances:** individual processes that receive requests from the HTTP adaptor and send responses back to it. To create a response, an instance can perform calculations, or save or retrieve data from a data source.
- **Data-source adaptor:** interface between an application instance and a data source. WebObjects includes a **JDBC** adaptor, allowing your applications to connect to any JDBC-compliant data store.

For JDBC connectivity, your database needs a JDBC driver, which you obtain from your database vendor. WebObjects applications can connect to databases that use Type 2 (partly Java) or Type 4 (all Java) JDBC drivers. The JDBC adaptor included with WebObjects Deployment has been certified to work adequately with Type 4 drivers. Type 2 drivers may require special configuration for them to work properly with the adaptor. If your database provides a Type 2 driver, consult with your database vendor to determine how it needs to be configured to work properly with a JDBC adaptor.

- **Data store:** the mechanism that your applications use to store persistent data. Consult with your database vendor to obtain configuration and optimization details.

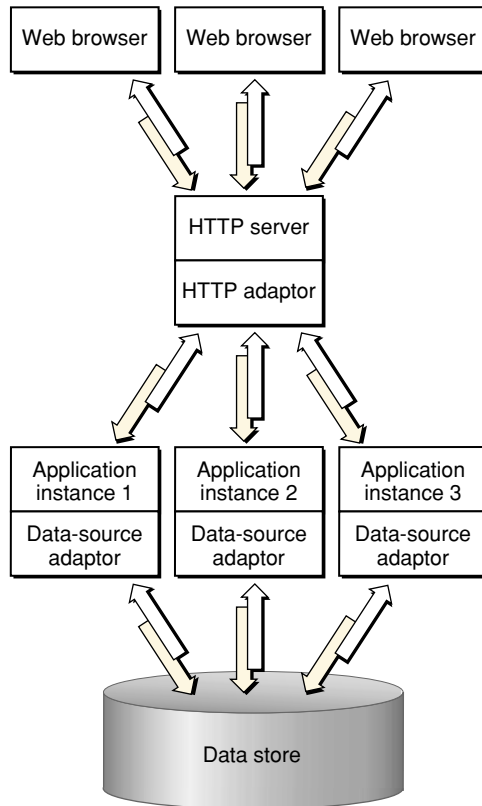
When an application user sends a request through a Web browser to your HTTP server, the server forwards the request to the HTTP adaptor. The adaptor then determines which application instance should process the request and forwards the request to it. When the application instance receives the request, it performs all the necessary processing to produce a response (a new Web page). The instance then sends the response page to the adaptor, which forwards it to the HTTP server. The HTTP server then forwards the response page to the user's Web browser. This process is illustrated in [Figure 2-1](#).

**Figure 2-1** WebObjects deployment model




Notice that both the application instance and the HTTP server contribute to the response page's content. The instance uses templates and logic to generate the HTML code for dynamic pages, while the HTTP server provides the content of images contained in those pages. The server can also dispense static pages.

The number of instances of your application necessary to support its users depends on the number of users that connect to your application concurrently. In some cases a single instance is adequate. When one instance is not able to process requests in a timely manner, additional instances can solve the problem. This way, the amount of user-state information that a single instance stores is reduced. In addition, with less state to keep track of, an instance can process requests faster. [Figure 2-2](#) shows a site with one host running multiple instances of an application.

**Figure 2-2** WebObjects deployment model—multiple instances of an application

However, adding instances of your application to a host may not be the most effective solution. Eventually a point of diminishing returns will be reached, where adding instances actually decreases your application's performance. In such a case, you should consider adding additional application hosts that run the extra instances required to handle the increased traffic to your site. [Figure 2-3](#) shows how a site with two computers, one acting as an HTTP server and application host, and the other just as an application host would look.

**Figure 2-3** Deployment using two computers

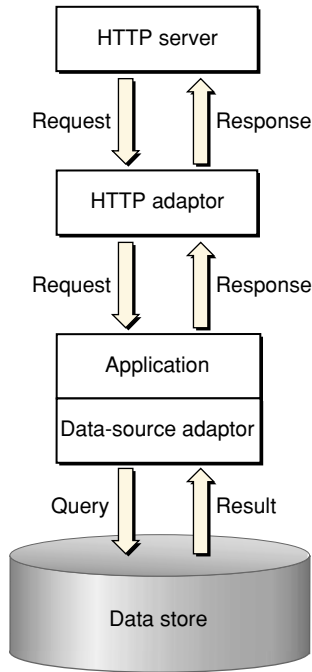
## The WebObjects Deployment Environment

You need to master two important aspects of WebObjects Deployment: the communication paths of client requests and server activity, and the deployment tools you use to configure your site.

### Communication Paths

Communication among the elements that make up a deployment occurs in two paths: the data path and the control path.

A client HTTP request takes the data path after it reaches your HTTP server. [Figure 2-4](#) shows how an HTTP request that your HTTP server receives is passed to the elements that generate the response.

**Figure 2-4** The data path of a WebObjects deployment

**Monitor** requests take the control path to propagate configuration changes to application hosts and, ultimately, application instances. These include adding application instances and starting and stopping instances according to a schedule that you define. The HTTP adaptor can obtain site information by polling **wotaskd** (WebObjects task daemon) processes or by reading the adaptor configuration file. (See “[Deployment Tools](#)” (page 23) for information about Monitor and wotaskd.) [Figure 2-5](#) shows the control path.

**Figure 2-5** The control path of a WebObjects deployment

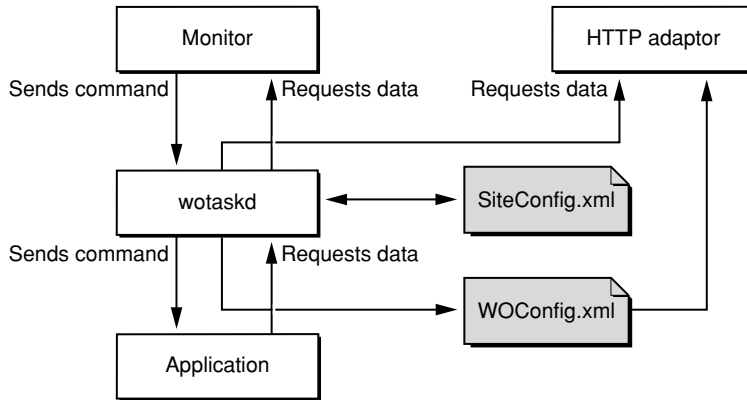



Figure 2-6 shows how the data path and control path are differentiated in the rest of the book.

**Figure 2-6** The symbols used to represent the data path and the control path



## Deployment Tools

The main tools you use to manage your site are **wotaskd** and **Monitor**. Normally, one **wotaskd** process runs on each application host. If you want to concurrently deploy multiple sites on the same hardware, you can configure a computer to run more than one **wotaskd** process. This essentially provides you with several independent application hosts per computer.

You manage a group of application hosts using **Monitor**, a tool that uses your Web browser as its user interface. **Monitor** lets you set, among other things, instance scheduling and the load-balancing scheme to be used for each application. Because each **Monitor** process maintains state information locally, you must run only one instance of **Monitor** per site. Figure 2-7 shows two application sites on one computer.

**Figure 2-7** Two sites deployed on one computer

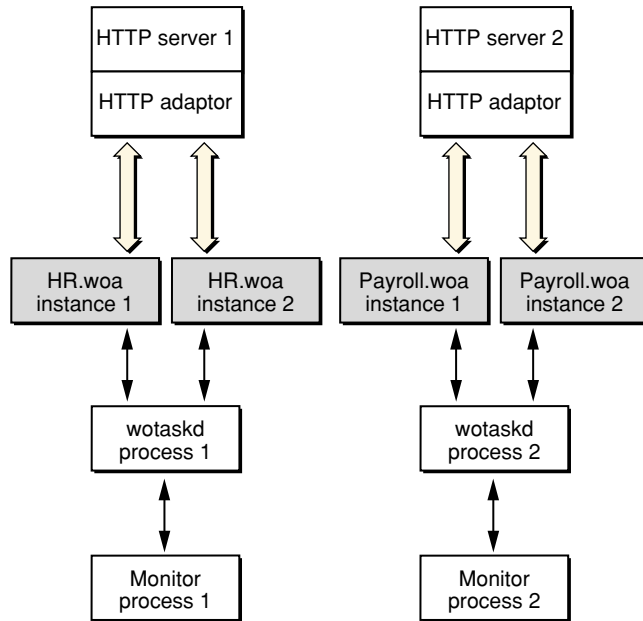
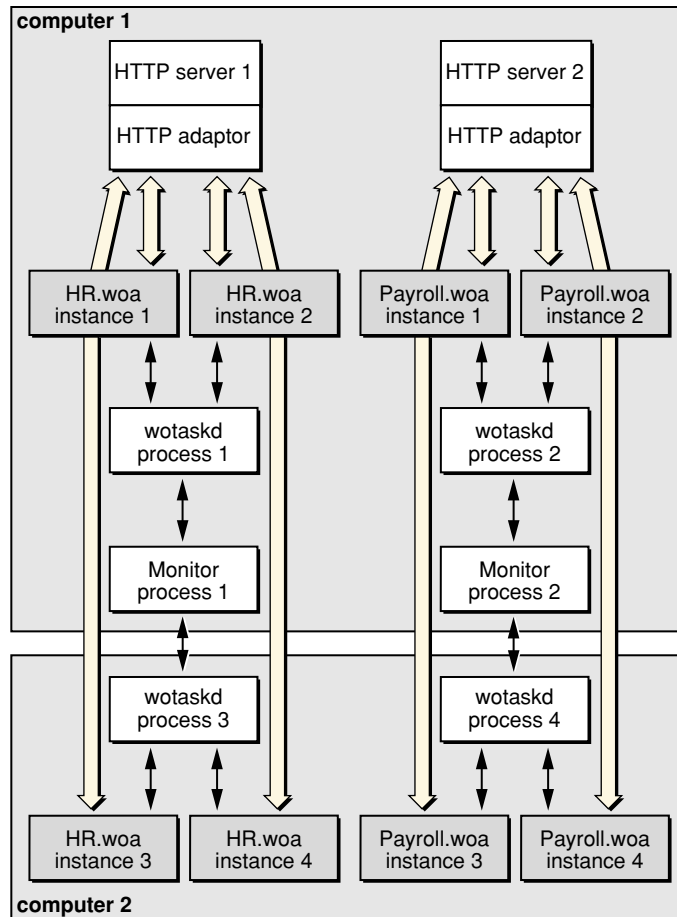


Figure 2-8 shows how you can distribute application instances among two computers.



**Figure 2-8** Two sites deployed on two computers



After you configure your site using Monitor, it enforces that configuration by performing tasks such as stopping and restarting application instances according to a schedule you set and sending email notifications when problems arise. The HTTP adaptor performs load balancing across the instances of each application on your site.

For detailed information on the subjects introduced above, see the following chapters or sections:

## Introduction to WebObjects Deployment

- “HTTP Adaptors” (page 35) shows you the different ways in which you can configure the WebObjects HTTP adaptor.
- “Deployment Tasks” (page 67) describes how you use Monitor to configure your site.
- “Setting Up Hosts” (page 68) describes how you use Monitor to add application hosts to your site.
- “Configuration Files” (page 57) shows you how the configuration you define in Monitor is distributed among the application hosts of your site.
- “wotaskd Processes” (page 61) explains how wotaskd processes communicate with and manage application instances.
- “Lifebeats” (page 60) explains how application instances communicate with a wotaskd process.
- “Deploying Multiple Sites” (page 97) explains how to configure your platform to deploy multiple sites concurrently.
- “Load Balancing” (page 96) describes how load balancing works and lists the algorithms that the HTTP adaptor can use to implement it.

## Keeping Your Site Secure

---

In a WebObjects deployment, you have several features at your disposal to enhance the security of your site:

- split-installation of applications (application files and HTTP-server resources). By installing application-related files in two locations, you can put sensitive information (such as business logic) into protected locations. Nonsensitive resources (such as image files) can be installed on the HTTP server’s `Document Root` directory. For more information, see “Installing Applications” (page 74).
- restricted access to deployment tools. “Monitor Password” (page 95) explains how you can password-protect access to Monitor and wotaskd through a single page.

### Introduction to WebObjects Deployment

- restricted access to development application instances. If your computing environment supports both the development and deployment of applications through the same HTTP server, access to development instances is restricted by the HTTP adaptor. See [“Viewing a Host’s Configuration”](#) (page 71) for details.
- ability to disallow direct connections (through host name and port number) to an application instance. With direct connect you can connect to an instance with the following URL:

```
http://myhost:1234
```

When you disallow direct connect for an instance, the only way to connect to it is through an HTTP server. For more information, see [“WODirectConnectEnabled”](#) (page 117).

- restricted access to application-instance statistics. Agents external to your organization can use the statistics that your application instances produce to get privileged information. To avoid this, access to the instance statistics page is restricted. See [“Setting a Password for the Instance Statistics Page”](#) (page 90) for details.
- restricted access to HTTP adaptor information (the WebObjects Adaptor Information page) by default. This closes another potential hacker-entry point. For details, see [“Setting Access to the WebObjects Adaptor Information Page”](#) (page 54).

## C H A P T E R 2

### Introduction to WebObjects Deployment

# Installing the Software

---

This chapter provides an overview of WebObjects Deployment installation issues, including what elements of the software need to be installed on a computer, taking the computer's purpose into account. However, it does not give platform-specific instructions. For detailed installation instructions, see the *Installation Guide* included with your package.

The following topics are addressed:

- “Choosing What to Install” (page 29)
- “Environment Variables in UNIX Platforms” (page 31)
- “Confirming That wotaskd Is Active” (page 31)
- “Default Adaptor Installation” (page 33)
- “Building Adaptors From Source” (page 34)

## Choosing What to Install

---

When you perform a complete installation of WebObjects Deployment on a computer, two types of files are copied to its hard disk: HTTP adaptor files and deployment files.

## Installing the Software

## HTTP Adaptor Files

---

These are the files that allow your HTTP server to communicate with WebObjects application instances. The executable files of the HTTP adaptors are placed in the `/System/Library/WebObjects/Adaptors` directory.

The source files for the HTTP adaptors are placed in the `/Developer/Examples/WebObjects/Source/Adaptors` directory. Also included are HTML files containing instructions for building and installing each adaptor. The code of the adaptors is open source, so you can install them anywhere.

## Deployment Files

---

These files are divided into two groups:

- **Runtime environment.** The runtime environment of WebObjects is implemented in JAR (Java archive) files contained within **framework** (`.framework`) bundles. Frameworks are installed in the `/System/Library/Frameworks` directory. These frameworks are used by any WebObjects application, including the deployment tools of WebObjects.
- **Deployment tools.** WebObjects Deployment includes two deployment tools you use to configure and monitor your site. The files that make up these tools are placed in the `/System/Library/WebObjects/JavaApplications` directory. [Table 3-1](#) shows the purpose of each tool. For more information on the deployment tools, see “[Managing Application Instances](#)” (page 57).

---

**Table 3-1** The WebObjects deployment and administration tools

Filename	Application name	Purpose
<code>JavaMonitor.woa</code>	Monitor	Site configuration and administration
<code>wotaskd.woa</code>	wotaskd	Instance management

**Note:** Installing WebObjects Deployment on a computer requires a WebObjects Deployment license. Please read the license agreement before installing this package.

## Types of WebObjects Deployment Installations

---

Depending on the purpose of the computer you're installing the software on, there are three types of WebObjects Deployment installations you can perform:

- **HTTP server only.** On a computer that you want to use as the HTTP server computer but on which you do not intend to run application instances (including the deployment tools), you need to install only the HTTP adaptor files.
- **Application host only.** On computers that you intend to use only as application hosts, you need to install only the deployment tools. (If you do not plan to run Monitor on that computer, you can delete its files.)
- **HTTP server and application host.** When one computer can satisfy all your deployment needs or when you want an HTTP server computer to also run application instances, you need to install the adaptor files and the deployment tools.

**Note:** On Windows 2000 systems it's not possible to perform tailored WebObjects Deployment installations.

## Environment Variables in UNIX Platforms

---

On Solaris and other UNIX platforms, you must set the value of the `NEXT_ROOT` environment variable. It must point to the directory into which WebObjects Deployment is installed.

## Confirming That wotaskd Is Active

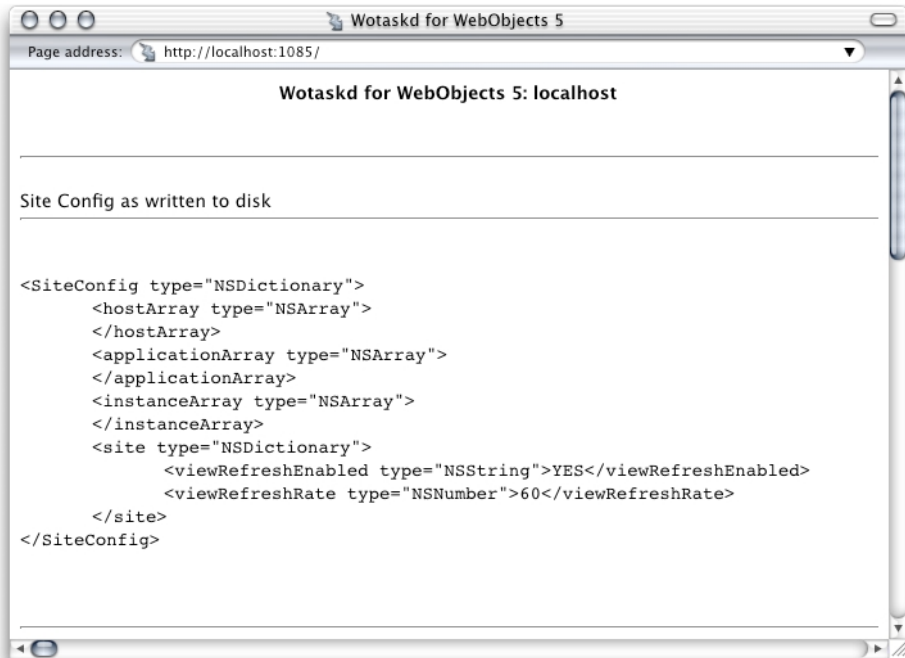
---

After the installation is complete (and if you chose to start **WOServices** in Solaris), a `wotaskd` process should be running on your computer. (WOServices ensures that `wotaskd` is running at all times.) To confirm that `wotaskd` is running, launch a Web

## Installing the Software

browser and enter this address: `http://localhost:1085`. You should see a page like the one in Figure 3-1. For more information, see “Viewing a Host’s Configuration” (page 71).

**Figure 3-1** The configuration page of a wotaskd process



If a wotaskd process isn’t running, these are possible reasons:

- The version of J2SE is not 1.3 or later.  
Upgrade to J2SE 1.3 (version 1.3.1 in Windows 2000).
- Port 1085 is already in use.  
Change the port of the process that’s using port 1085.
- The `NEXT_ROOT` environment variable is not set.



## Installing the Software

In platforms other than Mac OS X Server and Windows 2000, make sure that the `NEXT_ROOT` environment variable is set for the user under which you're trying to run `wotaskd`. This also applies to application instances.

## Default Adaptor Installation

---

Depending on your deployment platform, several HTTP adaptors are installed by default. Table 3-2 lists the adaptors that are installed on each platform.

**Table 3-2** The adaptors installed in each platform

	Apache	CGI	ISAPI	NSAPI
Mac OS X Server	x	x		
Solaris	x	x		x
Windows 2000		x	x	x

In Mac OS X Server and Solaris, the Apache adaptor is active by default. Requests in the form `http://.../cgi-bin/WebObjects/` are handled by the Apache adaptor. If you disable the Apache module, then such requests are handled by the CGI adaptor.

To install the ISAPI adaptor, copy the `WebObjects.dll` file to your HTTP server's `scripts` directory. The ISAPI adaptor is loaded into the HTTP server the first time a request of the form `http://.../scripts/WebObjects.dll/` is received. It then remains active until the server is stopped.

## Building Adaptors From Source

---

If you want to use a different HTTP adaptor from the one installed by default on your platform, you can find building and installation instructions for all the adaptors included in *WebObjects Deployment* in `/Developer/Examples/WebObjects/Source/Adaptors/InstallationInstructions.html`. **To build adaptors from the source code provided, refer to `BuildingInstructions.html` in the same directory.**

# HTTP Adaptors

---

This chapter provides detailed information about the HTTP adaptors that are included in a WebObjects Deployment installation. The HTTP adaptor is an important piece of an application site. It sits between your HTTP server and your application instances. It forwards requests from the HTTP server to the appropriate application instance and responses from the instance back to the HTTP server. It also performs load balancing between instances of an application.

This chapter addresses the following topics:


- [“Adaptors, Applications, and Hosts”](#) (page 36) provides a high-level view of the interaction between the HTTP adaptor, application instances, and application hosts.
- [“Types of Adaptors”](#) (page 38) explains the differences among the two types of HTTP adaptors that you can use on your site.
- [“State Discovery”](#) (page 39) describes how to configure the HTTP adaptor to obtain your site’s state dynamically or using a configuration file. It also explains how to change a dynamic configuration into a static one.
- [“The WebObjects Adaptor Information Page”](#) (page 50) shows an example of the Web browser page that displays information about an HTTP adaptor.
- [“Customizing the HTTP Adaptor”](#) (page 52) summarizes all the settings available for HTTP adaptors.

## Adaptors, Applications, and Hosts

---

The HTTP adaptor forwards requests from an HTTP server to application instances and returns responses from instances back to the server. You may need to have more than one instance of a given application to support a large number of concurrent users. [Figure 4-1](#) illustrates a simple site, implemented with one computer. It serves two applications, with two instances for each application.

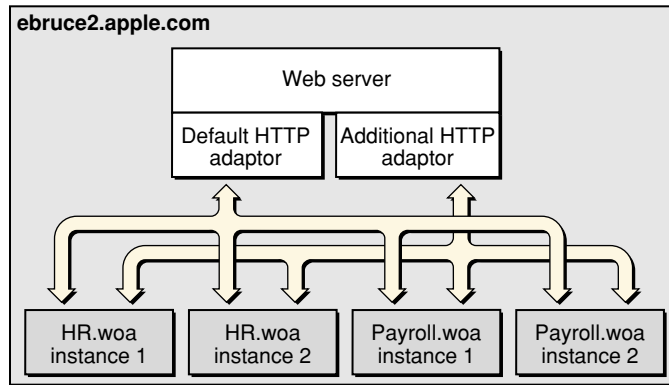
**Figure 4-1** Deployment on one computer, using one adaptor



Although the WebObjects installation provides several HTTP adaptors, only one is active by default (see [“Default Adaptor Installation”](#) (page 33) for details). However, an application instance can communicate with an adaptor other than the active adaptor.

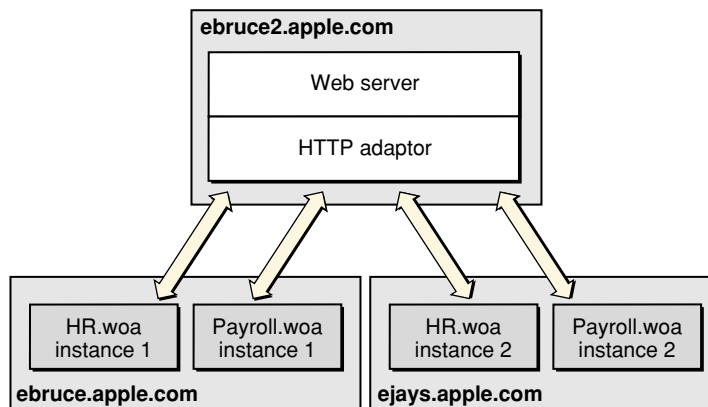
[Figure 4-2](#) depicts an application site running on one computer, using two adaptors.

**Figure 4-2** Deployment on one computer using two adaptors



Most sites require multiple computers to ensure that an instance of a particular application is always available. In this kind of deployment, usually one computer runs the HTTP server and the HTTP adaptor, while one or more additional computers serve as application hosts. [Figure 4-3](#) illustrates an application site using three computers, one running the HTTP server and the adaptor, and the other two running application instances.

**Figure 4-3** Deployment using three computers using one adaptor



## HTTP Adaptors

The HTTP adaptor needs to periodically determine your site's state—which application instances are running. There are two ways in which the adaptor can obtain this information:

- **Dynamically:** The adaptor determines your site's state by asking each application host for its state. The adaptor can use a multicast request to find out which hosts are available or you can define a host list for it. Using this method, you avoid having to configure new hosts as you add them to your site.
- **From a static file:** An adaptor configuration file contains host and application information about your site; it includes information about every application instance you want to run. After the adaptor reads the file, it has all the information it needs to communicate with the application instances you want to run. Using this kind of configuration avoids multicast requests and host polling. However, when you add new hosts, you'll have to update the configuration file. For more information on the adaptor configuration file, see [“Using a Configuration File”](#) (page 44).

You configure your site using Monitor, a Web-browser-based tool.

## Types of Adaptors

---

There are two general types of HTTP adaptors, **CGI adaptors** and **API-based adaptors**. CGI adaptors are portable across many platforms. API-based adaptors are generally more efficient than CGI adaptors.

### CGI Adaptors

---

WebObjects Deployment includes a CGI adaptor, which is an executable file named `WebObjects`; in Windows 2000, it's named `WebObjects.exe`. The CGI adaptor resides in the HTTP server's `cgi-bin` or `scripts` directory. This adaptor works with any HTTP server that conforms to the CGI standard.

The major drawback of CGI adaptors is their performance. When the HTTP server receives a request from a Web browser, it creates a new process for the HTTP adaptor. When the adaptor is done processing the request, the process is terminated.

## HTTP Adaptors

The CGI adaptor is installed by default on all platforms, but it may not be the active one on your platform. See “[Default Adaptor Installation](#)” (page 33) for more information.

## API-Based Adaptors

---

API-based adaptors are founded on API specific to a particular HTTP server. They allow CGI-like tasks to run as part of the main server process, avoiding the creation and termination of a process for each request. [Table 4-1](#) lists the API-based adaptors included with WebObjects Deployment and the platforms on which they are supported.

---

**Table 4-1** API-based adaptors and supported platforms

<b>Adaptor</b>	<b>API</b>	<b>Supported platforms</b>
Apache	Apache’s module API	Mac OS X Server Solaris
ISAPI	Microsoft’s Internet Information Server API	Windows 2000
NSAPI	Netscape Server 3.5 API	Solaris Windows 2000

## State Discovery

---

Your site’s state is represented by

- a list of application hosts
- a list of running application instances on each host

## HTTP Adaptors

The HTTP adaptor captures your site's state at regular intervals, which you set when you configure the adaptor. You also define the method that the adaptor uses to gather state information by configuring the adaptor itself. For details, see ["Customizing the HTTP Adaptor"](#) (page 52).

The adaptor can obtain the state of your site using one of three methods:

- **Multicast request.** The adaptor sends a multicast request to find out what application hosts are available. After the host list is compiled, the adaptor polls each host to get its list of running application instances.
- **Host list.** This method requires that you configure the host list in the adaptor itself. As with the first method, the adaptor polls the hosts on the list for their lists of running application instances.
- **Configuration file.** The adaptor obtains the site's configuration by reading an XML (Extensible Markup Language) document.

The method that requires the least administration on your part is the multicast request. If an application host goes down, the adaptor automatically removes the application instances running on it from its list of active instances. When the host is brought back up, the adaptor adds the instances back to its list. You should use this method if your site has many application hosts. See ["Using a Multicast Request"](#) (page 41) for more information.

The second method, defining a host list for your adaptor, eliminates the multicast request. Use this method if you do not want the adaptor to send regular multicast requests out on your network or if you seldom add or remove application hosts from your site. This is the method that is active by default. However, the host list contains only one host, `localhost`. For details, see ["Using a Defined Host List"](#) (page 44).

In the third method, using a configuration file, the HTTP adaptor obtains your site's configuration by reading a file. This file can be static or it can be dynamically updated as you configure your site with Monitor. For details, see ["Using a Configuration File"](#) (page 44).

You can write the adaptor configuration file in one of two ways:

- **Manually.** The information in the configuration file is stored in an XML document. For details, see ["The HTTP Adaptor Configuration File"](#) (page 44).



## HTTP Adaptors

- **Using Monitor and wotaskd.** After configuring your site to your liking using Monitor, you can have a file created for you or you can copy-and-paste the information. See “[Creating the HTTP Adaptor Configuration File](#)” (page 47) for more information.

## Using a Multicast Request

---

When you configure an adaptor to obtain your site’s state using a multicast discovery request, the adaptor obtains the list of active application hosts by broadcasting a message to which each computer configured as a WebObjects application host responds. After the adaptor compiles the list of available hosts, it polls each one to obtain its state (the list of running application instances).

There are drawbacks to using the multicast method:

- It increases network traffic. By default, the HTTP adaptor sends a multicast request every 100 seconds.
- A host may become unavailable between discovery requests if the multicast request or a wotaskd process’s response is lost (multicast is an inherently unreliable protocol).
- Normally, multicast broadcasts are limited to a subnet. However, you can configure your routers to pass on the multicast request to other subnets if you wish.

By default, wotaskd does not respond to multicast requests. To be able to use the multicast request method, you must configure wotaskd processes in your application hosts to respond to multicast requests.

## Multicast Request

---

To discover available hosts, the adaptor sends a host-discovery request on the multicast channel (a nonrouting IP address and a port number), which is set to IP address 239.128.14.2 and port 1085 by default. The frequency of each multicast request is ten times as long as the adaptor’s configuration refresh interval. For details on how to change the multicast channel, see “[Setting the Multicast Address and Port](#)” (page 52), “[WOPort](#)” (page 119), “[WOMulticastAddress](#)” (page 120), and

## HTTP Adaptors

“[WOREspondsToMulticastQuery](#)” (page 121). When a wotaskd process starts, it creates a UDP (User Datagram Protocol) socket that listens to the multicast channel through which it receives multicast requests.

**Note:** If you change the address and port that adaptors use to send multicast requests, you must also change the address and port that wotaskd processes use to receive multicast requests.

When each wotaskd process receives the multicast request, it replies with its URL, such as `http://host1.site.com:1085`. The adaptor in turn compiles a list of these URLs.

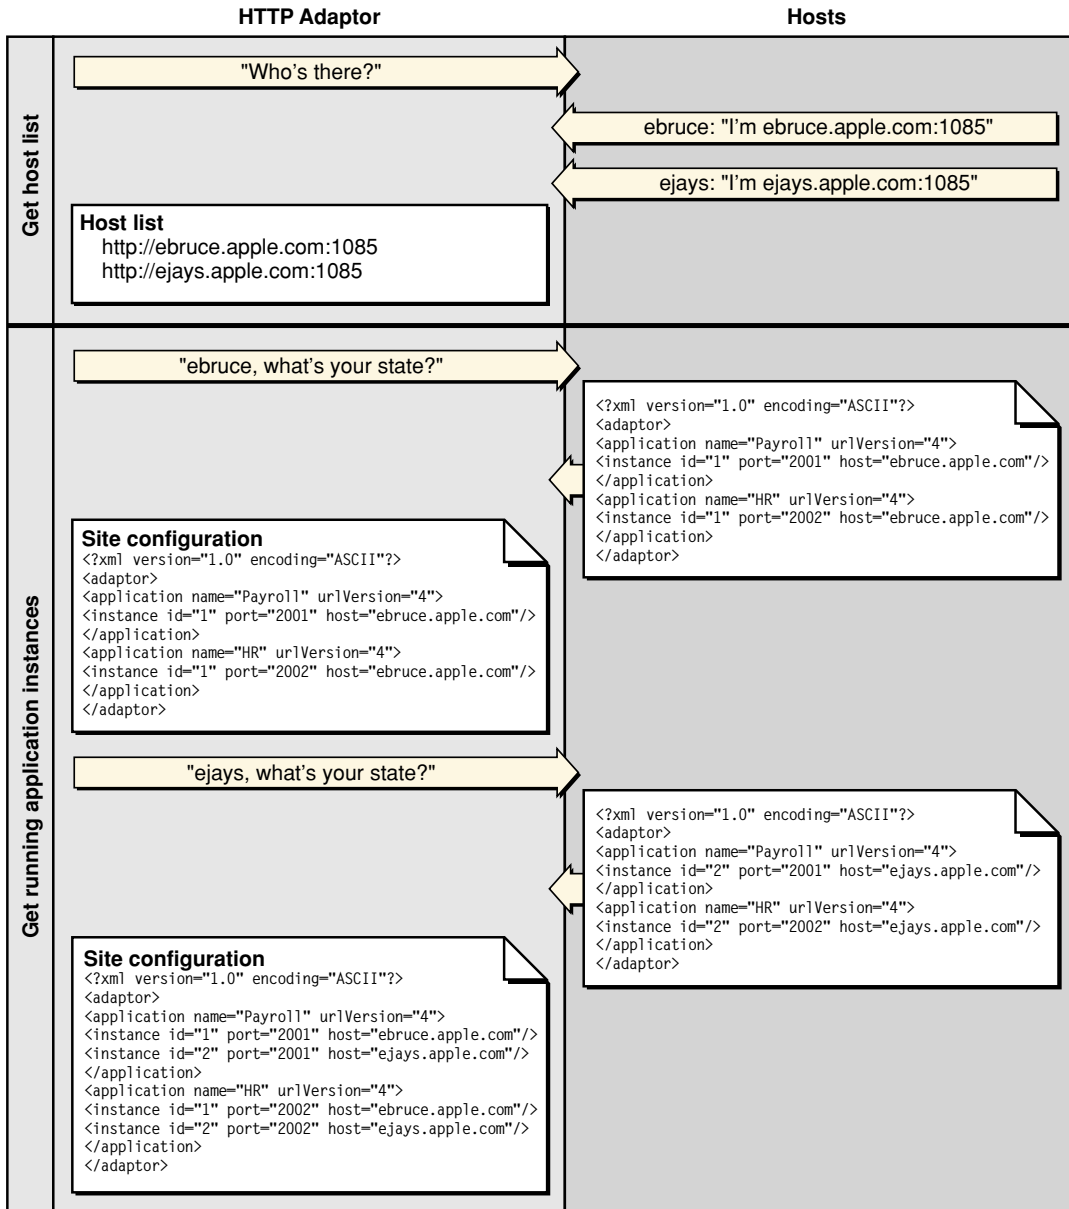
Sending a multicast request on an entire subnet is an expensive process. If your available hosts never change, consider using a defined host list instead.

## Host Polling

---

After the HTTP adaptor constructs the host list, it polls each application host on the list for information on the active application instances running on it. Each wotaskd process, in turn, sends its state information using the format in [Listing 4-2](#) (page 45). Host polling to obtain information on active instances occurs at the interval indicated in the configuration refresh interval setting for the HTTP adaptor. [Figure 4-4](#) illustrates the process used to determine the configuration of the site in [Figure 4-3](#) (page 37).

Figure 4-4 Dynamic site configuration using multicast request and polling



## Using a Defined Host List

---

This method is similar to the one described in [“Using a Multicast Request”](#) (page 41). The only difference is that the HTTP adaptor skips the first part, the multicast request. The host polling process occurs at the interval set in the adaptor’s configuration refresh interval setting.

You must explicitly define a host list for each adaptor. See [“Setting the Host List”](#) (page 53) for details on defining the host list for each of the adaptors provided.

## Using a Configuration File

---

Using an HTTP adaptor configuration file is useful when you want to have a static site configuration (one in which application instances are not stopped after they are started) or if you want to use Monitor to configure your site and have the adaptor read your configuration changes immediately. (The adaptor reads the configuration file every 10 seconds to determine which application instances are active.)

This method also provides a way of having more than one configuration of your site available. You can switch among different configurations by placing the appropriate configuration file in the configuration directory.

[“The HTTP Adaptor Configuration File”](#) (page 44) explains how the file is structured and lists the properties that it defines. For instructions on creating the configuration file and configuring the HTTP adaptor to use it, see [“Creating the HTTP Adaptor Configuration File”](#) (page 47).

## The HTTP Adaptor Configuration File

---

You can set up the HTTP adaptor to get your site’s configuration by reading an HTTP adaptor configuration file (called `W0Config.xml` by default) in the configuration directory (`/Library/WebObjects/Configuration` by default). You should have only one adaptor configuration file per HTTP server so that it can perform load balancing effectively. (See [“Load Balancing”](#) (page 96) for details.) In addition, in a site with multiple HTTP servers, if two servers share the configuration file, instead of deploying two sites you would be deploying the same site twice. [Listing 4-1](#) shows a configuration file that defines a site with two application hosts (`ebruce.apple.com` and `ejays.apple.com`), each running two application instances, one of the Payroll application and the other of the HR application.

**Listing 4-1** A WebObjects adaptor configuration file

```
<?xml version="1.0" encoding="ASCII"?>
<adaptor>
  <application name="Payroll" urlVersion="4">
    <instance id="1" port="2002" host="ebruce.apple.com"/>
    <instance id="2" port="2001" host="ejays.apple.com"/>
  </application>
  <application name="HR" urlVersion="4">
    <instance id="1" port="2001" host="ebruce.apple.com"/>
    <instance id="2" port="2002" host="ejays.apple.com"/>
  </application>
</adaptor>
```

The HTTP adaptor configuration file provides the HTTP adaptor with information about your site's registered application instances. The structure of the configuration file is provided in [Listing 4-2](#) (you can also view it by opening the `woadaptor.dtd` file, located in the `/Developer/Examples/WebObjects/Source/Adaptors` directory). For information on the properties defined in the configuration file, consult [Table 4-2](#) (page 46).

**Listing 4-2** Structure of the HTTP adaptor configuration file

```
<?xml version="1.0" encoding="ASCII"?>

<!DOCTYPE WebObjectsAdaptorConfiguration SYSTEM "woadaptor.dtd">

<adaptor>
  <application name=STRING
    retries=NUMBER
    scheduler=["RANDOM"|"ROUNDROBIN"|"LOADAVERAGE"]
    dormant=NUMBER
    protocol="http"
    redir=URL
    poolSize=NUMBER
    urlVersion=["3"|"4"]
    additionalArgs="unspecified"
  >
  <instance id=NUMBER port=NUMBER host=STRING
```

HTTP Adaptors

```

        sendTimeout=NUMBER
        recvTimeout=NUMBER
        cnctTimeout=NUMBER
        sendBufSize=NUMBER
        recvBufSize=NUMBER
        additionalArgs="unspecified"
    >
</instance>
</application>
</adaptor>

```

**Table 4-2** The properties of the HTTP adaptor configuration file

<b>Property</b>	
name	This property is used by the adaptor to implement load balancing. The adaptor can load-balance only between instances with the same application name. The property can be used to create groups of instances, even when the instances share the same executable file. This argument is set automatically for instances started by wotaskd.
retries	The number of times a request is retried (trying several instances) if a communications failure occurs before an error page is returned to the HTTP server.
scheduler	The load-balancing scheme used by the adaptor for instances of the application. The options provided by WebObjects are Round Robin, Random, and Load Average. You can also use a custom load balancer; see “Load Balancing and Adaptor Settings” (page 82) for details.
dormant	The number of times the adaptor skips an instance of the application before trying again.
redir	The URL that the user is redirected to when an instance fails to respond to a direct request.
poolSize	The maximum number of simultaneous connections the adaptor should keep open for each configured instance.
urlVersion	The WebObjects version to use for URL parsing and formatting. All WebObjects 4, 4.5, and 5 applications use version 4 URLs by default.

**Table 4-2** The properties of the HTTP adaptor configuration file (continued)

<b>Property</b>	
<code>additionalArgs</code>	Additional information to send to the instance when it's started.
<code>id</code>	The instance's identification number. Must be unique for the load-balancing process to operate correctly.
<code>port</code>	The port on which the instance runs.
<code>host</code>	Specifies the network interface that an instance binds to. This argument should only be used on hosts with multiple network interfaces (IP addresses).
<code>sendTimeout</code>	The length of time, in seconds, that the adaptor attempts to send data to an instance of the application before giving up.
<code>recvTimeout</code>	The length of time, in seconds, that the adaptor waits for a response from an instance of the application before giving up.
<code>cnctTimeout</code>	The length of time, in seconds, before the adaptor gives up connecting to an instance.
<code>sendBufSize</code>	The size, in bytes, of the TCP/IP socket send buffer that's used for adaptor-to-instance communication.
<code>recvBufSize</code>	The size, in bytes, of the TCP/IP socket receive buffer that's used for adaptor-to-instance communication.

## Creating the HTTP Adaptor Configuration File

You can define your site's configuration by writing the HTTP adaptor configuration file by hand. However, Monitor provides you with an easy-to-use interface that facilitates that task.

"[Deployment Tasks](#)" (page 67) shows you how to configure your site using Monitor. When you are satisfied with the configuration, you can save the settings into a configuration file by copying and pasting or by telling `wotaskd` to write the file.

To use the copy-and-paste method, follow these steps:

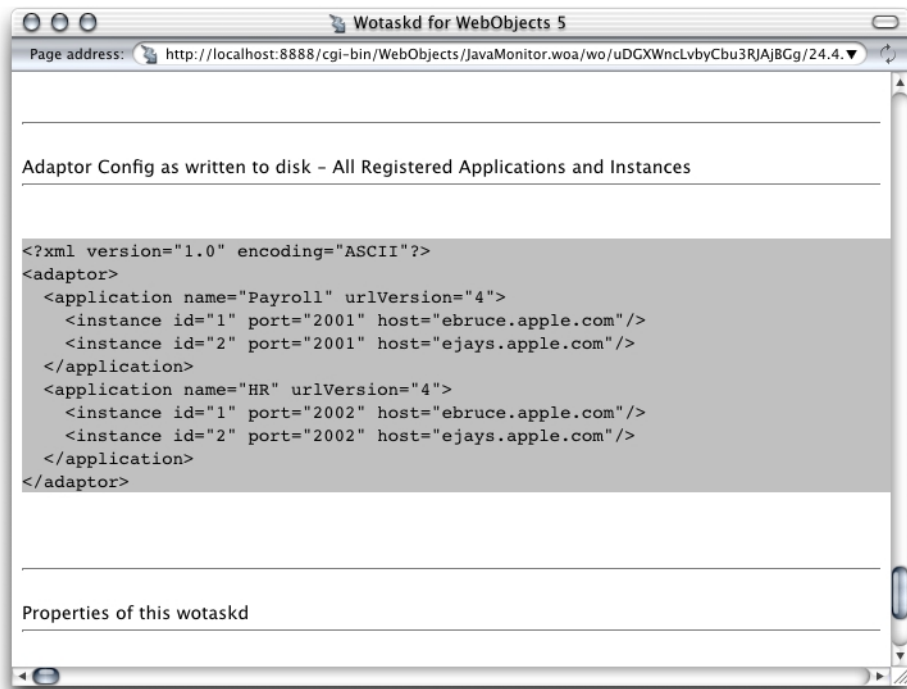
1. In Monitor, display the Hosts page.
2. Click YES for any host.

## HTTP Adaptors

The host configuration page is displayed in a new Web browser window.

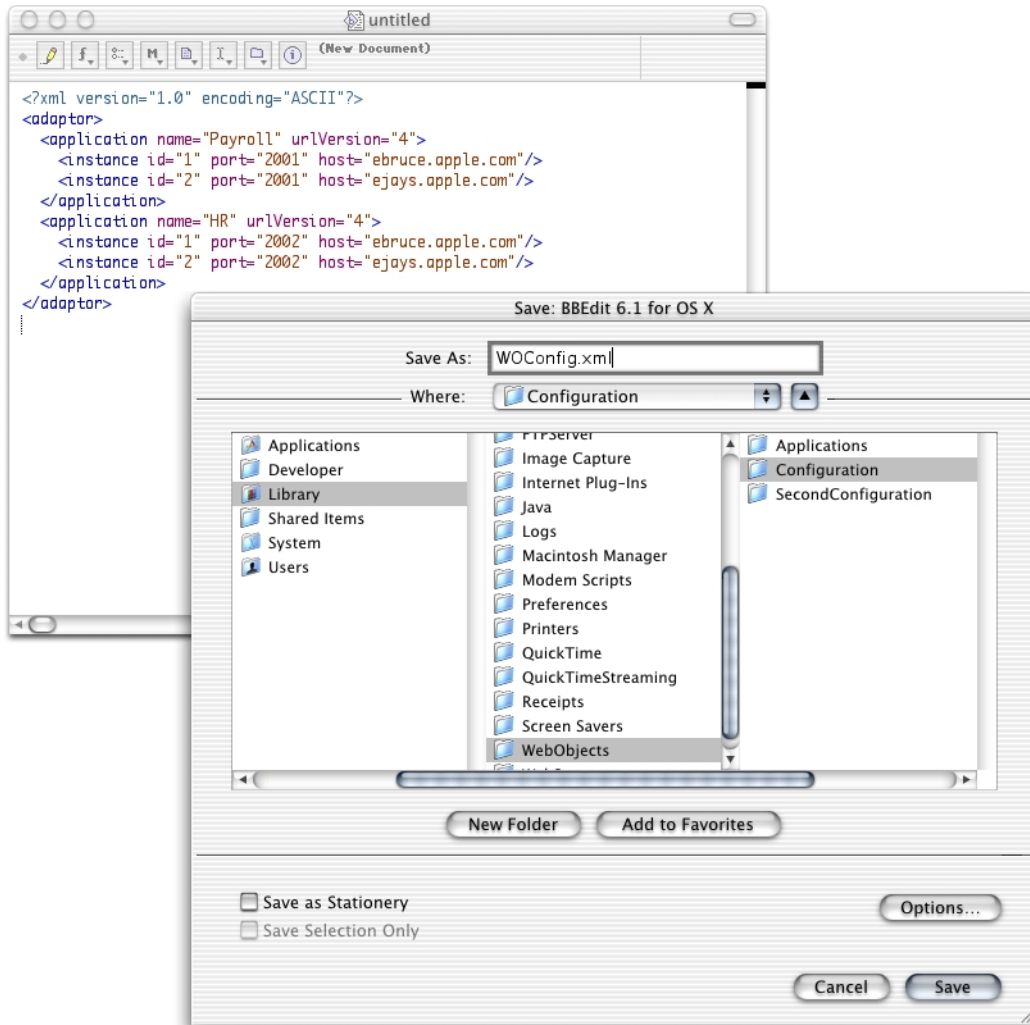
3. Copy the contents of the section “Adaptor Config as written to disk—All Registered Applications and Instances,” as shown in [Figure 4-5](#).

**Figure 4-5** Copying the information that makes up the HTTP adaptor configuration file



4. Using a text editor, create a new file and paste the contents of the clipboard into it.
5. Save the file as `W0Config.xml` (or any other name you choose) in the configuration directory.



**Figure 4-6** Creating and saving the HTTP adaptor configuration file

If instead of copying and pasting you want `wotaskd` to create the configuration file for you, you can start a `wotaskd` process specifically to create the file or you can tell `wotaskd` to continually maintain the configuration file.

## HTTP Adaptors

To start a wotaskd process specifically to create the file, you must first stop the process that corresponds to the site you configured if it's already running on the HTTP server computer.

To start a wotaskd process that writes the configuration file to the default location, execute the following two commands using your command shell editor:

```
cd /System/Library/WebObjects/JavaApplications/wotaskd.woa
./wotaskd -WOPort <port> -WOSavesAdaptorConfiguration true
```

To specify a different location for the HTTP adaptor configuration file, follow the instructions in “[WODeploymentConfigurationDirectory](#)” (page 120). If you want to give the file a different name, “[Setting the Name of the HTTP Adaptor Configuration File](#)” (page 54) shows you how.

To tell wotaskd to maintain the configuration file on a permanent basis, add the following to the WOServices script line that starts the wotaskd process:

```
-WOSavesAdaptorConfiguration true
```

When you restart your HTTP server, the HTTP adaptor configuration file is updated every time you make a change to your site's configuration through Monitor. The changes are picked up by the HTTP adaptor the next time it reads the configuration file.

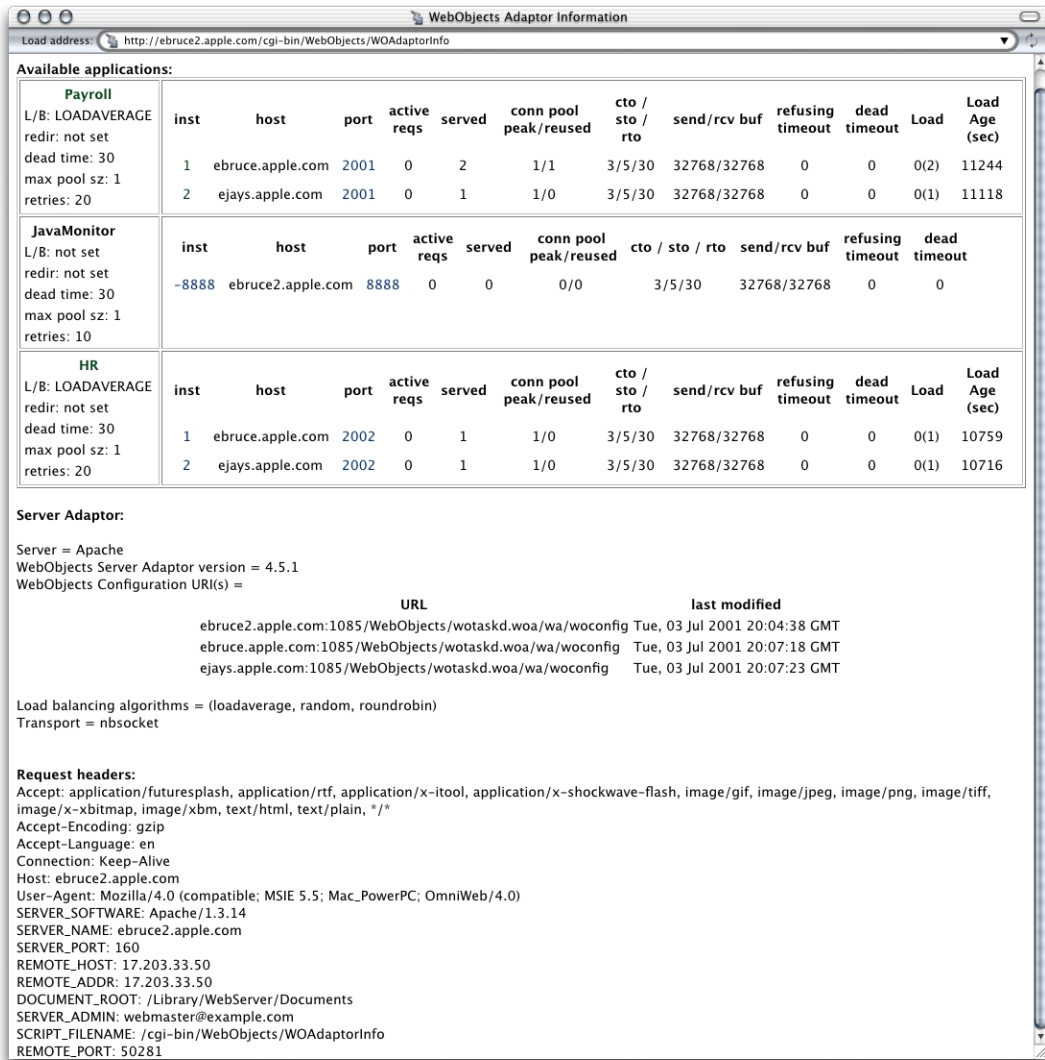
To configure the HTTP adaptor to read the configuration file instead of using a multicast request or a host list, follow the instructions in “[Setting the Name of the HTTP Adaptor Configuration File](#)” (page 54).

## The WebObjects Adaptor Information Page

---

The WebObjects Adaptor Information page displays information about an HTTP adaptor. Access to this page is disabled by default so you must modify the adaptor configuration file to allow access. See “[Setting Access to the WebObjects Adaptor Information Page](#)” (page 54) for details. [Figure 4-7](#) shows an example of the WebObjects Adaptor Information page.

Figure 4-7 The WebObjects Adaptor Information page



## Customizing the HTTP Adaptor

---

For the most part, you shouldn't need to modify the default values of settings in the configuration file. However, if you want to change the way the HTTP adaptor obtains your site's state information, for example, you'll need to perform some of the procedures explained here.

These are the tasks explained in this section:

- “Setting the Multicast Address and Port” (page 52)
- “Setting the Host List” (page 53)
- “Setting the Name of the HTTP Adaptor Configuration File” (page 54)
- “Setting Access to the WebObjects Adaptor Information Page” (page 54)
- “Setting an Alias for cgi-bin in the WebObjects URL” (page 55)
- “Setting the Document Root Path of the HTTP server” (page 56)

### Setting the Multicast Address and Port

---

The following list explains how to set the multicast address, port, and configuration refresh interval (in seconds) in the supported HTTP adaptors. The default values for each of these properties are 239.128.14.2, 1085, and 10 respectively. The adaptor uses the configuration interval to determine the amount of time that passes between state discoveries on your site. The host-discovery process occurs 10 times less frequently than the time indicated by the configuration refresh interval. With the configuration refresh interval set to 10, the discovery process occurs every 100 seconds.

- **Apache:** Set the value of the `WebObjectsConfig` variable in the `apache.conf` file to the desired values, using the format shown below:

```
WebObjectsConfig webobjects://<address>:<port> <configuration_interval>
```

- **ISAPI:** Add two keys to the Registry, `CONF_URL` and `CONF_INTERVAL`, choose `REG_SZ` as their data type, and set their values as follows:

## HTTP Adaptors

```
\\HKEY_LOCAL_MACHINE\\SOFTWARE\\Apple\\WebObjects\\Configuration\\CONF_UR
L webobjects://<address>:<port>
```

```
\\HKEY_LOCAL_MACHINE\\SOFTWARE\\Apple\\WebObjects\\Configuration\\CONF_I
NTERVAL <configuration_interval>
```

- **NSAPI:** Add the following line to the `magnus.conf` file:

```
Init fn="WebObjects_init" root="/opt/ns-home/docs" config="webobjects://
<address>:<port> confinterval="<configuration_interval>"
```

- **CGI:** Set the `WO_CONFIG_URL` environment variable to `webobjects://<address>:<port>`. Make sure your HTTP server is configured to pass the variable to the adaptor (consult your HTTP server's documentation for instructions).

## Setting the Host List

---

The following list explains how to set a host list for a site with two hosts, `host1` and `host2`, in the supported adaptors with a configuration interval of 10 (the configuration interval cannot be set in the CGI adaptor).

- **Apache:** Set the `WebObjectsConfig` variable in the `apache.conf` file to the desired list of hosts. By default it's set to `http://localhost:1085 10` (the 10 is the configuration refresh interval). Separate each host with a comma, as shown in the following example:

```
WebObjectsConfig http://host1:1085,http://host2:1085 10
```

- **ISAPI:** Add two keys to the Registry, `CONF_URL` and `CONF_INTERVAL`, choose `REG_SZ` as their data type, and set their values as follows:

```
\\HKEY_LOCAL_MACHINE\\SOFTWARE\\Apple\\WebObjects\\Configuration\\CONF_UR
L http://host1:1085,http://host2:1085
```

```
\\HKEY_LOCAL_MACHINE\\SOFTWARE\\Apple\\WebObjects\\Configuration\\CONF_I
NTERVAL 10
```

- **NSAPI:** Set the `WebObjects_init` function's arguments in the `magnus.conf` file as follows:

```
Init fn="WebObjects_init" root="/opt/ns-home/docs" config="http://
host1:1085,http://host2:1085" confinterval="10"
```

## HTTP Adaptors

- **CGI:** Set the `WO_CONFIG_URL` environment variable to `http://host1:1085,http://host:1085`. Make sure the HTTP server is configured to pass the variable to the adaptor (consult your HTTP server's documentation for instructions).

## Setting the Name of the HTTP Adaptor Configuration File

---

- **Apache:** Set the value of `WebObjectsConfig` variable in the `apache.conf` file to the path of the adaptor configuration file.

```
WebObjectsConfig file://<path-to-an-xml-config-file> 10
```

- **ISAPI:** Add the `CONF_URL` key to the Registry, choose `REG_SZ` as its data type, and set the adaptor configuration file path as its value, as the following example shows:

```
\\HKEY_LOCAL_MACHINE\\SOFTWARE\\Apple\\WebObjects\\Configuration\\CONF_URL file://<path-to-an-xml-config-file>
```

- **NSAPI:** Set the arguments of the `WebObjects_init` function in the `magnus.conf` file as follows:

```
Init fn="WebObjects_init" root="/opt/ns-home/docs" config="file://<path-to-an-xml-config-file>"
```

- **CGI:** Set the `WO_CONFIG_URL` environment variable to `file://<path-to-an-xml-config-file>`. Make sure the HTTP server is configured to pass the variable to the adaptor (consult your HTTP server's documentation for instructions).

## Setting Access to the WebObjects Adaptor Information Page

---

You can provide access to the WebObjects adaptor information page to a specific user or to everyone. To provide access to a single user, you set the values of the `username` and `password` properties. To provide public access, set the `username` attribute to `public`. The following list explains how to provide access to the information page to a user named Joe in the supported adaptors. For the changes to take effect, you need to restart the HTTP server.

## HTTP Adaptors

- **Apache:** Add the following lines to the `apache.conf` file, located in the `/System/Library/WebObjects/Adaptor/Apache` directory:

```
WebObjectsAdminUsername joe
WebObjectsAdminPassword secret
```

- **ISAPI:** Add two keys to the Registry, `WOUSername` and `WOPASSWORD`, choose `REG_SZ` as their data type, and set their values as follows:

```
\\HKEY_LOCAL_MACHINE\\SOFTWARE\\Apple\\WebObjects\\Configuration\\WOUSER
NAME joe
```

```
\\HKEY_LOCAL_MACHINE\\SOFTWARE\\Apple\\WebObjects\\Configuration\\WOPASS
WORD secret
```

- **NSAPI:** Add the following line to `magnus.conf`:

```
Init fn="WebObjects_init" root="/opt/ns-home/docs" config="webobjects://
<address>:<port>" username="joe" password="secret"
```

- **CGI:** Set the `WO_ADAPTOR_INFO_USERNAME` and `WO_ADAPTOR_INFO_PASSWORD` environment variables to the appropriate values. Make sure the HTTP server is configured to pass the variables to the adaptor (consult your HTTP server's documentation for instructions).

## Setting an Alias for cgi-bin in the WebObjects URL

---

The following list explains how to change the `cgi-bin` part of the URL used to connect to an application instance to `Store` in the Apache, NSAPI, and CGI adaptors:

- **Apache:** In the `apache.conf` file, change the line

```
WebObjectsAlias /cgi-bin/WebObjects
```

to

```
WebObjectsAlias /Store/WebObjects
```

- **NSAPI:** In the `magnus.conf` file, change the lines

```
NameTrans from="/cgi-bin/WebObjects" fn="WebObjectsNameTrans"
name="webobjects"
NameTrans from="/cgi-bin" fn="pfx2dir" dir="/opt/ns-home/cgi-bin"
name="cgi"
```

### HTTP Adaptors

to

```
NameTrans from="/Store/WebObjects" fn="WebObjectsNameTrans"
name="webobjects"
NameTrans from="/Store" fn="pfx2dir" dir="/opt/ns-home/cgi-bin"
name="cgi"
```

## Setting the Document Root Path of the HTTP server

---

- **Apache:** In the `apache.conf` file, change the line

```
WebObjectsDocumentRoot /Library/WebServer/Documents
```

to

```
WebObjectsDocumentRoot <document-root-path>
```

- **ISAPI:** Add the following Registry entry:

```
\\HKEY_LOCAL_MACHINE\\SOFTWARE\\Apple\\WebObjects\\Configuration\\DOCUME
NT_ROOT <document-root-path>
```

- **NSAPI:** In the `magnus.conf` file, change the value of the `root` variable in every line that defines it to the desired path. For example, the line

```
Init fn="WebObjects_init" root="/opt/ns-home/docs" config="http://
localhost:1085"
```

needs to be changed to

```
Init fn="WebObjects_init" root="<document-root-path>" config="http://
localhost:1085"
```

- **CGI:** Set the value of the `CGI_DOCUMENT_ROOT` environment variable to the desired path. Make sure that your HTTP server is configured to pass the variable to the adaptor (consult your HTTP server's documentation for instructions).



# Managing Application Instances

---

This chapter provides a detailed description of Monitor and wotaskd, two tools you use to manage the application instances running on your site, and it explains how instances communicate with wotaskd:

- “[Configuration Files](#)” (page 57) provides an overview of the configuration files that can be used in your site.
- “[Lifebeats](#)” (page 60) introduces lifebeats, the mechanism used by wotaskd processes to determine the status of the application instances that they manage.
- “[wotaskd Processes](#)” (page 61) explains how wotaskd processes ensure that the application instances they manage are always running. It also describes how to restrict access to wotaskd using Monitor.
- “[Starting WebObjects Services](#)” (page 62) shows you how to configure a computer to start wotaskd and Monitor during its startup process.

## Configuration Files

---

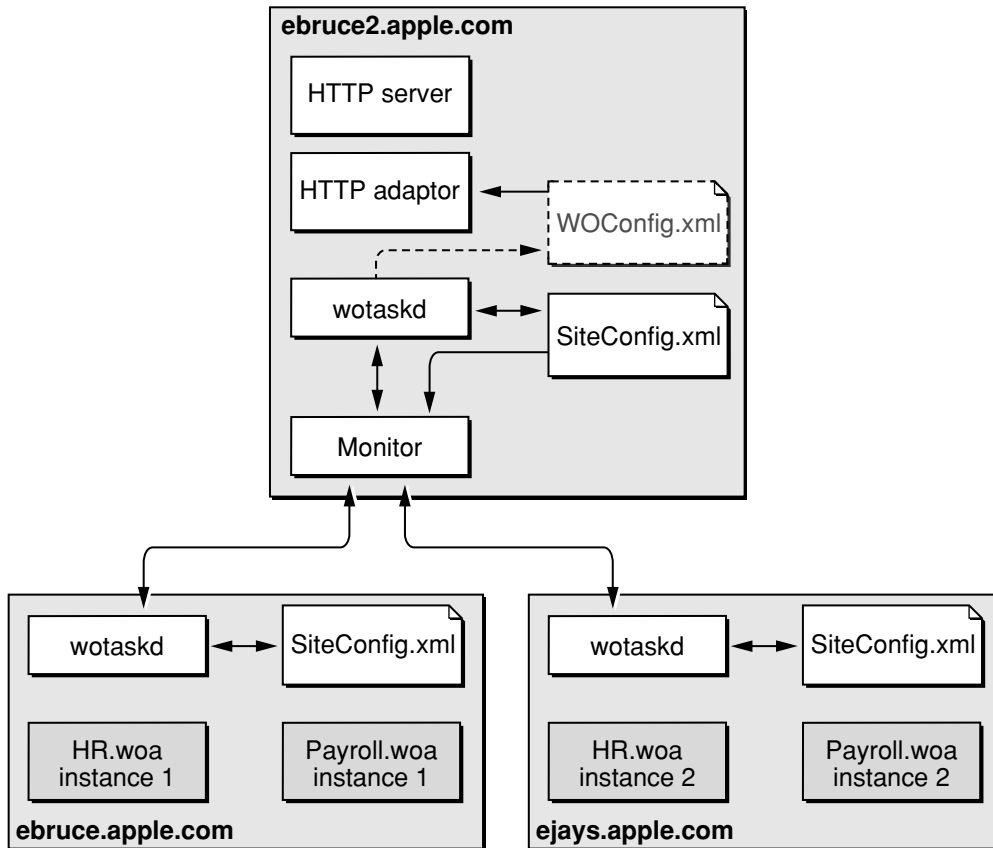
You use Monitor to configure your site. With it you configure hosts, applications, and application instances. You also define schedules for restarting instances and choose the algorithm used to balance the user load among the instances of an application. For details on the tasks that you can perform using Monitor, see “[Deployment Tasks](#)” (page 67).

### Managing Application Instances

The `SiteConfig.xml` file, which is maintained on each host's deployment-configuration directory by `wotaskd`, stores the configuration choices you make in Monitor. You should not modify the contents of this file directly. The user under which `wotaskd` runs must have read and write privileges to `SiteConfig.xml` and the user under which Monitor runs must have read privileges.

You can create another configuration file (the HTTP adaptor configuration file), which is called `WOConfig.xml` by default. This is the file the HTTP adaptor uses to obtain your site's configuration when you choose the configuration file method for the adaptor.

[Figure 5-1](#) shows how the configuration files are distributed in a deployment with one HTTP server and two application hosts. For information on how to generate the HTTP adaptor configuration file, see [“Creating the HTTP Adaptor Configuration File”](#) (page 47).

**Figure 5-1** WebObjects configuration-file distribution

Monitor reads the `SiteConfig.xml` file when it starts up but never writes to it. `wotaskd` writes the `SiteConfig.xml` file when instructed by Monitor to do so—such as when existing entities are configured, added, or deleted. Furthermore, Monitor tells `wotaskd` to update the `SiteConfig.xml` file only when you use it to change an aspect of your site's configuration.

## Lifebeats

---

After adding an application to your site, you need to add instances of it as well. This allows the application's users connect to it. One of the main goals of WebObjects Deployment is to provide you with tools that help you deploy a resilient site. To accomplish that, wotaskd restarts application instances that become unresponsive.

A **lifebeat** is a status message that an application instance sends to a wotaskd process to keep it informed of the instance's status. There are four kinds of lifebeats:

- has started
- is alive
- will stop
- will crash

An application instance is configured by default to send lifebeats to a wotaskd process through TCP (Transmission Control Protocol) sockets. (A **socket** is a mechanism through which two processes communicate.) Instances can also send lifebeats using UDP (User Datagram Protocol) sockets. UDP sockets use fewer resources than TCP sockets.

The lifebeat mechanism is how the state of your site is constantly updated. Lifebeats are sent from a separate execution thread in a WebObjects application and do not interfere with, nor are they affected by, normal request processing.

By default, application instances try to send a lifebeat every 30 seconds. If a failure occurs (there is no wotaskd process listening on the port indicated by `WOLifebeatDesinationPort`), the instance does the following:

1. Sends up to 10 TCP lifebeats until a response is received. After the tenth unanswered lifebeat, the instance goes to stage 2.
2. Sends UDP lifebeats (low-resource-lifebeat mode) until a response is received. If the instance cannot create a UDP socket, it goes back to stage 1. When the instance receives an acknowledgement, it resumes sending TCP lifebeats.

## Managing Application Instances

The two-stage mechanism allows application instances to go into low-resource lifebeat mode if wotaskd isn't running when the instance starts.

## wotaskd Processes

---

WebObjects Deployment uses wotaskd to manage the application instances running on your application hosts. Its main task is to start up instances when hosts are restarted. To accomplish this, wotaskd itself has to be restarted when the host starts up. This is done by configuring wotaskd as a service started when the computer boots. By default, a wotaskd process running on port 1085 is configured as a service on all supported platforms. The implementation of this feature is platform-specific. See [“Starting WebObjects Services”](#) (page 62) for details.

You need to run a wotaskd process on every computer that you want to use as an application host. These processes constantly receive lifebeats from the application instances they manage. Lifebeats communicate the instance's state and allow wotaskd to determine the state of the instances it oversees. A wotaskd process assumes that an application instance is dead if it does not receive a lifebeat within a certain period. For details, see [“WOAssumeApplicationIsDeadMultiplier”](#) (page 120).

When you restrict access to Monitor with a password, wotaskd processes running on hosts configured in Monitor are also protected: They do not respond to `http://<hostname>:<wotaskd-port>`, as described in [“Confirming That wotaskd Is Active”](#) (page 31).

To access the state information of a particular wotaskd process, you use Monitor's Host page. See [“Viewing a Host's Configuration”](#) (page 71) for more information.

## Starting WebObjects Services

---

You use Monitor and wotaskd to build and troubleshoot your site. WebObjects provides scripts that allow you to control some of the behavior of Monitor and wotaskd. Specifically, you can determine whether they are started automatically when a computer starts up. That way, you'll have one less item to worry about if a computer goes down. On the other hand, if you prefer a more hands-on approach to site management, you can start and stop WebObjects services manually.

### Starting WebObjects Services Automatically

---

The default installation of WebObjects Deployment in Mac OS X Server adds a startup script that automatically starts a wotaskd process during system startup. In addition, if the wotaskd process dies, it's automatically restarted.

You can configure a computer to keep wotaskd and Monitor processes active at all times. To accomplish that, you have to edit the `WebObjects` script file. In Mac OS X Server, it's located in the `/System/Library/StartupItems/WebObjects` directory; in Solaris, you'll find it in the `/etc/init.d` directory. For example, to always have a Monitor process running on port 8888 in your site-administration computer, you add the following line after the line that starts wotaskd in the `WebObjects` script file:

```
"$WOSERVICE" -appPath /System/Library/WebObjects/JavaApplications/  
JavaMonitor.woa/JavaMonitor -WOPort 8888 -WOAutoOpenInBrowser false>>/var/  
log/webobjects.log 2>&1 &
```

For more on the command-line arguments available, see [“Command-Line Arguments”](#) (page 116).

### Starting Monitor Manually

---

To start Monitor, enter the following commands in your command shell editor:

```
cd ($NEXT_ROOT)/System/Library/WebObjects/JavaApplications/JavaMonitor.woa  
./JavaMonitor
```

## C H A P T E R 5

### Managing Application Instances

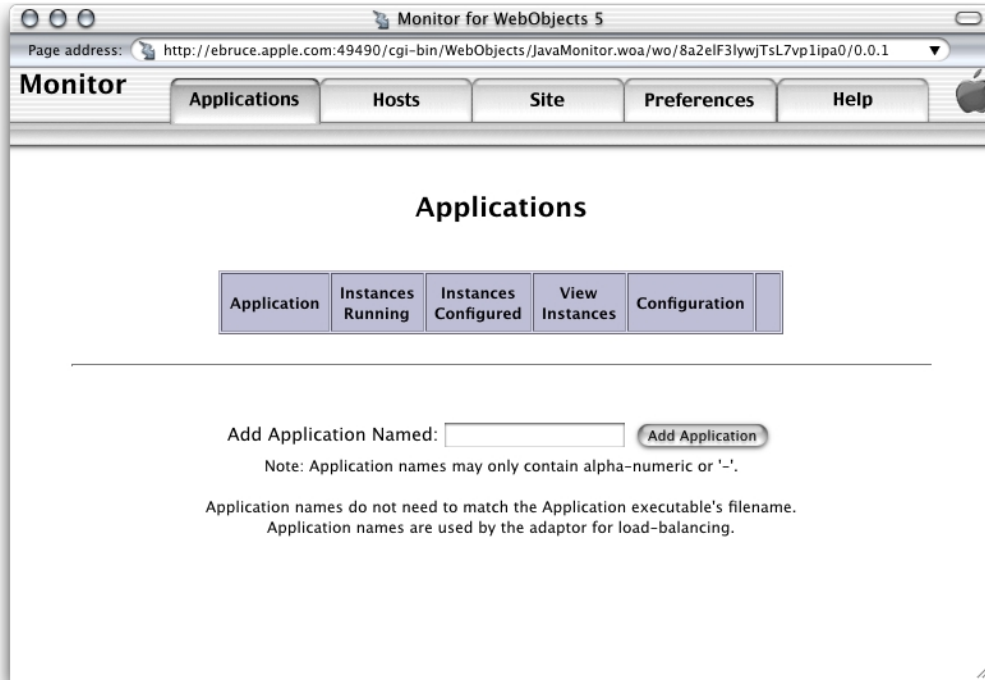
You should see output similar to that in [Listing 5-1](#).

---

#### **Listing 5-1** Starting Monitor

```
Reading MacOSClassPath.txt ...
Launching JavaMonitor.woa ...
...
Creating LifebeatThread now with: JavaMonitor 49490 ebruce.apple.com/
17.203.33.19 1085 30000
Opening application's URL in browser:
http://ebruce.apple.com:49490/cgi-bin/WebObjects/JavaMonitor
Waiting for requests...
```

A page like the one in [Figure 5-2](#) should display in your Web browser. If your browser is not launched automatically, you can copy the URL from your shell and paste it into your browser's address field.

**Figure 5-2** Monitor—empty Applications page

## Controlling WebObjects Services in Solaris

The `WOServices` script file, located in the `$NEXT_ROOT/Library/WebObjects/Executables` directory, provides you with four options to manage WebObjects services (such as `wotaskd` and `Monitor`) from the command line: `start`, `stop`, `enable`, and `disable`.

- `WOServices start`  
Starts WebObjects services.
- `WOServices stop`  
Stops WebObjects services.
- `WOServices enable [-altJVMPATH path]`



## Managing Application Instances

Starts WebObjects services every time the computer is restarted. The `altJVMPATH` argument specifies a Java executable different from the one present during installation. It is used to set the `PATH` variable used when starting `wotaskd`. If you use the `altJVMPATH` argument, make sure to specify a full path.

■ `WOServices disable`

Disables `wotaskd` processes from starting automatically during system startup.

## Windows 2000

---

During installation on Windows 2000, `wotaskd` is configured to start automatically at boot time (in the Services control panel, it's listed as Apple WebObjects Task Daemon). If it doesn't start, check the Services control panel to ensure that `wotaskd`'s startup mode is set to `Automatic`.

Monitor is also configured as a service (listed as Apple WebObjects Monitor in the Services control panel) and you can configure it to start automatically during the boot process by changing its startup mode to `Automatic`. Note that although this causes Monitor to start automatically, you have to manually start your Web browser and connect to the Monitor process manually (or by putting a shortcut to your Web browser in your Startup program group). The URL for Monitor can be verified by checking the Windows 2000 Event Viewer (choose Start > Programs > Administrative Tools > Event Viewer) and is similar to the following:

```
http://localhost:1027/cgi-bin/WebObjects.exe/JavaMonitor
```

If you don't have Monitor configured to start automatically, you can launch it by choosing Start > Programs > WebObjects > Monitor. In that case, it runs on port 56789.

## C H A P T E R 5

### Managing Application Instances

# Deployment Tasks

---

Now that you are acquainted with the deployment tools of WebObjects, you are ready to put that knowledge to work by deploying your applications.

This chapter explains how Monitor allows you to perform most of the tasks required to maintain a WebObjects application site with point-and-click ease. The sections below guide you through the process of adding and configuring hosts, applications, and instances. Also explained is the load-balancing mechanism used by the HTTP adaptor, which performs load balancing among the instances of an application (which can be distributed across more than one host) and how to set up email notifications so that you and your colleagues are notified when problems arise.

This chapter addresses the following subjects:

- *“Setting Up Hosts”* (page 68) explains how you add and configure hosts for application deployment. Note that this is different from configuring hosts in the HTTP adaptor.
- *“Installing Applications”* (page 74) shows you where to place application files and HTTP server resources on an application host before deploying an application.
- *“Setting Up Applications”* (page 75) details the ways you can customize a deployment. These include choosing a load-balancing algorithm and scheduling instances to restart at regular intervals.
- *“Configuring Sites”* (page 91) describes the site-wide properties available. These include configuring Monitor to use your SMTP (Simple Mail Transfer Protocol) server to send email notifications.
- *“Setting Monitor Preferences”* (page 94) shows you the Monitor-specific preferences you can use to tailor the tool’s behavior.

## Deployment Tasks

- “Load Balancing” (page 96) explains how load balancing distributes user load among the running instances of an application in your site.
- “Deploying Multiple Sites” (page 97) shows how you can maintain multiple sites on one set of computers.

## Setting Up Hosts

---

An application host is a computer that runs application instances. For Monitor to be able to identify a host, it has to be running a wotaskd process. See “Choosing What to Install” (page 29) and “Confirming That wotaskd Is Active” (page 31).

### Adding a Host

---

Before you can deploy applications, you need to tell Monitor which hosts you want to use for deployment. (See “Load Balancing” (page 96) for additional information regarding load balancing and hosts added in Monitor.) *Figure 6-1* shows Monitor’s Hosts page, which you use to add and configure hosts.

Figure 6-1 The Hosts page



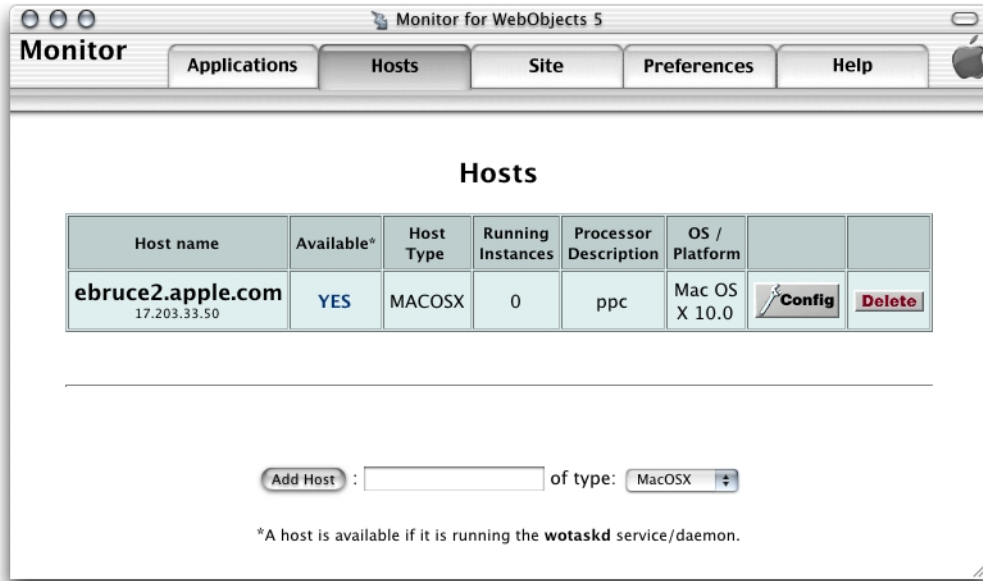
1. In Monitor, click the Hosts tab.
2. Enter the name or IP address of the host you want to add in the host text input field.

The computer must be running a wotaskd process in the port that Monitor sends its lifebeats to.

Avoid using a **loopback** address (a connection that does not go over the network), such as `localhost` or `127.0.0.1`. If you do, it must be the only application host in your site.

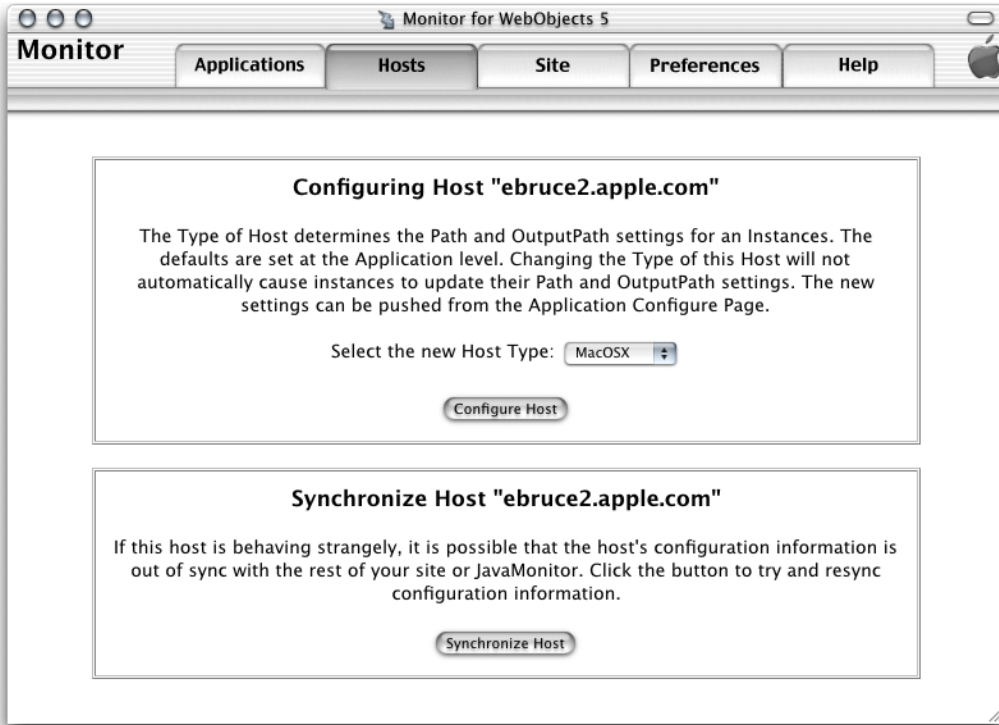
3. Choose the appropriate platform from the pop-up menu.
4. Click Add Host.

Figure 6-2 shows the Hosts page after a host has been added.

**Figure 6-2** Newly added host in Monitor

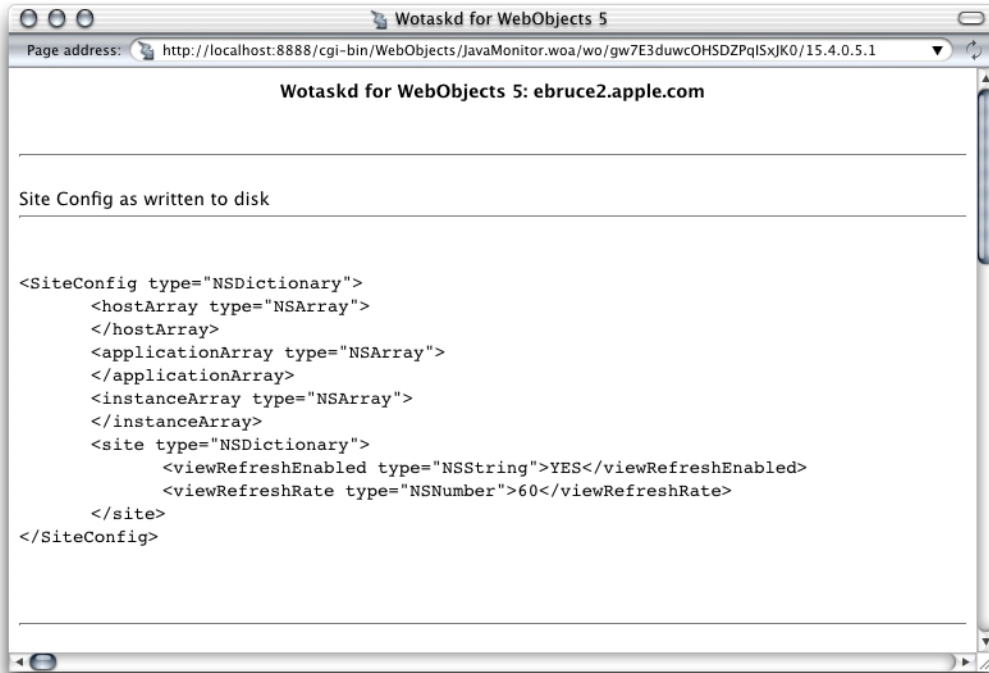
## Configuring a Host

To change the configuration of an application host, click the Config button on the Hosts page. A page similar to the one in [Figure 6-3](#) appears. It allows you to set the type of the host and to resynchronize configuration information if needed.

**Figure 6-3** Host configuration page

## Viewing a Host's Configuration

To display the configuration for a host, click YES, on the Hosts page. A page similar to the one in Figure 6-4 is displayed in a separate Web browser window.

**Figure 6-4** Host configuration information page

This page displays the current state of the host in several sections.

The section visible in [Figure 6-4](#) (page 72) shows the contents of the host's `SiteConfig.xml` file. For more information on the `SiteConfig.xml` file, see ["Configuration Files"](#) (page 57).

The next section shows the adaptor configuration sent to local HTTP adaptors, which lists all running application instances that wotaskd knows about (this includes Monitor processes).

```
<?xml version="1.0" encoding="ASCII"?>
<adaptor>
  <application name="JavaMonitor">
    <instance id="-8888" port="8888" host="ebruce2.apple.com"/>
  </application>
```



## Deployment Tasks

```
</adaptor>
```

Note that the instance ID of the Monitor process is negative. When `wotaskd` receives a lifebeat from an unregistered application instance (one that has not been added to your site through Monitor), it discloses the instance to the HTTP server with a negative ID number. This allows developers (internal users) to connect to development instances through the HTTP adaptor for testing purposes. To address security concerns, external users can connect to instances with negative ID numbers only if they know the instance's port number. However, this behavior can be disallowed by setting `WODirectConnectEnabled` to `false`. See [“WODirectConnectEnabled”](#) (page 117) for details.

To connect to a development instance through the HTTP server, the instance must run on the same computer that the HTTP server runs on, and the computer must be the `localhost`.

Next is the local adaptor configuration sent to remote HTTP adaptors. This section lists all instances that are active, registered (configured through Monitor), and available to external users.

```
<?xml version="1.0" encoding="ASCII"?>
<adaptor>
  <application name="Payroll" urlVersion="4">
    <instance id="1" port="2001" host="eBruce.apple.com"/>
  </application>
  <application name="HR" urlVersion="4">
    <instance id="1" port="2002" host="eBruce.apple.com"/>
  </application>
</adaptor>
```

The next section shows the contents of the HTTP adaptor configuration file. If you tell `wotaskd` to write the HTTP adaptor configuration file, it lists all the registered application instances on the site, whether they are running or not. (This file is identical across all the site's application hosts.) See [“The HTTP Adaptor Configuration File”](#) (page 44) for more information and an example of the file's contents.

The last section lists information on the `wotaskd` process's environment, including its port and multicast address.

```
The Configuration Directory is: /Library/WebObjects/Configuration/
Wotaskd is NOT writing WOConfig.xml to disk
```

## Deployment Tasks

```
The multicast address is: 239.128.14.2
This wotaskd is running on Port: 1085
Wotaskd is NOT responding to Multicast
WOAssumeApplicationIsDeadMultiplier is 4
The System Properties are: ...
```

## Installing Applications

---

Before you can deploy applications on your site, you have to install them on the hosts on which you want to run instances of them. Most applications have files of two types:

- **application files**, which store the application's logic
- **HTTP-server resources**, which store resources that can be shared among applications

The developer tool Project Builder (used to develop WebObjects applications) can be used to install an application on a host. With it, you can create an application bundle (a directory in the file system) with source code and resources the application needs to run. All the application files for the application can be included in the application bundle; however, doing so exposes the application's source to outside agents.

Instead, you should perform a split installation, installing most of the application in a directory that is not accessible to the outside world. The nonsensitive resources can be placed on the HTTP server's `Document Root` directory.

To perform a split installation of your application, navigate to your project's directory and execute the following command as the `root` user using your shell editor:

```
pbxbuild install -buildstyle WebServer DSTR00T=/
```

This places the application files in the `/Library/WebObjects/Applications` directory and the HTTP server resources in the `Document Root/WebObjects` directory.

## Deployment Tasks

You can place application files in any directory of the application host. HTTP-server resources, however, must be placed in the `Document Root` directory of the application host. Make sure that you use the following organization:

```
<HTTP server Document Root>/
  WebObjects/
    AppName.woa/
      Contents/
        WebServerResources/
          <resource files>
```

## Setting Up Applications

---

To deploy an application on your site, it must be installed in the appropriate directories on the hosts that run instances of it. For more information, see [“Installing Applications”](#) (page 74).

Setting up an application in your site involves three main steps:

- **Adding the application**, which is described in [“Adding an Application”](#) (page 76).
- **Configuring the application**, which includes the following subtasks:
  - defining a default configuration for new instances of the application
  - defining the recipients of email notifications
  - defining a schedule (only available for existing instances)
  - choosing a load-balancing algorithm

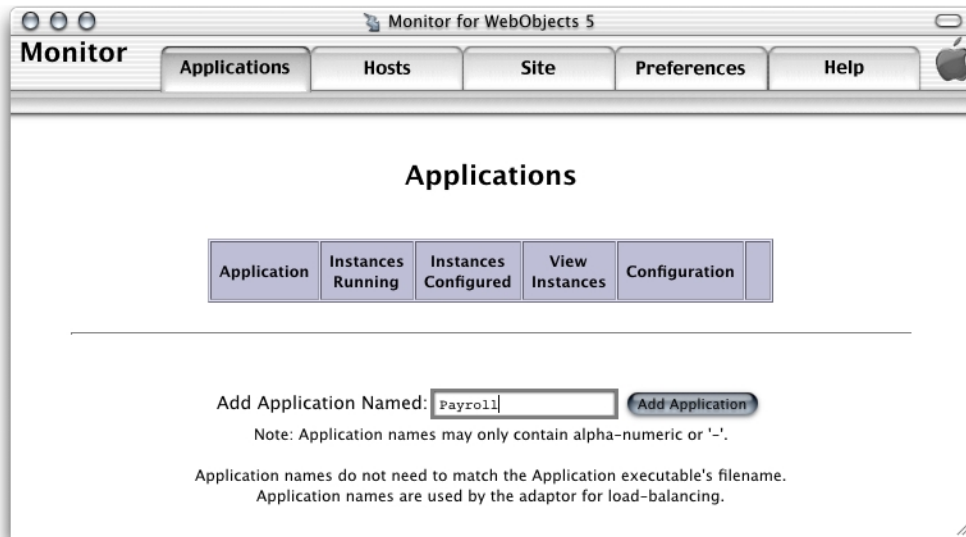
These steps are detailed in [“Configuring an Application”](#) (page 76).

- **Adding application instances**, which is described in [“Adding Application Instances”](#) (page 83).

## Adding an Application

You add applications to your site using Monitor's Applications page, shown in Figure 6-5.

**Figure 6-5** Adding an application using Monitor's Applications page



Follow these steps to add an application:

1. Enter the application's name (without the extension) in the Add Application Named text field.
2. Click Add Application.

## Configuring an Application

After you add an application, the application configuration page is displayed. This page has five major sections, which you can show and hide using the disclosure triangles:

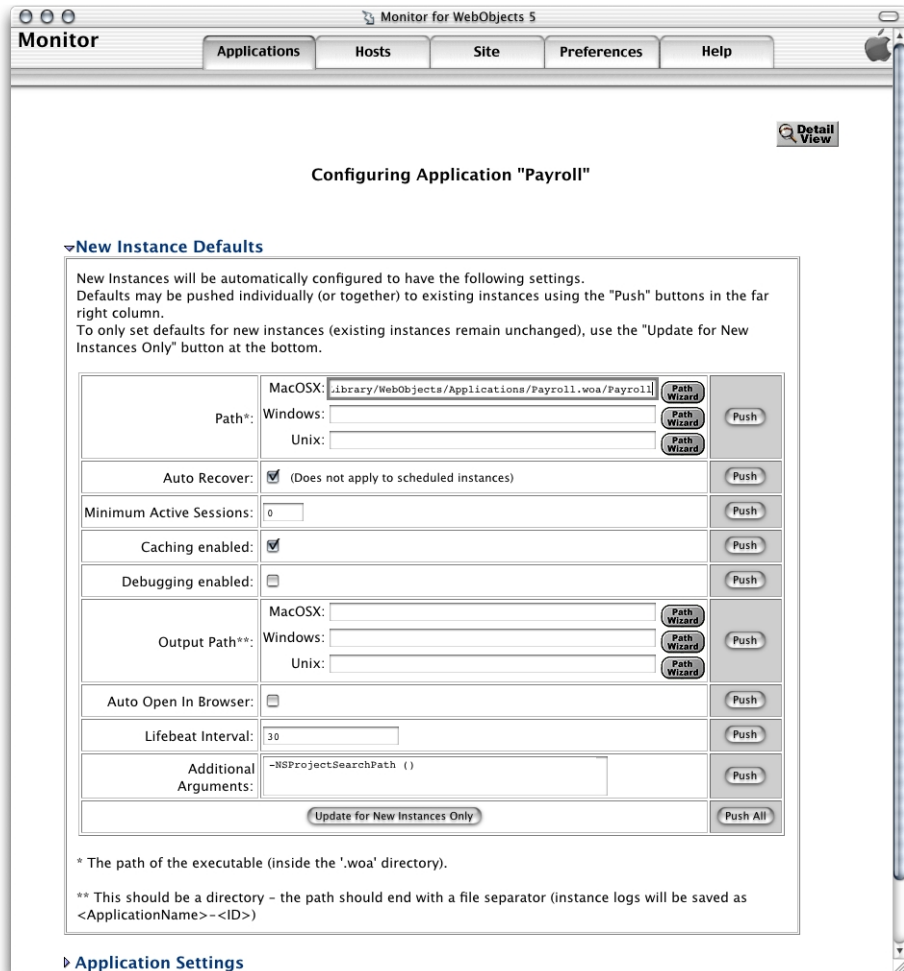
## Deployment Tasks

- The [New Instance Defaults](#) section lets you set the default values for the instance settings for application instances you add afterward. See “[New Instance Defaults](#)” (page 77) for details.
- The [Application Settings](#) section contains properties that apply to all the instances of the application. For more, see “[Application Settings](#)” (page 79).
- The [Scheduling](#) section allows you to individually schedule instances to restart at specific intervals. See “[Scheduling](#)” (page 80) for details.
- The [Email Notifications](#) section is where you specify the list of email addresses you want email notifications to be sent to. For details, see “[Email Notifications](#)” (page 81).
- The [Load Balancing and Adaptor Settings](#) section lets you choose the algorithm that the HTTP adaptor uses to perform load balancing among the instances of the application. See “[Load Balancing and Adaptor Settings](#)” (page 82).

## New Instance Defaults

---

[Figure 6-6](#) shows the section of the application configuration page that allows you to set defaults for the application instances you create afterward and for current ones (which are updated after you restart them). For details on each of the properties shown on this page, see “[Instance Settings](#)” (page 111).

**Figure 6-6** The New Instance Defaults section of the application configuration page

Here's an explanation of the buttons you see on the page:

- **Push** updates the value of the property for new and configured (registered) instances of the application. The changes take effect after the instances are restarted.

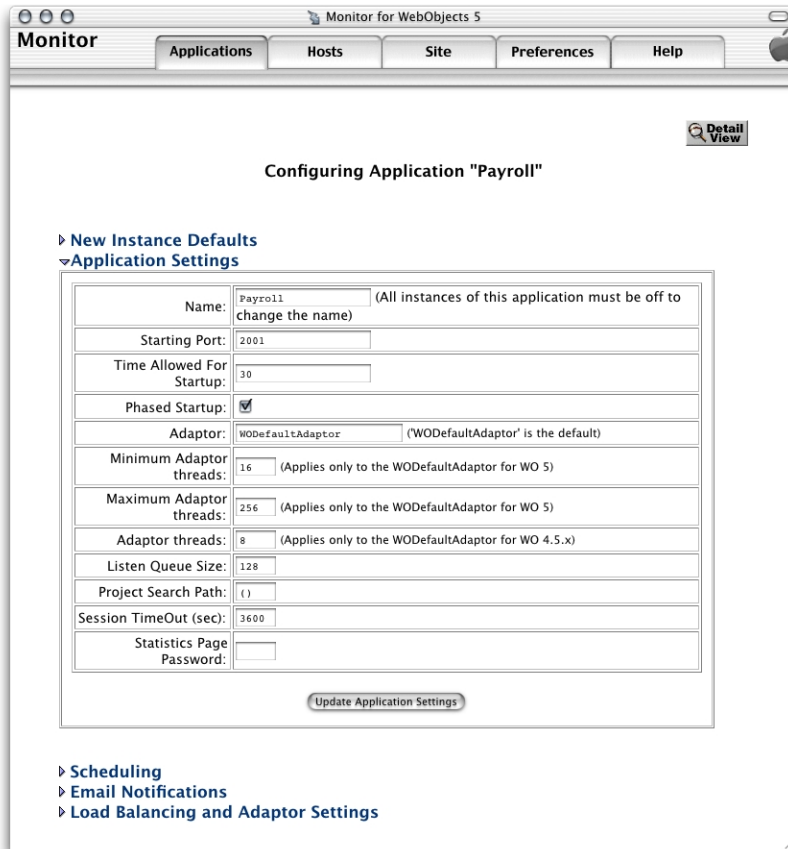
### Deployment Tasks

- **Push All** updates the value of the all the properties for new and configured instances of the application. As with Push, the changes become effective when the instances are restarted.
- **Update for New Instances Only** sets the defaults to be used when you create new instances of the application. The properties of existing instances are not changed.
- **Path Wizard** opens a tool that allows you to navigate through a host's file system.

### Application Settings

---

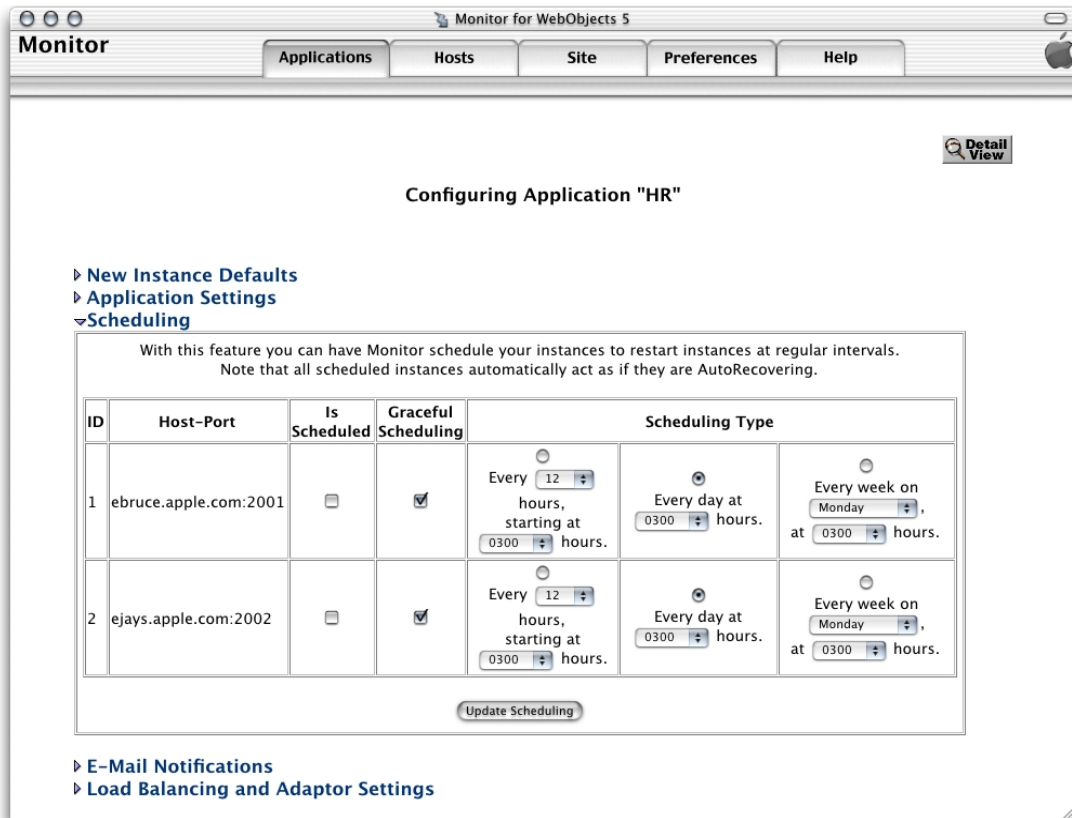
Figure 6-7 shows the Application Settings section of the application configuration page. In it, you define application settings that apply to all the instances of the application. For details of the properties shown in this section, see “Application Settings” (page 110).

**Figure 6-7** The Application Settings section of the application configuration page

## Scheduling

Figure 6-8 shows the Scheduling section of the application configuration page. After you add instances of an application, you can schedule them individually here. For details, see “Scheduling Settings” (page 116).

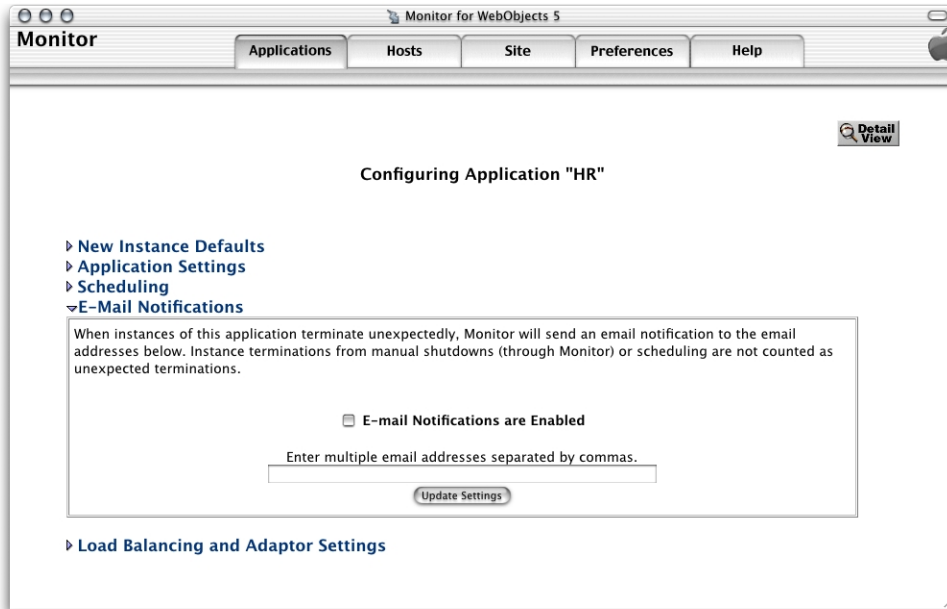


**Figure 6-8** The Scheduling section of the application configuration page

## Email Notifications

**Note:** Before you can configure email notifications, you have to tell Monitor which SMTP server to use. See “Configuring Sites” (page 91).

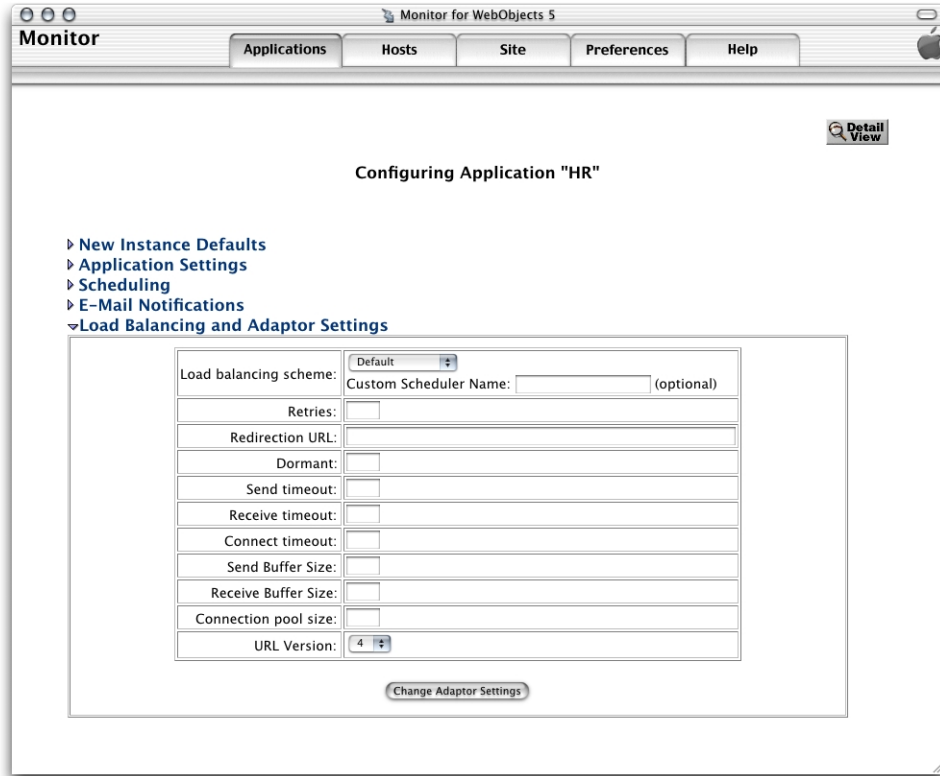
Figure 6-9 shows the **Email Notification Settings** section of the application configuration page. In it you can enter the email addresses of people that are to be notified when instances of the application terminate unexpectedly.

**Figure 6-9** The Email Notifications section of the Application Configuration page

## Load Balancing and Adaptor Settings

Figure 6-10 shows the Load Balancing and Adaptor Settings section of the application configuration page. This is where you enter values for the HTTP adaptor's configuration properties, including the load-balancing algorithm the adaptor will use to balance user load among the application's instances. These settings override the values entered in the HTTP Adaptor Settings section of the Site page. For information on the properties you can set, see "Load Balancing and Adaptor Settings" (page 114).

**Figure 6-10** The Load Balancing and Adaptor Settings section of the application configuration page

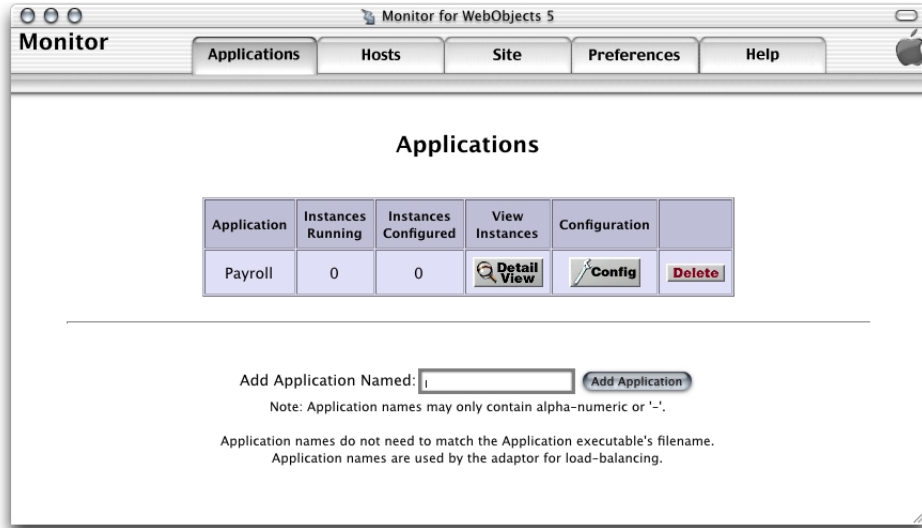


## Adding Application Instances

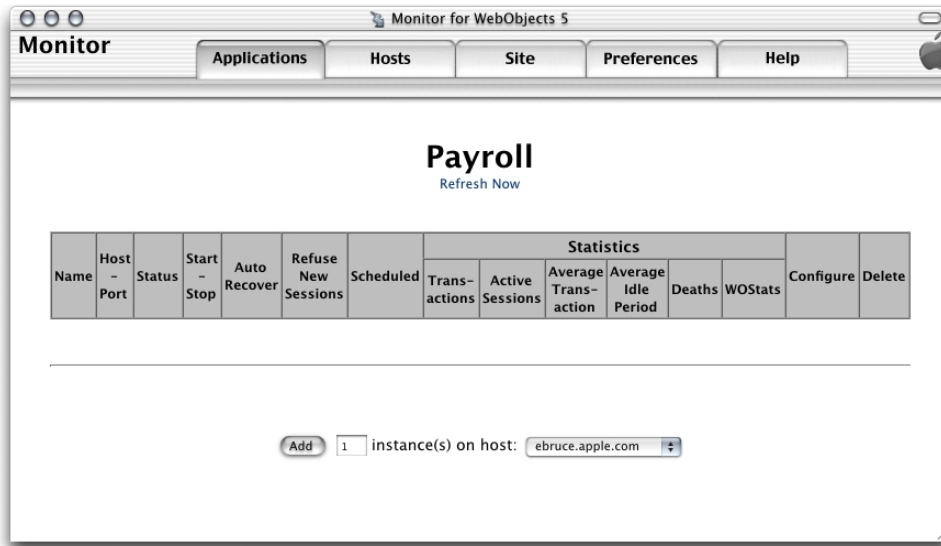
After you have configured an application in Monitor, you can create application instances with ease.

1. In Monitor, click the Applications tab. The Applications page is displayed, as shown in Figure 6-11.

**Figure 6-11** The Applications page with one application



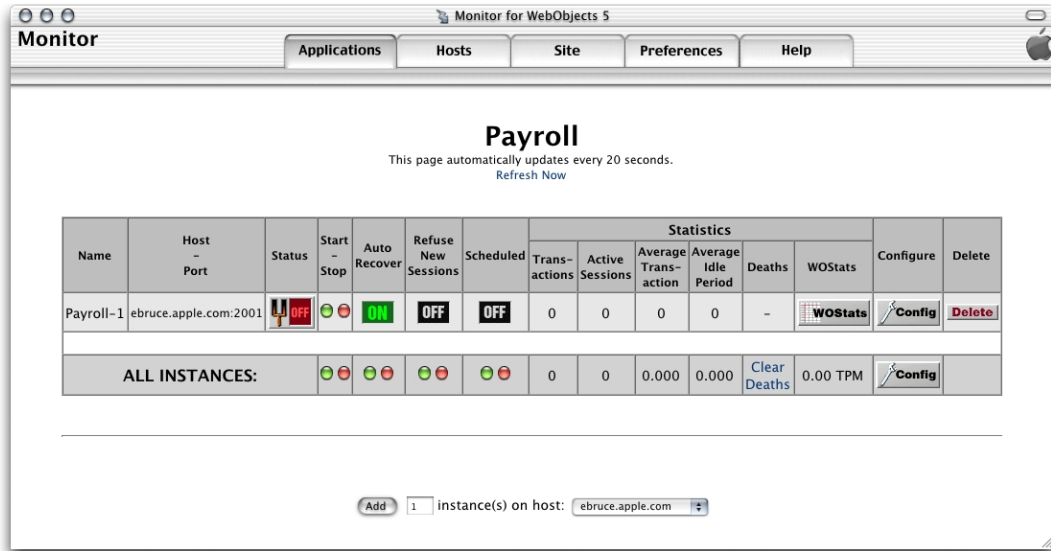
2. Click the Detail View button under View Instances. The application detail page is displayed, as shown in [Figure 6-12](#).

**Figure 6-12** The application detail page

3. Enter the number of instances you want to add in the text input field.
4. Choose the application host you want the instances to run on from the pop-up menu.

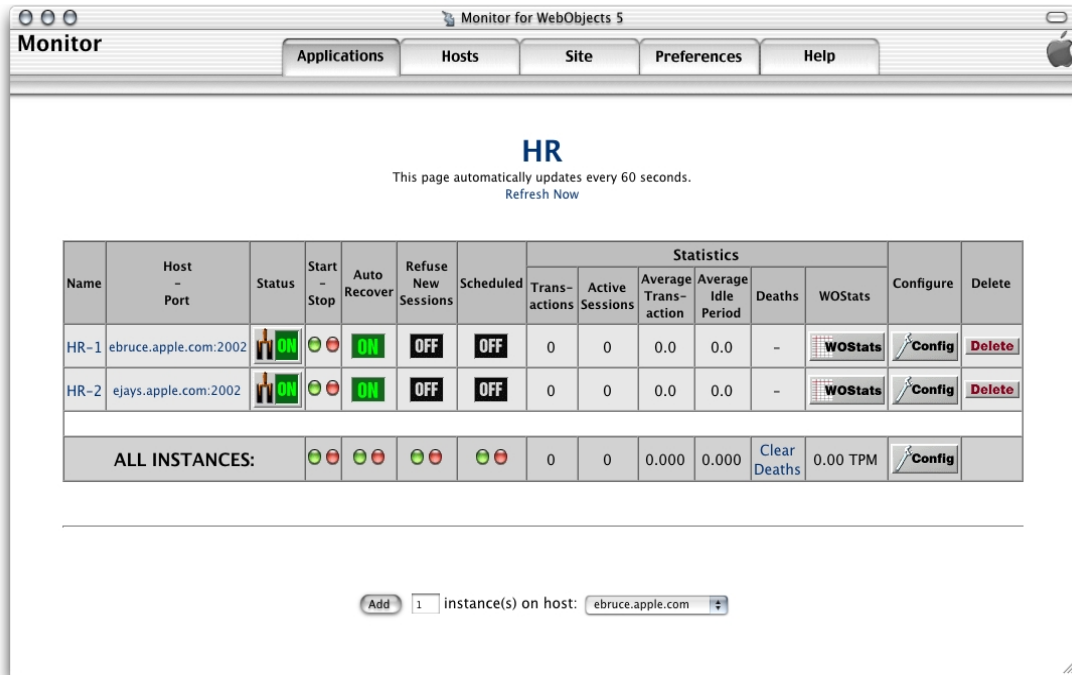
The application must be installed on the host you choose; otherwise, an error message is displayed when you try to start the instance. See [“Installing Applications”](#) (page 74) for details.

5. Click Add. Your Web browser displays a page like the one in [Figure 6-13](#).

**Figure 6-13** The application detail page after an instance has been added

The Status column indicates whether the instance is on or off. The first time the page is displayed, the newly added instances are off. After a moment (or if you click Refresh Now), and if [Auto Recover](#) is enabled for the instance, the page refreshes, showing that the instance is active. (You can change the length of the interval between the automatic updates of the application detail page; see “[Setting Monitor Preferences](#)” (page 94) for details.)

[Figure 6-14](#) shows the application detail page of the HR application, with two instances configured.

**Figure 6-14** The application detail page with two instances added

The following list describes the instance configuration information that appears in the application detail page:

- **Page heading:** A link to the application through the HTTP adaptor. When you click it, the adaptor uses load balancing to determine which of the application's instances receives the request. Then, your Web browser displays a new window showing your application's entry page. For this to work, the HTTP adaptor URL has to be set ("Configuring Sites" (page 91) shows you how to do this).
- **Name:** A link to the instance through the HTTP adaptor. When you click it, the request goes through the adaptor, but it's not load balanced. The URL is derived in the same way the URL of the page heading is derived, but with the addition of the instance number. Such a URL looks similar to the following:  

```
http://eBruce2.apple.com/cgi-bin/WebObjects/HR.woa/1
```
- **Host-Port:** A direct link to the instance; does not go through the HTTP adaptor.

## Deployment Tasks

- **Status:** Tells whether the instance is on, off, starting, or stopping.
- **Start-Stop:** Click the green button to turn the instance on, or the red button to turn it off.
- **Auto Recover:** Click to toggle between ON and OFF. This is available only if the instance is not scheduled. See “[Auto Recover](#)” (page 112).
- **Refuse New Sessions:** Click to toggle between ON and OFF. When ON, the instance does not accept new users. This is available only if the instance is not scheduled.
- **Scheduled:** Click to toggle between ON and OFF. When ON, the schedule defined for the instance is used.
- **Configure:** Click the Config button to go to the instance configuration page for the instance.
- **Delete:** Click the Delete button to delete the instance. You’ll see a confirmation page before the deletion takes place.

For an explanation of the columns under Statistics, see “[The Application Detail Page](#)” (page 101).

The row with the caption ALL INSTANCES contains buttons that perform some of the functions listed above on all the instances of the application. Clicking Config displays the application configuration page.

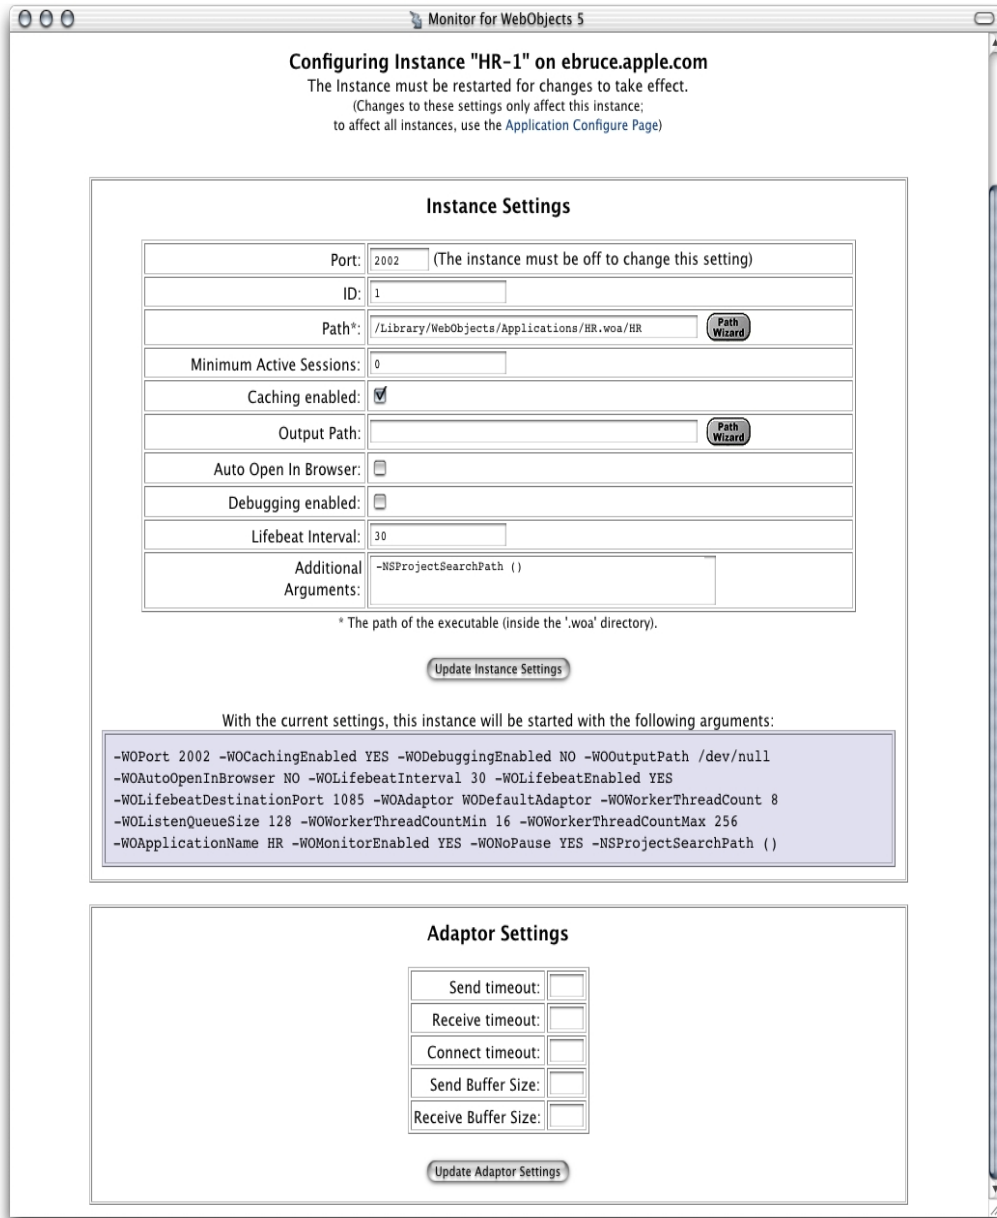
## Configuring Instances

---

After you have added an instance, you can change its configuration in the instance configuration page, shown in [Figure 6-15](#). You can access this page through the instance’s Config button in the application detail page. It contains two sections: [Instance Settings](#) and [Adaptor Settings](#).



Figure 6-15 Instance configuration page



### Deployment Tasks

#### Instance Settings

---

This section is very similar to the [New Instance Defaults](#) section of the application configuration page. It has two additional properties: **ID** and **Port**, which can only be changed after an instance has been added. For details, see “[ID](#)” (page 112) and “[Port](#)” (page 113).

#### Adaptor Settings

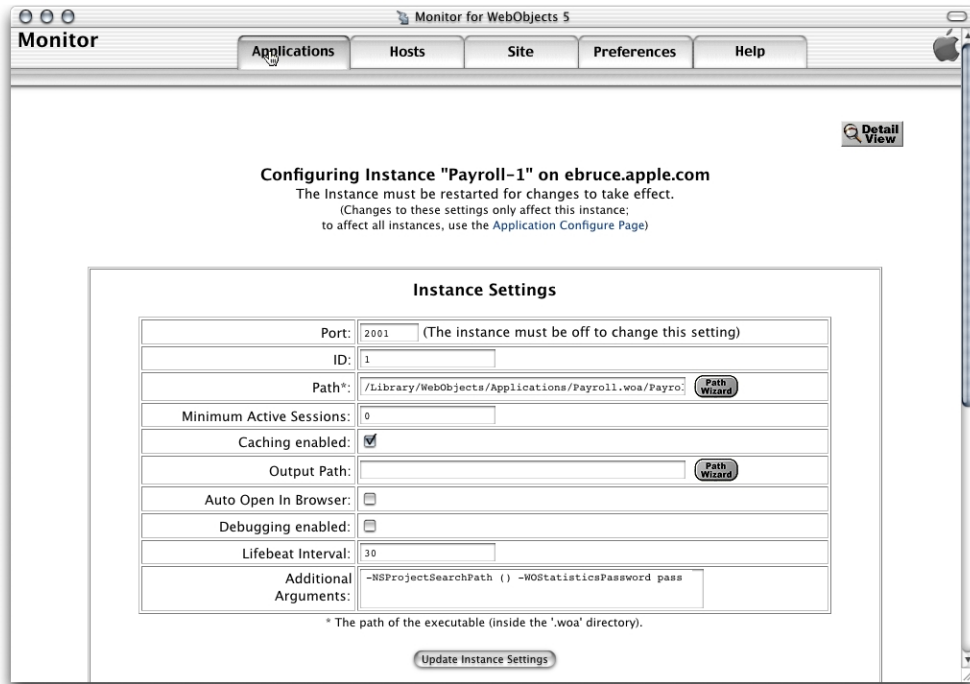
---

In this section you can change a subset of the properties available in the [Load Balancing and Adaptor Settings](#) section of the application configuration page. For details, see “[Load Balancing and Adaptor Settings](#)” (page 114).

#### Setting a Password for the Instance Statistics Page

---

For each instance of your application, there’s a statistics page that displays information such as its running time and memory usage. See “[The Instance Statistics Page](#)” (page 102) for more information on this page. If you want to prevent outside agents from gaining access to the instance statistics page, you can set a password in the Instance Settings section of the instance configuration page. Add the following to the Additional Arguments property: `-W0StatisticsPassword password`. [Figure 6-16](#) shows an example where the `W0StatisticsPassword` argument has been added to the Additional Arguments field.

**Figure 6-16** Setting a password for an instance's statistics page

## Configuring Sites

When you click the Site tab in Monitor, the site configuration page is displayed. It contains three sections:

- **HTTP Adaptor URL:** This is where you tell Monitor how to compose an application's URL, which used in the application detail page to connect you to instances of your application.

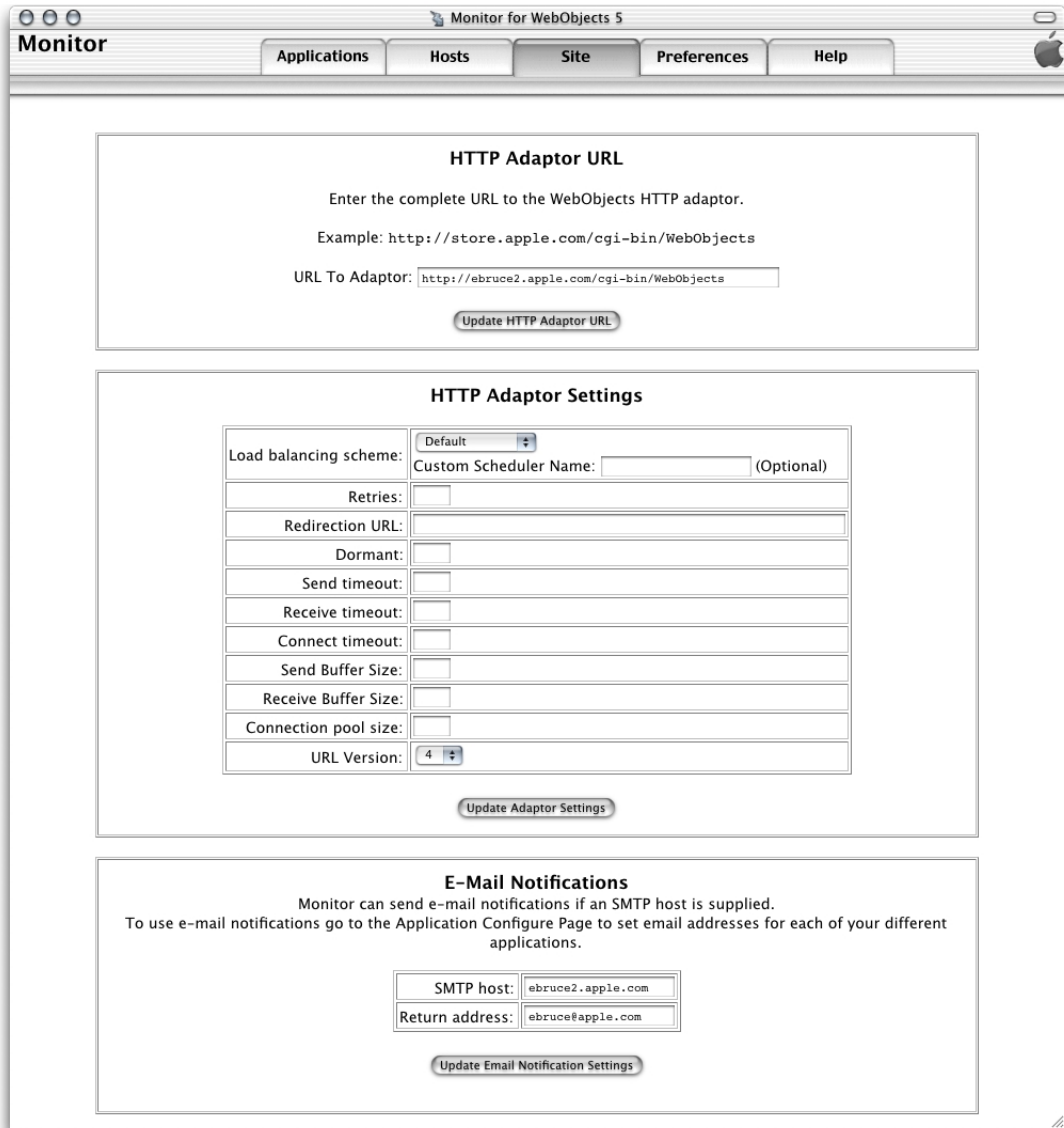
To set the URL for your application, enter the URL in the URL to Adaptor field.

### Deployment Tasks

- **HTTP Adaptor Settings:** This is where you set default HTTP adaptor settings for all your deployed applications. They can be overridden by each application. For more information, see “Load Balancing and Adaptor Settings” (page 82) and “Load Balancing and Adaptor Settings” (page 114).
- **Email Notifications:** Here’s where you specify the SMTP host and the return address that Monitor uses for email notifications.

Figure 6-17 shows the site configuration page:

**Figure 6-17** The site configuration page

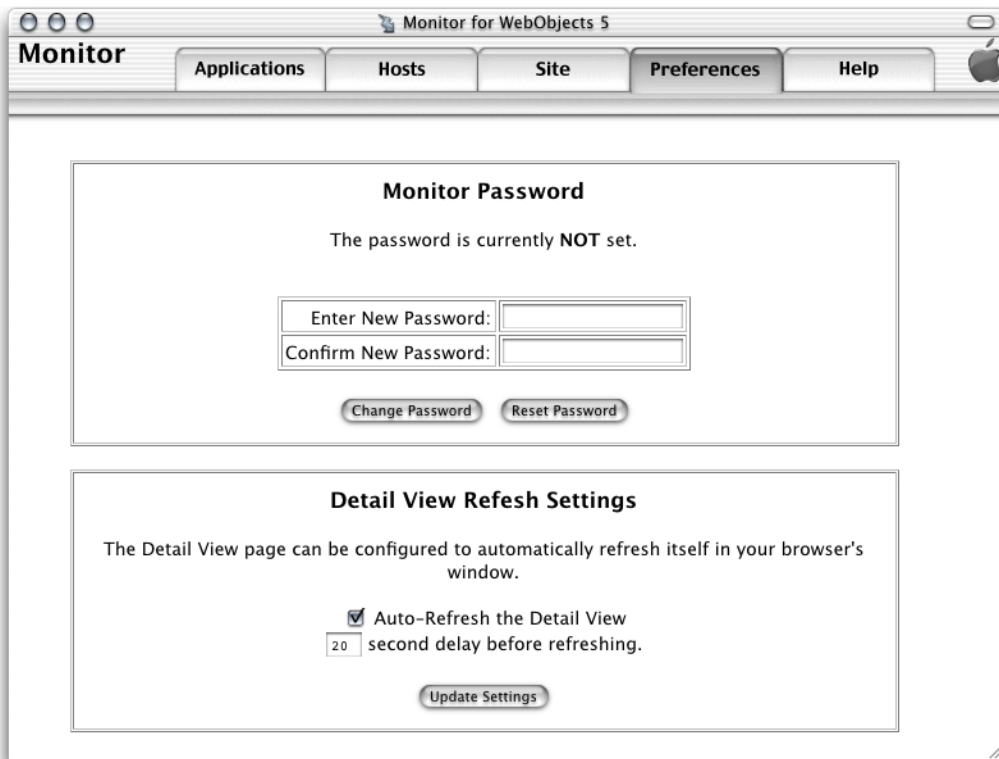


## Setting Monitor Preferences

---

When you click the Preferences tab in Monitor, the page in [Figure 6-18](#) is displayed. It contains two sections: Monitor Password and Detail View Refresh Settings.

**Figure 6-18** The preferences page



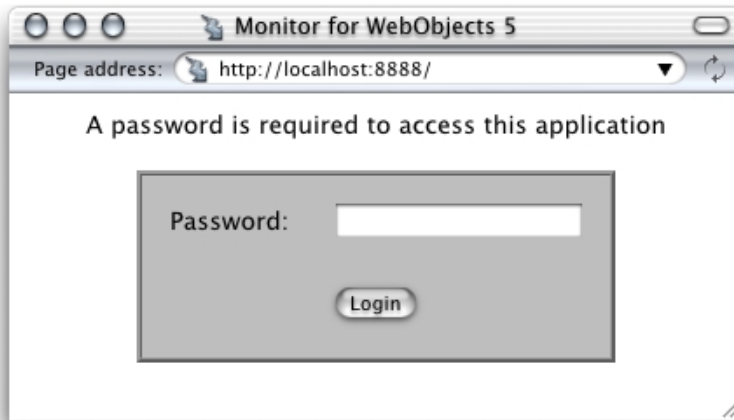
## Monitor Password

---

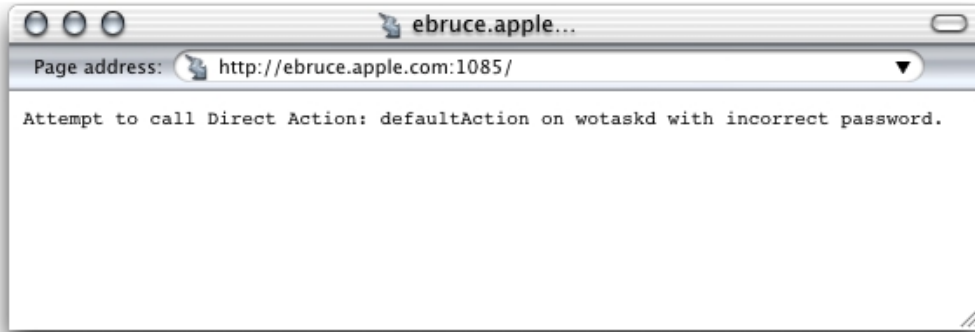
You can restrict access to Monitor by requiring its users to enter a password before they can use it. When a site is protected this way, the site's wotaskd processes are also protected; that is, you cannot directly obtain a wotaskd process's information by connecting to its port, as described in "Confirming That wotaskd Is Active" (page 31).

Figure 6-19 shows the login page that Monitor displays after you password-protect your site.

**Figure 6-19** Login page displayed by Monitor on a password-protected site



When you try to view the configuration of an application host on a site that you've password-protected by connecting to the appropriate wotaskd process's port, you'll see a page similar to the one shown in Figure 6-20 (the page's content varies according to your deployment platform).

**Figure 6-20** Page returned by wotaskd when the site is password-protected

On password-protected sites, you'll have to use Monitor to view an application host's configuration.

## Detail-View Refresh Settings

---

This section allows you to tell Monitor if you want it to refresh the application detail page and how often to do it.

## Load Balancing

---

**Note:** Load balancing occurs only between the hosts that the HTTP adaptor knows about. Adding a host in Monitor is not enough. For more information on how to configure hosts in the HTTP adaptor, see ["State Discovery"](#) (page 39).

Load balancing is a mechanism by which user-load is spread out among the instances of an application; these instances can be running on different hosts. Load balancing ensures that your site's hardware resources are used efficiently and with



## Deployment Tasks

the highest level of performance. The default load-balancing scheme used is Random. The following list describes how user load is distributed under each of the provided algorithms:

- **Random** assigns a user to an arbitrarily-chosen instance.
- **Round Robin** assigns users among instances sequentially.
- **Load Average** balances load by distributing users evenly among instances.

You can choose a load-balancing algorithm at two levels:

- **Site level:** You set a site-wide load-balancing algorithm in the HTTP Adaptor Settings section of the Site page. From then on, the applications in your site will use that load-balancing algorithm.
- **Application level:** You can override the site-wide load-balancing algorithm in the Load Balancing and Adaptor Settings section of the application configuration page.

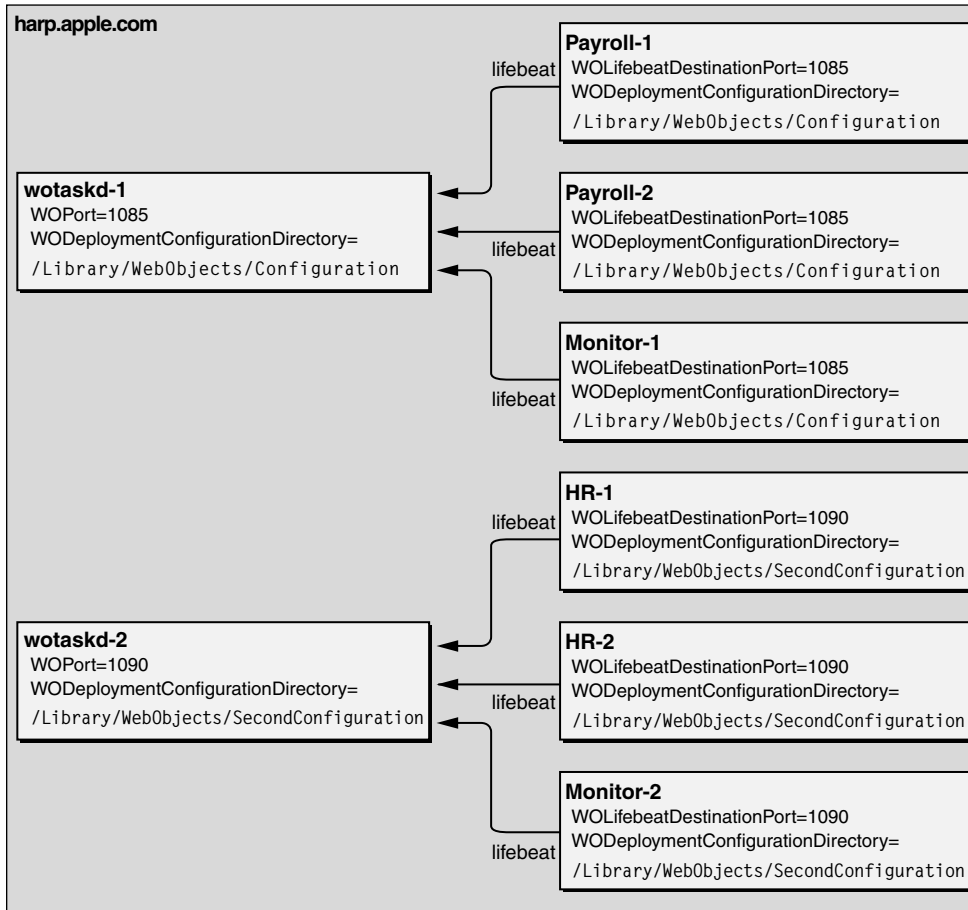
## Deploying Multiple Sites

---

You can deploy and configure separate sites on a set of computers by running multiple HTTP servers, each with its own HTTP adaptor. Such a deployment requires a separate group of wotaskd processes running on the same port. You also need an additional Monitor process to configure each site.

The default installation of the WebObjects Deployment package provides you with one site—one wotaskd process per host, running on port 1085. To create a second site, using the same hardware, you'll have to add an additional wotaskd process to each of the hosts you want to use.

What separates the environments from each other are the `WOPort` and `WOLifebeatDestinationPort` settings of each wotaskd process and the configuration directory used for each site. The application instances send their lifebeats to their `WOLifebeatDestinationPort`, while wotaskd processes listen for them in their `WOPort`. [Figure 6-21](#) illustrates two sites on one host.

**Figure 6-21** Multiple application environments on one computer

Because Monitor is not started by a wotaskd process, its `WOLifebeatDestinationPort` argument needs to be set to match wotaskd's `WOPort` setting.

For details on how to set the different command-line argument values required when starting wotaskd and Monitor processes for separate application environments, see “`WOPort`” (page 119), “`WOLifebeatDestinationPort`” (page 118), and “`WODeploymentConfigurationDirectory`” (page 120).

# Application Administration

---

After deploying applications, you should monitor their performance to find out, among other things, if you need to add instances to an application to improve response times. “[Monitoring Activity](#)” (page 99) shows you the kind of performance data you can collect about an application.

There are several steps you can take to improve your site’s performance; most of them have been explained in the preceding chapters. In the section “[Improving Performance](#)” (page 107) you’ll find a list of recommended measures that will help optimize your site’s operation.

## Monitoring Activity

---

There are several ways to obtain information about the applications deployed on your site. You can

- use Monitor
- analyze logs from application instances and adaptors
- view instance statistics (WOSTats) pages

## Monitoring Application Performance

---

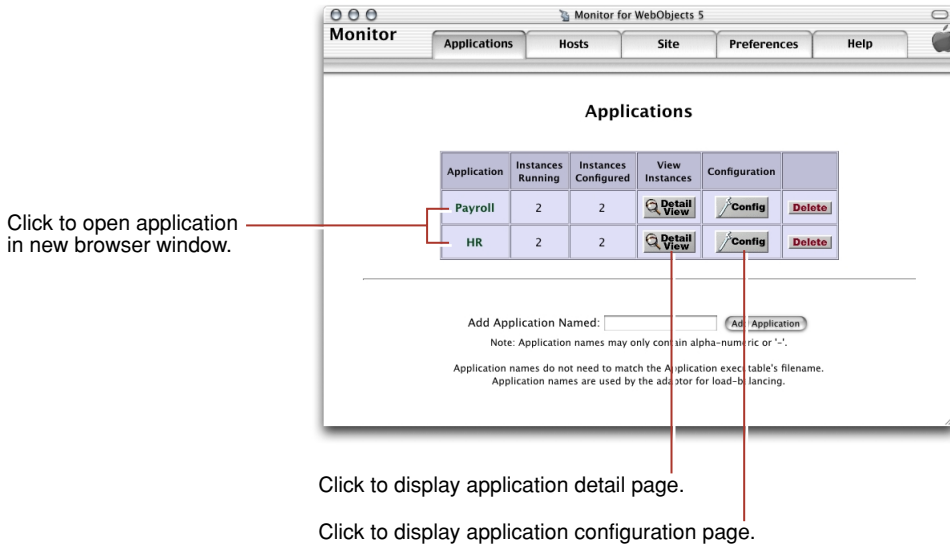
Monitor’s Applications page gives you an overall view of a site. [Figure 7-1](#) illustrates the kind of information you can access through the Applications page:

- configured applications

Application Administration

- number of configured instances per application
- number of running instances per application

**Figure 7-1** The Applications page



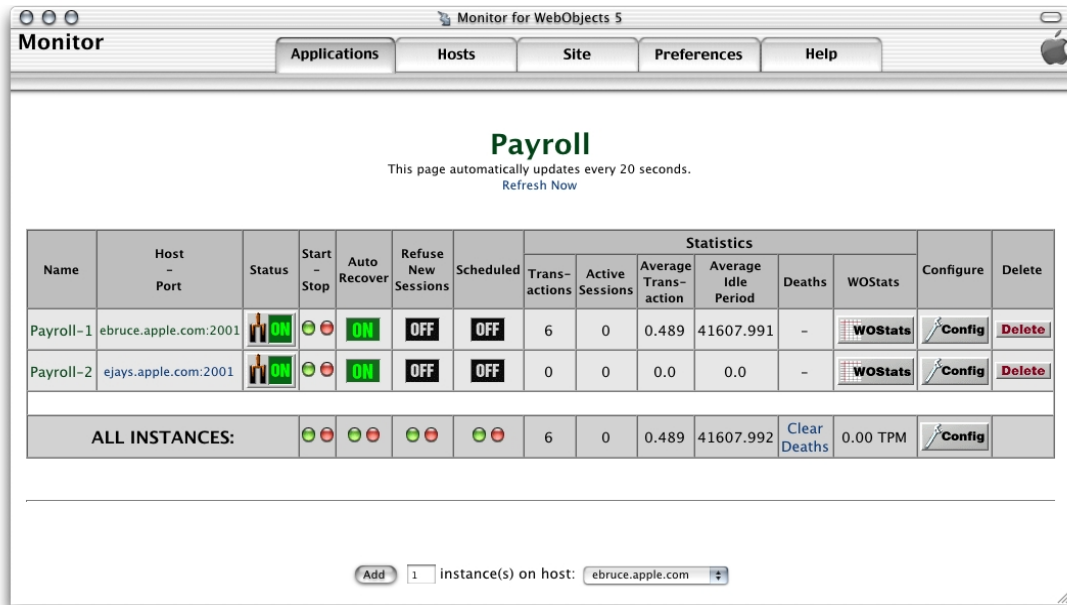
From this page, you can perform several tasks:

- **Connect to an instance of an application.** The entry page of the application is displayed in a new Web browser window.
- **Display the application detail page of an application.** For details on the information provided, see “Adding an Application” (page 76).
- **Display the application configuration page.** See “Configuring an Application” (page 76) for more information.
- **Delete an application.** You confirm that you really want to delete the application (including all of its instances) in a confirmation page before the deletion takes place.

## The Application Detail Page

Figure 7-2 depicts the application detail page.

**Figure 7-2** The application detail page



The application's name is displayed centered and in bold letters. When there are active instances, it is a link to the application through the HTTP adaptor; the request is load balanced among the application hosts configured in the adaptor. (For this to work, the HTTP adaptor URL property must be set; see "Configuring Sites" (page 91) for details.) If no value has been entered for the property, the default URL (`http://localhost/cgi-bin/WebObjects/`) is used instead.) The URL used to connect to the application looks similar to the following:

```
http://localhost/cgi-bin/WebObjects/Payroll
```

The following list describes the performance-related information shown on the application detail page in the Statistics column group:

## Application Administration

- **Transactions:** The number of requests the instance has received since it was started.
- **Active Sessions:** The number of active sessions (users) currently maintained by the instance.
- **Average Transaction:** Average time, in seconds, the instance has taken to process requests.
- **Average Idle Period:** The average time, in seconds, that the instance is idle (average time between requests).
- **Deaths:** The number of unexpected failures or deaths the instance has had since it was started. These exclude scheduled shutdowns or manual shutdowns through Monitor.
- **WOSTats:** Click to display the statistics page for the instance in a new Web browser window. See [“The Instance Statistics Page”](#) (page 102) for details.

Before you can view the instance statistics page, you have to enter the password you set on the instance configuration page. See [“Setting a Password for the Instance Statistics Page”](#) (page 90) for details.

The row with the caption ALL INSTANCES contains application-wide performance information, including the average number of transactions processed per minute (TPM).

## The Instance Statistics Page

---

You can access the instance statistics (WOSTats) page of an instance through Monitor or directly through a Web browser.

To use Monitor, go to the application detail page and click WOSTats for the instance whose statistics you want to view.

To use your Web browser, access the following URL:

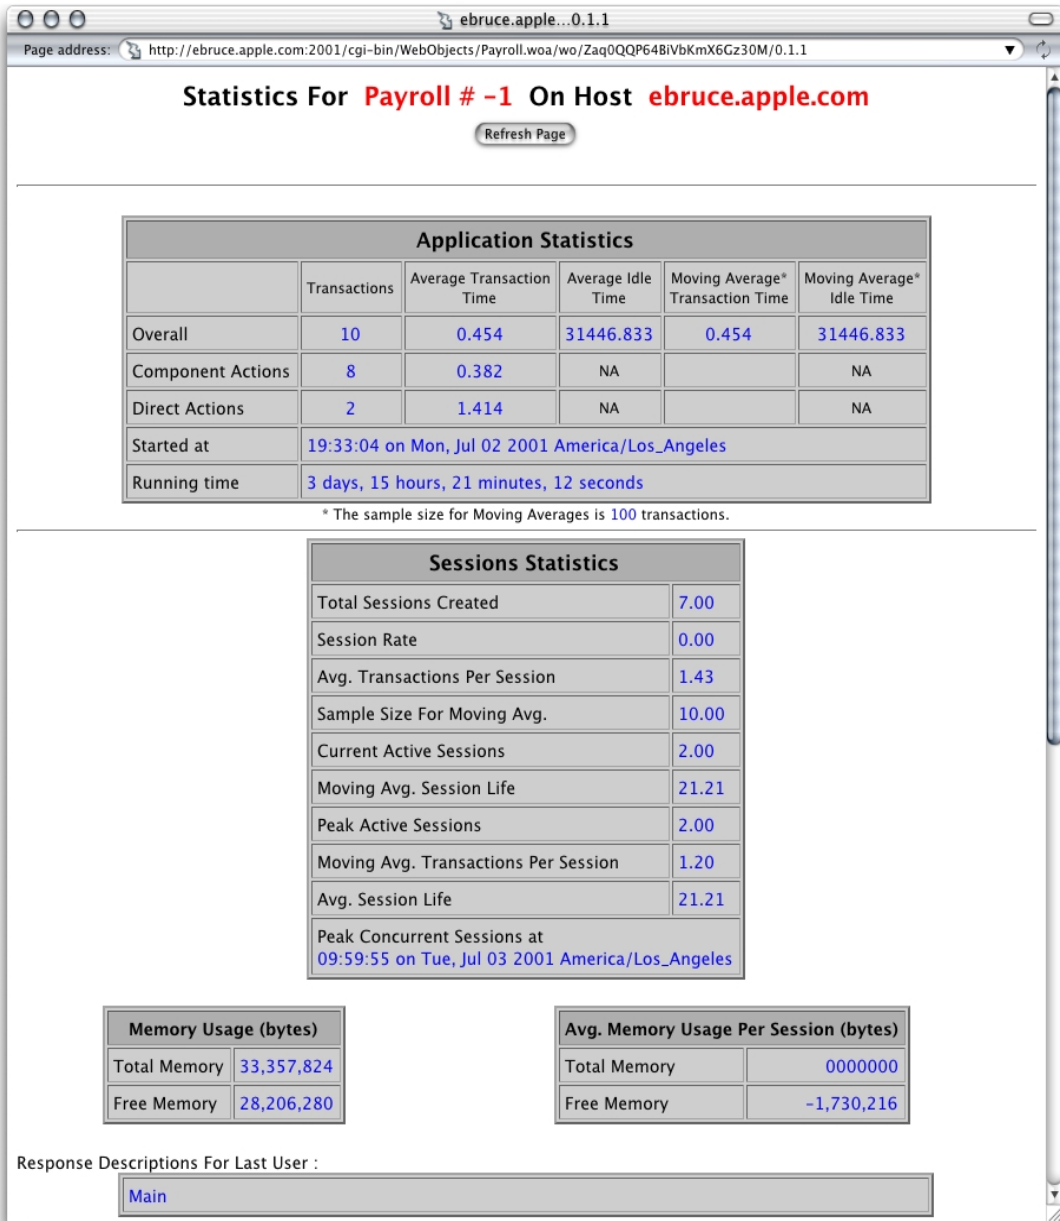
```
http://myhost/cgi-bin/WebObjects/MyApp.woa/wa/WOSTats
```

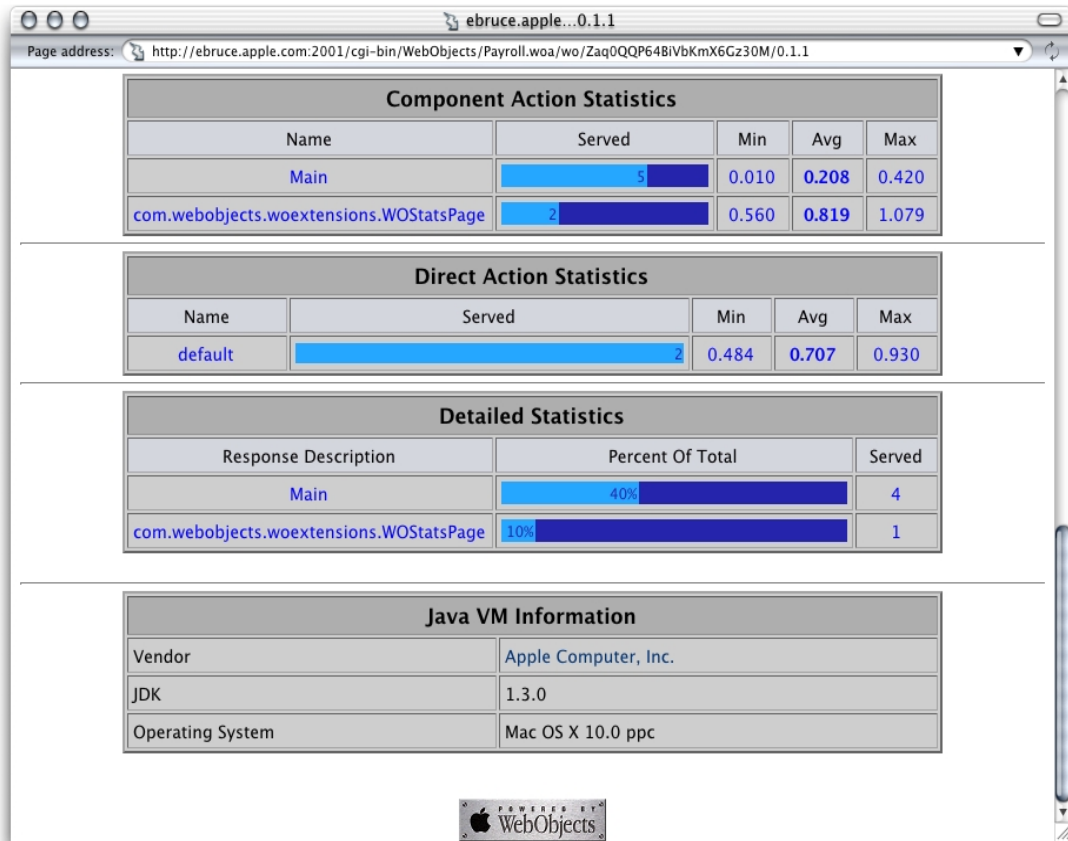
If there are multiple instances, specify the instance number as well:

```
http://myhost/cgi-bin/WebObjects/MyApp.woa/1/wa/WOSTats
```

[Figure 7-3](#) and [Figure 7-4](#) (page 104) show the information the instance statistics page provides.

**Figure 7-3** The instance statistics page—part 1 of 2



**Figure 7-4** The instance statistics page—part 2 of 2

## Logging and Analyzing Application Activity

WebObjects applications can record information in a log file that can be analyzed by a Common Log File Format (CLFF) standard analysis tool. Applications do not maintain this log file by default; log-file recording must be enabled through an application's code. When enabled, the application records a list of components accessed during each session. By default, only component names are recorded, but you may add more information.



## Logging and Analyzing Adaptor Activity

---

To enable adaptor logging, you create a file called `logWebObjects` in the temporary directory of the computer where the HTTP server runs. When logging is enabled, the adaptor logs its activity in a file called `WebObjects.log` in the temporary directory. Logging adaptor activity significantly decreases performance. Use this feature only as a troubleshooting aid; do not use it during regular deployment.

The location of the temporary directory depends on the platform:

- Mac OS X Server and Solaris: `/tmp`
- Windows 2000: The directory indicated by the `TEMP` environment variable.

### Creating the Adaptor Log File

---

On UNIX-based platforms, do the following to create the `logWebObjects` file (you must have root privileges):

1. Start a command-shell window.
2. Set the working directory to the temporary directory.
3. Enter the following command:

```
touch logWebObjects
```

On Windows 2000, create a blank file using a text editor and save it as `logWebObjects` in the temporary directory.

You can use the `tail` (UNIX) or `type` (Windows 2000) commands to display the adaptor log file in your console.

On UNIX-based systems, use the following:

```
tail -f WebObjects.log
```

On Windows 2000, type the following in a DOS prompt:

```
type WebObjects.log
```

## Analyzing the Adaptor Log File's Contents

---

You can analyze the information in the log to find out such things as which applications are being requested, which applications are being auto-started, and what the HTTP headers of the requests are. You can also use the log to verify that the HTTP adaptor is properly configured for load balancing.

The following excerpt includes an error message that indicates that an instance of Payroll wasn't running when a request for it came in:

```
Info: <WebObjects Apache Module> new request: /cgi-bin/WebObjects/Payroll
Debug: App Name: Payroll (7)
Info: Specific instance Payroll: not found. Reloading config.
```

After Payroll is started, the same request produces the following log-file entries:

```
Info: New response: HTTP/1.0 200 Apple WebObjects
Info: ac_newInstance(): added Payroll:2 (2001)
Info: V4 URL: /cgi-bin/WebObjects/Payroll
Info: loadaverage: selected instance at index 4
Info: Selected new app instance at index 4
Debug: Composed URL to '/cgi-bin/WebObjects/Payroll.woa/1'
Info: New request is GET /cgi-bin/WebObjects/Payroll.woa/1 HTTP/1.0

Info: Sending request to instance number 1, port 2001
Info: Trying to contact Payroll:1 on (2001)
Info: attempting to connect to ebruce.apple.com on port 2001
Info: Created new pooled connection [1] to ebruce.apple.com:2001
Info: Using pooled connection to ebruce.apple.com:2001
Info: Payroll:1 on (2001) connected [pooled: Yes]
Info: Request GET /cgi-bin/WebObjects/Payroll.woa/1 HTTP/1.0
    sent, awaiting response
Debug: ac_readConfiguration(): skipped reading config
Info: New response: HTTP/1.0 200 Apple WebObjects
Info: Payroll 1 load avg = 1
Info: received ->200 Apple
```

## Improving Performance

---

Performance is a major concern of website administrators. This section provides a list of areas to check to achieve the maximum performance possible.

- Configure your operating system so that it delivers the highest performance for your needs. Check your operating system's documentation and your HTTP server's documentation for performance-tuning information.

- When possible, use an API-based HTTP adaptor instead of a CGI adaptor.

- Make sure that the applications are written to perform optimally.

The WebObjects developer documentation contains coding suggestions that help improve the performance of WebObjects applications.

- Enable component-definition caching for all applications.

When applications are deployed, component-definition caching should be enabled so that each component's HTML and declarations files are parsed only once per session.

- Shut down and restart application instances periodically.

Scheduling instances to shut down and restart periodically increases application performance and reliability by reducing the effects of memory leaks. For more on scheduling, see "[Scheduling Settings](#)" (page 116).

If your applications use custom scheduling algorithms to shut down themselves, you should not use Monitor's scheduling feature. Instead, just use Monitor's auto-recover feature. For more information, see "[Auto Recover](#)" (page 112).

- Consider changing the physical configuration of your system.

Determine the size of a single application instance (you can find this data on the application's instance statistics page) and multiply it by the number of instances you intend to run on a given computer. (For information on the instance statistics page, see "[The Instance Statistics Page](#)" (page 102).) The result is the amount of physical memory needed for that application. You have to add the memory required by the operating system, HTTP server, and any other applications that run constantly on the computer. The result is the amount of physical memory that should be installed on the computer.

## Application Administration

- Try to reduce the size of the application instance by limiting the amount of state that it stores. Set the session timeout value to ensure that sessions expire after a reasonable length of time. See “[WOSessionTimeout](#)” (page 119) for details on setting the session timeout interval for application instances.
- Make sure that all static content is served by the HTTP server, not your application.

If you use WebObjects Deployment mainly to deploy applications that access a data store, you’ll achieve the best performance with a dedicated data-store server and a separate server for WebObjects applications.

# Deployment Settings Reference

---

Monitor offers an interface that you use to set the values of many configuration properties of your site. But if you prefer to use your shell editor, there are command-line arguments available that you can use to set properties for each application instance, including Monitor and wotaskd processes.

The following sections describe the deployment properties you can set through Monitor or the command line. It contains two major sections:

- “[Application Configuration Settings](#)” (page 109) explains the properties you can set through Monitor.
- “[Command-Line Arguments](#)” (page 116) describes the function of the command-line arguments used when you start application instances from the command line (or use an application instance’s `Additional Arguments` property).

## Application Configuration Settings

---

There are several settings that Monitor, wotaskd, and application instances use in a deployment. You can set the value of these properties in Monitor, in the command line, or in an application’s `Properties` file (which you would have to create and place in the `Contents/Resources` directory of the application). The following sections describe all the settings and their function.

## Application Settings

---

The following sections describe the properties that apply to all the instances of an application. In Monitor, you set the value of these properties in the Application Settings section of the application configuration page.

### Adaptor

---

The default adaptor class an instance uses. (This is not the HTTP adaptor used by the HTTP server.) You use this if the application defines a subclass of the WOAdaptor class to be used to create adaptor objects (instead of using the WOAdaptor class itself).

### Adaptor Threads

---

The number of worker threads that the adaptor creates to handle requests to the application. It applies only to adaptor objects of the WODefaultAdaptor class in WebObjects 4.5.

### Listen Queue Size

---

Determines the depth of the listen queue. While the instance handles a request, the socket buffer can hold as many additional requests as this setting indicates before it starts refusing them. If you expect spikes in the traffic level of a specific application, consider increasing the value of this property. While it does not necessarily improve performance, or allow the instance to process more requests at sustained high loads, it may reduce the number of times an application's user has to retry to send a particular request during high-traffic periods.

Keep in mind that if an application instance's listen queue size becomes full and a request is refused by it, the request does not get redirected to another instance with space left over in its queue. The client will have to resend the request.

### Maximum Adaptor Threads

---

The maximum number of worker threads the HTTP adaptor creates to handle requests to the application. It applies only to adaptor processes of the WODefaultAdaptor class in WebObjects 5. The purpose of these threads is to process TCP or UDP packets; they have nothing to do with request processing.

## Minimum Adaptor Threads

---

The initial number of worker threads the adaptor creates to handle requests to the application. It applies only to adaptor processes of the WODefaultAdaptor class in WebObjects 5.

## Name

---

This property is used by the HTTP adaptor to implement load balancing . The adaptor can load-balance only between instances with the same application name. The property can be used to create groups of instances, even when the instances share the same executable file. This argument is set automatically for instances started by wotaskd.

## Phased Startup

---

If this option is selected, when a wotaskd process starts up, the instances that it manages that have [Auto Recover](#) selected are started one-at-a-time instead of all at once. This option is selected by default. For more information, see “[Auto Recover](#)” (page 112).

## Starting Port

---

The first port number to try to assign to a new instance of the application. If the port is in use by another instance (of any application), the next nonassigned port number is used.

## Time Allowed for Startup

---

The number of seconds Monitor waits for an instance to start before determining that the instance failed to start.

## Instance Settings

---

The following sections explain the properties that can be set in Monitor in the New Instance Defaults section of the application configuration page and in the Instance Settings section of the instance configuration page.

### Additional Arguments

---

Additional information to send to the instance when it's started. You can use this property to set the values of properties for which Monitor doesn't provide a user interface.

### Auto-Open In Browser

---

Determines whether the instance automatically opens a Web browser window to the application's URL (starting up the browser if necessary).

### Auto Recover

---

Indicates whether instances are restarted automatically whenever they become unresponsive or are manually shut down.

### Caching Enabled

---

Determines whether the instance caches component definitions instead of parsing HTML code and declaration files each time a request is processed. (A component is a Web page or a portion of one.)

### Debugging Enabled

---

Determines whether application instances print debugging messages to the standard error stream during startup.

### ID

---

The instance's identification number. Must be unique for the load-balancing process to operate correctly. (This property is not available in the New Instance Defaults section.)

### Lifebeat Interval

---

Determines the interval, in seconds, between lifebeats.



## Minimum Active Sessions

---

The minimum number of active sessions an instance must have before it can be terminated by Monitor. (A session is an object of the `WOSession` class that stores user-state information.) See “Scheduling Settings” (page 116).

## Output Path

---

Allows you to redirect the instance’s standard output and standard error streams to the directory specified. The file generated is named as follows:

```
<application name>-<instance ID>
```

## Path

---

Location of the launch script file for the application. For example, `MyApp` for Mac OS X and Solaris, and `MyApp.cmd` for Windows 2000.

## Port

---

The port on which the instance runs. (This property is not available in the New Instance Defaults section.)

## Project Search Path

---

List of file-system paths that WebObjects searches in rapid-turnaround mode. These are the paths where your projects are located (relative paths are relative to your project’s launch script). WebObjects searches for your projects in the order in which you specify the paths. You must specify the paths using the following format:

```
("firstPath", "secondPath", ...)
```

## Session Timeout

---

Specifies the time, in seconds, that passes after the last request is processed before the session times out.

## Statistics Page Password

---

Specifies the password that must be entered to gain access to the instance statistics (WOStats) page of an application instance.

## Email Notification Settings

---

You can tell Monitor to send a message to one of more email addresses when instances of the application terminate unexpectedly. (Manual or scheduled instance shut downs are not considered unexpected terminations.) You specify the value of these settings in the Email Notifications section of the Site page.

## Load Balancing and Adaptor Settings

---

The following sections explain the properties related to load balancing and other adaptor settings. In Monitor, you specify the value of these settings in the HTTP Adaptor Settings section of the site configuration page and in the Load Balancing and Adaptor Settings section of the application configuration page.

### Connect Timeout

---

The length of time, in seconds, before the adaptor gives up connecting to an instance.

### Connection Pool Size

---

The maximum number of simultaneous connections the adaptor should keep open for each configured instance.

### Dormant

---

The number of times the adaptor skips an instance of the application before trying again.

### Load-Balancing Scheme

---

The load-balancing method used by the adaptor for instances of the application. The options provided by WebObjects are Round Robin, Random, and Load Average. You can also use a custom load balancer by choosing Custom in the

### Deployment Settings Reference

pop-up menu under “Load balancing scheme” in the Load Balancing and Adaptor Settings section of the application configuration page in Monitor and entering the load balancer’s name in the Custom Scheduler Name text field. See “[Load Balancing and Adaptor Settings](#)” (page 82) for more information.

---

#### Receive Buffer Size

The size, in bytes, of the TCP socket receive buffer that’s used for adaptor-to-instance communication.

---

#### Receive Timeout

The length of time, in seconds, the adaptor waits for a response from an instance of the application before giving up.

---

#### Redirection URL

The URL that the user is redirected to when an instance fails to respond to a direct request.

---

#### Retries

The number of times a request is retried (trying several instances) if a communications failure occurs before an error page is returned to the HTTP server.

---

#### Send Buffer Size

The size, in bytes, of the TCP socket send buffer that’s used for adaptor-to-instance communication.

---

#### Send Timeout

The length of time, in seconds, the adaptor attempts to send data to an instance of the application before giving up.

---

#### URL Version

The WebObjects version to use for URL parsing and formatting. All WebObjects 4, 4.5, and 5 applications use version 4 URLs by default.

## Scheduling Settings

---

These properties determine when Monitor restarts application instances. Restarting instances regularly helps you deploy highly reliable sites. You set the values of scheduling properties in the Scheduling section of the application configuration page.

### Is Scheduled

---

Determines whether the schedule defined for a particular instance is active.

### Graceful Scheduling

---

Determines whether an instance is shut down gracefully or immediately at its scheduled shut-down time. During a graceful shutdown, the instance does not create sessions for new users (they are automatically directed to other available instances, if any). Existing sessions remain active until they time out or are destroyed programmatically. When the number of active sessions drops to the value set for the Minimum Active Sessions property, the instance is restarted. When Graceful Scheduling is not selected for the instance, it is restarted immediately, terminating active sessions. See “[Minimum Active Sessions](#)” (page 113) for more information.

### Types of Schedule

---

There are three types of schedule available for instances: hourly, daily, and weekly.

- **Hourly:** The instance is restarted after a certain number of hours from a particular hour.
- **Daily:** The instance is restarted at a particular hour every day.
- **Weekly:** The instance is restarted a particular hour on a specific day of the week.

## Command-Line Arguments

---

The following sections describe the command-line options available to fine-tune your deployment. You specify command-line options using the following format:

## Deployment Settings Reference

```
<appName> -<optionName> <value> -<optionName> <value> ...
```

There are three types of command-line options: general, wotaskd and Monitor specific, and wotaskd specific.

## General Command-Line Arguments

---

The sections below describe the command-line options that apply to Monitor processes, wotaskd processes, and application instances.

### WOApplicationName

---

**Value format:** *string*.

**Default value:** Name of executable file.

**Description:** This setting is used by the HTTP adaptor to implement load balancing. The adaptor can load-balance only between instances with the same application name. This setting can be used to create groups of instances, even when the instances share the same executable file. This argument is set automatically for instances started by wotaskd.

### WODirectConnectEnabled

---

**Value format:** *string*.

**Default value:** `true` in development, `false` in deployment.

**Description:** Specifies whether a user can connect to an instance using direct connect, that is, by specifying a host name and a port number.

### WOHost

---

**Value format:** *string*.

**Default value:** None.

**Description:** Specifies the network interface that an instance binds to. This argument should only be used on hosts with multiple network interfaces (IP addresses).

## WOLifebeatDestinationPort

---

**Value format:** *number*.

**Default value:** 1085.

**Description:** Determines the port that instances send lifebeats to (should be the port that the wotaskd process that overlooks the instance runs on). The value of this property on an application instance must be the same as the value of the property for the Monitor process used to manage the instance. The wotaskd process sets this argument to the value of its `WOPort` for instances that it starts.

**See:** “`WOPort`” (page 119).

## WOLifebeatEnabled

---

**Value format:** {true, false}.

**Default value:** true.

**Description:** Determines whether the instance sends lifebeats.

## WOLifebeatInterval

---

**Value format:** *number*.

**Default value:** 30.

**Description:** Determines the interval, in seconds, between lifebeats.

## WONoPause

---

**Value format:** {true, false}.

**Default value:** false.

**Description:** In Windows 2000 launch scripts, if an error occurs during script execution, the process idles because the message “Press any key...” is displayed on the console and the system awaits a keypress. To avoid this behavior, set this property to true.

## WOOOutputPath

---

**Value format:** *path*.

Deployment Settings Reference

**Default value:** /dev/null.

**Description:** Allows you to redirect the instance's standard output and standard error streams to the directory specified. The file generated is named as follows:

<application name>-<instance ID>

## WOPort

---

**Value format:** *integer*.

**Default value:** 1085.

**Description:** The port that the instance runs on.

**See:** "WOLifebeatDestinationPort" (page 118).

## WOREcordingPath

---

**Value format:** *path*.

**Default value:** None.

**Description:** Specifies the path for the file that stores each request and response made in a session.

## WOSessionTimeOut

---

**Value format:** *number*.

**Default value:** 3600.

**Description:** Specifies the time, in seconds, that passes after the last request is processed before the session times out.

## WOStatisticsPassword

---

**Value format:** *text*.

**Default value:** None.

**Description:** Specifies the password that must be entered to gain access to the instance statistics (WOSTats) page of an application instance.

## Monitor and wotaskd Command-Line Arguments

---

There's one command-line option that applies to both Monitor and wotaskd processes but not to application instances.

### WODeploymentConfigurationDirectory

---

**Value format:** *path*.

**Default value:** /Library/WebObjects/Configuration.

**Description:** Each wotaskd process writes its configuration to a file called `SiteConfig.xml` in the directory specified here. (The HTTP adaptor configuration file is also written to this directory.) This argument, in conjunction with `WOLifebeatDestinationPort`, allows you to run multiple wotaskd processes on a single computer.

**See:** “`WOLifebeatDestinationPort`” (page 118).

## wotaskd Command-Line Arguments

---

These are command-line arguments that apply only to wotaskd processes.

### WOAssumeApplicationIsDeadMultiplier

---

**Value format:** *integer*.

**Default value:** 4.

**Description:** Used to determine the number of seconds that the wotaskd process waits without receiving a status message from an instance before considering it dead. It species a multiplier against `WOLifebeatInterval`. If `WOLifebeatInterval` is 30, a wotaskd process waits 120 seconds from the last status message before determining that an instance is dead.

### WOMulticastAddress

---

**Value format:** *ip-address*.

**Default value:** 239.128.14.2.

**Description:** Sets the IP address that the wotaskd process listens to for multicast requests from the HTTP adaptor.



### WORespondsToMulticastQuery

---

**Value format:** {true, false}.

**Default value:** false.

**Description:** Determines whether the wotaskd process responds to multicast queries from the HTTP adaptor.

### WOSavesAdaptorConfiguration

---

**Value format:** {true, false}.

**Default value:** false.

**Description:** Determines whether the wotaskd process generates an HTTP adaptor configuration file.

## C H A P T E R 8

### Deployment Settings Reference

# Special Deployment Issues

---

The following sections explain special issues to address when deploying WebObjects applications.

## Deployment Issues With Java Client Applications

---

A WebObjects application developer can produce applications of two types:

- **HTML-based applications** on which the user interface elements are produced using HTML code.
- **Java Client applications**, which use Sun's Swing technology to produce a user interface that is more appealing and more efficient than HTML-based interfaces. For more information on Java Client applications see the WebObjects Java Client documentation, available at [www.apple.com/developer](http://www.apple.com/developer).

There are two main issues that you should keep in mind when you deploy and administer Java Client applications:

- **session timeout**

Java Client applications offer the user an interface that is very similar to the one offered by regular desktop applications. Therefore, they expect Java Client applications to behave in a way similar to their desktop applications.

One of the main differences between a desktop application and a Java Client application is that Java Client applications open a connection to a server-side application. This connection expires after a certain period of inactivity. By default, the timeout period is 30 minutes. This may not be enough time for an application user that launches the application, goes to lunch, and returns to

## Special Deployment Issues

work 45 minutes later. When the user tries to use the application (which is still running on her computer), she will see a dialog that indicates that her session has timed out. In addition, any changes that were not saved are lost.

**■ traffic level**

In HTML-based applications, the packets sent between the Web browser and the HTTP server tend to be large. However, the size of the packets doesn't vary much (the server always sends the entire page to the browser). In Java Client applications, the packets sent by the server during application startup can be large (the entire application or part of it is downloaded); subsequent packets are relatively small (user-entered data and search results, for example).

For more information on WebObjects's Java Client technology, refer to *Inside WebObjects: WebObjects Desktop Applications*.

## Deploying WebObjects 4.5.1 and WebObjects 5 Applications

---

If you want to deploy WebObjects 4.5.1 and WebObjects 5 or WebObjects 5.1 applications on the same computer, you must

- install WebObjects 4.5.1 Deployment and either WebObjects 5 or WebObjects 5.1 Deployment on the computer
- use either the WebObjects 5, or 5.1 versions of Monitor and wotaskd

Note that you cannot use the 4.5.1 version of wotaskd together with the 5 or 5.1 version of Monitor. Furthermore, you cannot deploy 4.5.1, 5, and 5.1 applications on the same computer.

# Document Revision History

---

This document has the revisions listed in [Table B-1](#).

**Table B-1** Document revision history

Date	Notes
Jan 2002	Added index.
Nov 2001	Changed conceptual art.
	Added information on <code>WODirectConnectEnabled</code> command-line argument.
	Changes references to <code>obj.conf</code> to <code>magnus.conf</code> .
	Added information on how and when the <code>SiteConfig.xml</code> file is written.
	Changed application settings screen shot to reflect the addition of the Project Search Path, Session Timeout, and Statistics Page Password text input fields.

## A P P E N D I X B

### Document Revision History

# Glossary

---

**API-based adaptor** HTTP adaptor based on APIs specific to a particular HTTP server. It allows CGI-like tasks to run as part of the main server process, avoiding the creation and termination of a process for each request. See **API-based adaptor**; **CGI adaptor**; **HTTP adaptor**.

**application host** A computer capable of running application instances.

**bundle** In Mac OS X systems, a bundle is a directory in the file system that stores executable code and the software resources related to that code. The bundle directory, in essence, groups a set of resources in a discrete package.

**CGI (Common Gateway Interface)** A standard for communication between external applications and information servers, such as HTTP servers.

**CGI adaptor** HTTP adaptor that uses the Common Gateway Interface (CGI) to translate requests from an HTTP server into requests to an application instance, and responses from an application instance to responses to the HTTP server. The HTTP server creates a new CGI process to handle each request. See **API-based adaptor**; **HTTP adaptor**.

**component** An object (of the `WOComponent` class) that represents a Web page or a reusable portion of one.

**datasource adaptor** A mechanism that connects your application to a particular database server. For each type of server you use, you need a separate adaptor. WebObjects provides an adaptor for databases conforming to JDBC. See also **JDBC adaptor**.

**framework** A type of bundle that packages a dynamic shared library with the resources that the library requires, including header files and reference documentation.

**HTTP** The client-server TCP/IP protocol used on the Web for the exchange of HTML documents.

**HTTP adaptor** A process (or a part of one) that connects WebObjects applications to an HTTP server. See also **HTTP server**.

**Java Client** A WebObjects development approach that allows you to create graphical user interface applications that run on the user's computer and communicate with a WebObjects server application.

**JDBC (Java Database Connectivity)** An interface between Java platforms and databases.

**JDBC adaptor** A datasource adaptor that allows WebObjects applications to connect to JDBC-compliant database management systems. See **datasource adaptor**.

**heartbeat** Status message sent by WebObjects applications to wotaskd to report their activity. The four types of heartbeat messages are has started, is alive, will stop, and will crash.

**load balancing** Technique used to distribute user-load among the instances of an application. When multiple instances of an application are running and a new user accesses the application, the WebObjects adaptor uses one of several algorithms to determine which instance to forward the request to.

**loopback** Mechanism that allows you to open a connection to a computer that does not go over the network.

**Monitor** A tool used to configure and maintain deployed WebObjects applications capable of handling multiple applications, application instances, and applications hosts at the same time.

**Project Builder** A tool used to manage the development of a WebObjects application or framework.

**session** A period during which access to a WebObjects application and its resources is granted to a particular client (typically a browser). Also an object (of the WOSession class) representing a session.

**SMTP (Simple Mail Transfer Protocol)** A protocol used to transfer email between computers, usually over Ethernet.

**socket** Mechanism for creating a virtual connection between processes. It interfaces standard I/O with network communication facilities. A socket address consists of a port number and an IP address.

**UDP (User Datagram Protocol)** Lightweight and efficient connectionless datagram transport protocol. Used to send self-routing data throughout a network.

**HTTP server** An application that serves Web pages to Web browsers using the HTTP protocol. In WebObjects, the HTTP server lies between the browser and a WebObjects application. When the HTTP server receives a request from a browser, it passes the request to the WebObjects adaptor, which generates a response and returns it to the HTTP server. The HTTP server then sends the response to the browser. See also **HTTP adaptor**.

**WebObjects Deployment** Software package that allows you to deploy WebObjects applications on an intranet or the Web. You need to install a WebObjects deployment license on computers on which you want to install this package. See also **WebObjects Development**.

**WebObjects Development** Software package that allows you to develop WebObjects applications. It includes tools to design applications using an object-oriented approach. You need to install a WebObjects development license on computers on which you want to develop applications. See also **WebObjects Deployment**.



## G L O S S A R Y

**WOServices** WebObjects service that monitors wotaskd processes. Its main duty is to monitor wotaskd and restart it if it dies or when the host is restarted. The implementation of this service is platform-dependent.

**wotaskd (WebObjects task daemon)** WebObjects Deployment tool that manages the instances on an application host. It's used by Monitor to propagate site configuration changes throughout the site's application hosts. See also **Monitor**.

# G L O S S A R Y

## Glossary

# Index

---

## A

---

- access, restricted 26, 27, 50
- Adaptor setting 110
- Adaptor Threads setting 110
- adaptors
  - JDBC 18
- Additional Arguments setting 112
- `additionalArgs` property 47
- Apache HTTP adaptor 33
- API-based HTTP adaptor 107
- Apple WebObjects Monitor service 65
- Apple WebObjects Task Daemon service 65
- application configuration page 76–83
  - application properties 110
  - definition 76
  - HTTP-adaptor properties 82, 83, 114
  - instance properties 111
  - load balancing 82, 83
  - load-balancing scheme 115
  - scheduling instances 77, 80
  - setting application properties 77, 80
  - setting email notifications 77, 81
  - setting new-instance defaults 77, 78
  - setting the load-balancing scheme 77
- application detail page illustrated 85, 87, 101
- application files 26, 74
- application hosts
  - adding 68
  - configuring 70
  - definition 68
  - discovering 52
  - hardware 107
  - HTTP adaptors 37
  - load balancing between 18
  - Monitor 22, 57, 68
  - performance 20, 97
  - state of 72
  - viewing the configuration of 71
  - wotaskd 23, 61
- application instances 18–19
  - adding 83–88
  - auto-recovering 88, 107
  - configuring 88, 100
  - connecting to 91, 100, 101
  - death of 61, 102, 120
  - debugging 112
  - deleting 88
  - development 27
  - failing to respond 115
  - groups of 46, 111
  - HTTP adaptors 36
  - instance IDs in 73
  - launch script 113
  - lifebeats 60, 118
  - management of 30
  - Monitor 22, 57
  - performance 107
  - properties 90
  - Receive Timeout setting 115
  - refusing new sessions 88
  - refusing requests 110
  - registered 73
  - running 31, 100
  - scheduling 88, 116
  - send timeout 115
  - session timeout 123
  - settings 111–114
  - starting and stopping 88, 111, 112
  - state information 108
  - statistics 27, 114, 119
  - status 88
  - unregistered 73
  - unresponsive 60, 112
  - `WOLifebeatDestinationPort` property 118
- application performance 20
  - application detail page
    - Active Sessions 102

## INDEX

### application performance (*continued*)

- Average Idle Period 102
- Average Transaction 102
- Deaths 102
- Transactions 102
- WOSTats 102

- application hosts 20
- Listen Queue Size 110
- logs 105
- monitoring 99–104

### applications

- adding to a site 60, 76
- analyzing activity 104
- configuring 76
- deleting 100
- installing 74
- load-balancing scheme 97
- logging activity 104
- Monitor 57
- multiple environments 97–98
  - defined 97
  - illustrated 98
- naming 46, 101, 111
- properties
  - application configuration page 110
  - defined 110
- setting up 75
- settings 110–111
- traffic level 110, 124

### Applications page

- application performance 99
- illustrated 76, 84, 100

Auto Recover setting 112

Auto-Open In Browser setting 112

## B

---

browsers, Web 18

buffers 115

## C

---

Caching Enabled setting 112

CGI HTTP adaptor  
performance 107

CGI HTTP adaptors 33, 38

cgi-bin in URL, setting alias for 55

clients 17

`cnctTimeout` property 47

Common Log File Format (CLFF) 104

communications failure 115

component-definition caching 107, 112

configuration files, HTTP adaptor

HTTP adaptors 22

properties in 46

static application-site configuration using 38

structure of 45

wotaskd configuration page 73

Connect Timeout setting 114

Connection Pool Size setting 114

control path 21–23

## D

---

Daily schedule 116

data path 21–23

data stores 18

data-source adaptors 18

debugging applications 112

Debugging Enabled setting 112

deployment, parts of 17

deployment-configuration directory 58

development instances 73

direct connect 27, 117

`Document Root`

directory 74

setting the path 56

`dormant` property 46

Dormant setting 114

## INDEX

### E, F

---

Email Notification settings 114  
email notifications  
    application configuration page 77  
    Monitor 25  
    site configuration page 92  
email notifications application configuration page 81  
error page 115

### G

---

Graceful Scheduling setting 116

### H

---

hardware, application host 107  
Host configuration information page 72  
host property 47  
Hosts page 69  
Hourly schedule 116  
HR application 44  
HTTP adaptors 35–56  
    analyzing activity 105  
    Apache 33  
    API-based 39, 107  
    application configuration page 82  
    application hosts 40  
    application instances 38, 40  
    building 30  
    CGI 33, 38  
    configuration refresh interval  
        host lists 44, 53  
        multicast request 41, 52  
    connecting to application instances 101  
    customizing 52  
    definition 17  
    deployment tool 25  
    dynamically, obtaining state 38  
    executable files 30

files of 31  
host lists  
    application hosts, fixed 42  
    setting 38, 53  
    site configuration method 40  
host polling 44  
installing 30  
instance configuration page 90  
ISAPI 33  
logging 105  
logging activity 105  
multicast requests 38, 40, 41, 120, 121  
performance 105, 107  
platform, default installation by 33  
settings 92  
site configuration page 93  
site state 39  
source files 30  
static configuration file  
    deployment-configuration directory 120  
    naming the 54  
    obtaining state from a 38, 40, 50  
static configuration file, writing the 40  
URL 91, 101  
WebObjects Adaptor Information page 54  
worker threads 110  
wotaskd 22  
HTTP servers  
    CGI adaptors 38  
    configuration files, HTTP adaptor 44  
    deployment model 17, 19  
    development instances 73  
    logging HTTP-adaptor activity 105  
    multiple-computer site 37  
    performance 107  
    running multiple 97  
HTTP-adaptor  
    properties 114–115  
HTTP-server resources 26, 74

## INDEX

### I

---

id property 47  
ID setting 112  
installation, software 29–33  
    application hosts 31  
    HTTP servers 31  
    tools, deployment 31  
instance configuration page 88  
    displayed 89  
    instance properties 111  
instance statistics page 102–104  
    accessing 102  
    illustrated 103  
    password-protecting 90, 91, 102  
    Statistics Page Password property 114  
    WOLStatisticsPassword argument 119  
Is Scheduled setting 116  
ISAPI HTTP adaptors 33

### J, K

---

J2SE 32  
JDBC adaptors 18  
JDBC drivers 18

### L

---

license 30  
Lifebeat Interval setting 112  
lifebeats  
    definition 60  
    multiple-site deployments 97  
    WOLifebeatDestinationPort property 118  
Listen Queue Size setting 110  
load balancing  
    application detail page 101  
    configuring applications 77, 82  
    definition 96  
    deployment model 18  
    HTTP-adaptor configuration file 44

HTTP-adaptor log file 106  
    name property 46  
Name setting 111  
properties 114–115  
scheme  
    application level 97  
    custom 114  
    default 97  
    HTTP-adaptor configuration file 46  
    Load Average 97  
    Monitor 23  
    Random 97  
    Round Robin 97  
    site level 97  
    site configuration page 92  
Load-Balancing Scheme setting 114  
load-balancing, ID setting in 112  
logging HTTP-adaptor activity 105  
logs  
    adaptor activity 105  
    application activity 104  
    creating HTTP-adaptor log file 105  
    HTTP headers 106  
    WebObjects.log file 105  
logWebObjects file 105

### M

---

Maximum Adaptor Threads setting 110  
Minimum Active Sessions setting 113  
Minimum Adaptor Threads setting 111  
Monitor  
    access to 61, 95  
    application-instance deaths 102  
    communication paths 22  
    configuration files, HTTP adaptor 44, 47  
    deployment tool 23, 25  
    Minimum Active Sessions setting 113  
    preferences 94–96  
    properties 120  
    refresh settings for the application detail page 96  
    scheduling application instances 107, 116

## INDEX

### Monitor (continued)

- SiteConfig.xml file 59
  - sites, configuring 57
  - SMTP server 81
  - starting 62
  - Time Allowed for Startup setting 111
  - user privileges 58
  - WOLifebeatDestinationPort property 118
- Monitor state information 23
- multicast channel
- multicast request 41, 120
  - setting the IP address 52
  - setting the port 52

## N

---

- name property 46
- Name setting 111
- naming applications 46, 101, 111
- network interfaces, multiple 47, 117
- NEXT\_ROOT environment variable 31, 32
- nonsensitive resources 26, 74

## O

---

- operating system performance 107
- operating systems supported 15
- Output Path setting 113

## P, Q

---

- packets, TCP and UDP 110
- password-protect (Monitor and wotaskd) 26, 61, 95
- Path setting 113
- Payroll application 44, 106
- pbxbuild tool 74
- performance
  - application-instance state information 108
  - auto-recovering application instances 107

- CGI HTTP adaptor 107
  - component-definition caching 107
  - data-store server 108
  - HTTP adaptors 107
  - HTTP servers 107
  - operating system 107
  - physical system configuration 107
  - restarting application instances 107
  - session timeout 108
  - static content 108
- Phased Startup setting 111
- poolSize property 46
  - port property 47
  - Port setting 113
  - preferences page
    - illustrated 94
    - Monitor 94
    - password 95
  - Project Builder application 74
  - Project Search Path setting 113
  - properties
    - additionalArgs 47
    - cnctTimeout 47
    - dormant 46
    - id 47
    - name 46
    - poolSize 46
    - port 47
    - recvBufSize 47
    - recvTimeout 47
    - redir 46
    - retries 46
    - scheduler 46
    - sendBufSize 47
    - sendTimeout 47
    - urlVersion 46
    - WOApplicationName 117
    - WOAssumeApplicationIsDeadMultiplier 120
    - WODeploymentConfigurationDirectory 120
    - WODirectConnectEnabled 73, 117
    - WOHost 117
    - WOLifebeatDesinationPort 60
    - WOLifebeatDestinationPort 118
    - WOLifebeatEnabled 118
    - WOLifebeatInterval 118

## INDEX

### properties (*continued*)

- WOMulticastAddress 120
- WONoPause 118
- WOOutputPath 118
- WOPort 119
- WORecordingPath 119
- WORespondsToMulticastQuery 121
- WOSavesAdaptorConfiguration 50
- WOSessionTimeOut 119
- WOStatisticsPassword 90, 119

## R

---

- rapid-turnaround mode 113
- Receive Buffer Size setting 115
- Receive Timeout setting 115
- recvBufSize property 47
- recvTimeout property 47
- redir property 46
- Redirection URL setting 115
- request processing 17–23
- response pages 19
- retries property 46
- Retries setting 115
- root user 74
- runtime environment 30

## S

---

- scheduler property 46
- Scheduling Type setting 116
- scheduling, application-instance
  - configuring applications 77, 80, 88
  - custom algorithms 107
  - daily 116
  - hourly 116
  - Monitor 23, 57, 116
  - settings 116
  - Weekly 116
- scripts directory 33
- security 26, 73, 95

- Send Buffer Size setting 115
- Send Timeout setting 115
- sendBufSize property 47
- sendTimeout property 47
- services 62–65
  - Apple WebObjects Monitor 65
  - Apple WebObjects Task Daemon 65
- session timeout in Java Client applications 123
- Session Timeout setting 113
- settings
  - application 110–111
    - Adaptor 110
    - Adaptor Threads 110
    - Listen Queue Size 110
    - Maximum Adaptor Threads 110
    - Minimum Adaptor Threads 111
    - Name 111
    - Phased Startup 111
    - Starting Port 111
    - Time Allowed for Startup 111
  - Debugging Enabled 112
  - HTTP-adaptor
    - Connect Timeout 114
    - Connection Pool Size 114
    - Dormant 114
    - Load-Balancing Scheme 114
    - Receive Buffer Size 115
    - Receive Timeout 115
    - Redirection URL 115
    - Retries 115
    - Send Buffer Size 115
    - Send Timeout 115
    - URL Version 115
- instance 111–114
  - Additional Arguments 112
  - Auto Recover 112
  - Auto-Open In Browser 112
  - Caching Enabled 112
  - ID 112
  - Lifebeat Interval 112
  - Minimum Active Sessions 113
  - Output Path 113
  - Path 113
  - Port 113
  - Project Search Path 113



## INDEX

- instance (*continued*)
  - Session Timeout 113
  - Statistics Page Password 114
  - Time Allowed for Startup 111
- scheduling
  - Graceful Scheduling 116
  - Is Scheduled 116
  - Scheduling Type 116
- site configuration page
  - application instances 91
  - email notifications 92
  - HTTP-adaptor settings 91, 114
  - illustrated 93
  - Monitor 91
- SiteConfig.xml file 58, 72, 120
- sites, application
  - administration 30
  - configuration 91–93
    - copy-and-paste method 47
    - Monitor 30
    - static 44
  - configuration page 93
  - deploying multiple 97
  - HTTP adaptors 37
  - load-balancing scheme 97
  - multiple 97–98
  - protecting 95
  - state in 38, 60
- sockets
  - definition 60
  - TCP/IP 47, 60, 115
  - UDP 42, 60
- split installation 74
- split-installation 26
- Starting Port setting 111
- Statistics Page Password setting 114
- streams
  - standard error 112, 113, 119
  - standard output 113, 119

## T

---

- Time Allowed for Startup setting 111
- tools, deployment 30
- traffic level 110, 124

## U, V

---

- URL
  - parsing and formatting 46
  - redirection 115
- URL Version setting 115
- urlVersion property 46
- user load 18

## W, X, Y, Z

---

- WebObjects Adaptor Information page 50, 54
- WebObjects script file 62
- WebObjects.log file 105
- Weekly schedule 116
- Windows 2000 118
- WOAdaptor class 110
- woadaptor.dtd file 45
- WOApplicationName property 117
- WOAssumeApplicationIsDeadMultiplier property 120
- WOConfig.xml file 44, 48, 58
- WODefaultAdaptor class 110, 111
- WODeploymentConfigurationDirectory property 120
- WODirectConnectEnabled property 73, 117
- WOHost property 117
- WOLifebeatDesinationPort property 60
- WOLifebeatDestinationPort property 97, 118, 120
- WOLifebeatEnabled property 118
- WOLifebeatInterval property 118, 120
- WOMulticastAddress property 120
- WONoPause property 118
- WOOutputPath property 118

## INDEX

- WOPort property 97, 118, 119
- WOWRecordingPath property 119
- WORespondsToMulticastQuery property 121
- worker threads 110, 111
- WOSavesAdaptorConfiguration property 50
  - properties
    - WOSavesAdaptorConfiguration 121
- WOServices 31, 50
- WOServices script file 64
- WOSessionTimeout property 119
- WOStatisticsPassword property 90, 119
- WOStats. See instance statistics page
- wotaskd 68, 73
  - application instances 61, 72, 118
  - configuration files, HTTP adaptor 50
  - configuration page 32
  - configuring 61
  - death of 62
  - deployment tool 23
  - deployment-configuration directory 120
  - environment of 73
  - multicast requests 41, 42
  - multiple-site deployments 97
  - Phased Startup setting 111
  - properties 120–121
  - protecting 95
  - running 31
  - SiteConfig.xml file and 59
  - starting 62
  - user privileges 58