WebAccelerator Getting Started

version 5.1

MAN-0209-00

Service and Support Information

Product Version

This manual applies to product version 5.1 of the WebAccelerator™.

Legal Notices

Copyright

Trademarks

Patents

This producted by U.S. Patents 6,327,242; 6,505,230; 6,640,240; 6,772,203. Other patents pending.

Table of Contents

Preface viii
Conventions Used in this Book ix

Chapter 1 Release Information 1
F5 WebAccelerator 2
Clusters And Applications 3
F5 WebAccelerator Remote 4
Installation Requirements 6
Disk Partitioning Requirements 6
Software Requirements 6
Network Requirements 7
Changes and Enhancements 7
Version 5.1 7
Improved User Interface 7
Clusters 7
Applications 7
Performance Reports 8
Enhanced Response Processing 8
Application Matching Extended 8
F5 WebAccelerator Remote 9
Simplified Installation with Appliances 9
Corporate Name Changes 9
F5 WebAccelerator Product Documentation 10

Chapter 2 Planning Your Installation 11
Installation Overview 12
Installation Planning Worksheets 12
WebAccelerator Machine 12
Accelerator Machines 15
Standby Management Console Machines 16

Chapter 3 Installing Your WebAccelerator 17
Installation Process 18
Post-Installation Steps 20
Meeting Minimum Software Requirements 21

Uninstalling WebAccelerator Software 22
Installing a Standby Management Console 23
Installation Process 23

Chapter 4 Configuring Your WebAccelerator 25
The Admin Tool 26
Host Mapping 26
Valid Hostnames 27
Defining a Destination Host 28
Protocol Security 28
Timeout Properties 29
Express Connect Options 29
Unmapped Requests 30
Configuration Process 30
Changing Your Configuration 38
Changing Your Preferences 39
Clearing Your Cache 39
Monitoring the WebAccelerator 40
Log File Rotation 41
Finding the MAC Address 42
Advanced Configuration 42

Chapter 5 Installing Accelerators 43
Installation Process 44
Uninstalling Accelerators 46

Chapter 6 Upgrading Your WebAccelerator 47
Upgrade Restrictions 48
Upgrade Process 48

Chapter 7 Confirming Your Installation 51
Check Processes 52
Check Proxy 52

Chapter 8 Configuring Your Network 55
Network Changes 56
Defining WebAccelerator Hostnames 56
Routing Requests to the WebAccelerator 57
Accessing your Origin Servers 57
Routing Requests to the WebAccelerator Remote 58
123 Corporation Usage Scenario 59
Configuring Your Firewalls 62
WebAccelerator Ports 62
WebAccelerator Ports 63

Appendix A Linux-Specific Information 65
Linux Installation Planning 66
Operating System Installation 66
Setting up Red Hat up2date 67
Managing CD-ROMs 68
Mounting CD-ROMs: Method 1 68
Mounting CD-ROMs: Method 2 69
Mounting CD-ROMs: Method 3 69
Unmounting CD-ROMs 69

Preface

The F5 WebAccelerator is a distributed system that is designed to improve your site's performance while off-loading traffic from your site's origin servers.

This guide describes how to install, configure, and manage your F5 WebAccelerator software. In particular, this guide provides information on:

- core concepts and features of interest to the WebAccelerator administrator
- pre-install planning, including information on platform requirements such as hardware and OS versions, firewall configuration, and DNS requirements
- installation if you purchased a software license for the F5 WebAccelerator and are installing it on your own hardware
- configuration procedures
- log file management

Conventions Used in this Book

This section explains the conventions used in this book.

`Monospaced font` – This font is used for examples, text that appears on the screen, command line utility names, and filenames.

`<bracketed text>` or *`italic text`* represents elements in a path or example that are intended to be replaced with information specific to your installation or procedural requirements.

Text of this color indicates a link in PDF or HTML that you can click on to navigate to a related section.

Note: Notes mark important information. Make sure you read this information before continuing with the task.

Technical Support Information

Phone	(+1) 206-272-6888
Fax	(+1) 206-272-6802
Web	http://tech.f5.com
FTP	ftp.f5.com
Email	support@f5.com

Chapter 1

Release Information

F5 WebAccelerator
Installation Requirements
Changes and Enhancements
F5 WebAccelerator Product Documentation

The F5 WebAccelerator provides technologies to accelerate your application's response to HTTP requests. The WebAccelerator must be installed on your corporate network between your applications' users and the server or servers on which your application runs.

Each WebAccelerator component, the Management Console, the WebAccelerator, and the WebAccelerator Remote, can be purchased through a software license or as an appliance. If you purchase appliances, the software is pre-installed and you simply set up the box, following instructions that come with the hardware. You do not need to read the chapters of this guide that deal with software installation. Instead, focus on these sections:

- "F5 WebAccelerator" on page 2
- "Changes and Enhancements" on page 7
- "F5 WebAccelerator Product Documentation" on page 10
- Chapter 4, Configuring Your WebAccelerator
- "Check Proxy" on page 52
- Chapter 8, Configuring Your Network

If you purchase the software license option, you provide the hardware and install the software yourself, following the instructions in this guide. This option is useful if your enterprise has certain hardware or purchasing requirements. You might find it easier to purchase a machine locally or re-purpose a machine and download the software from the Internet, rather than have an appliance shipped to the location where you want to install it.

F5 WebAccelerator

The WebAccelerator is a software package that comes pre-installed on one or more appliances running Redhat Linux. In its most basic configuration, both parts of the WebAccelerator are on the same machine:

- Management Console

  The Management Console controls and manages the WebAccelerator, and provides internal communications between the WebAccelerators and other WebAccelerator processes. The WebAccelerators receive most of their configuration from the Management Console. Every WebAccelerator installation has one active Management Console. You cannot have two Management Consoles actively running in a single WebAccelerator installation.

- the WebAccelerator

  The WebAccelerator performs the core content caching and serving functions, getting its caching rules from the Management Console.

This diagram illustrates this basic configuration:

Figure 1 Single Machine Configuration

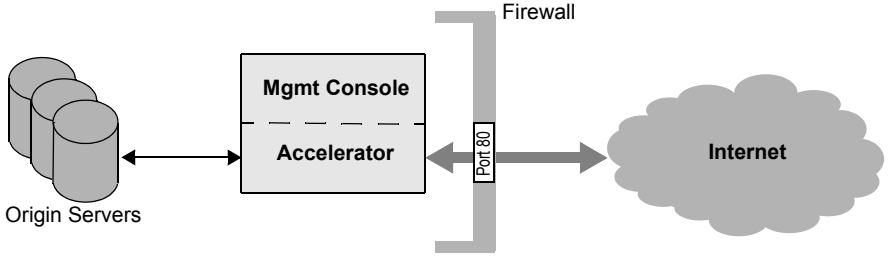

In this diagram, you can see the incoming traffic to the WebAccelerator is heavier than the traffic between the WebAccelerator and the origin servers. When the WebAccelerator first starts up, the WebAccelerator proxies every request to the origin servers, and caches the responses. Over time, more and more requests can be serviced from the WebAccelerator's cache. The load on the origin servers decreases as the WebAccelerator handles the requests and only proxies to the origin servers for first requests for particular content or to refresh content.

You might have multiple origin servers, each described to the WebAccelerator in the host map. The host map is created when you first configure the WebAccelerator. Basically, you list the hosts you expect to see on incoming HTTP requests, the domain names that would be in the Host header. You map each of these hosts to the domain name of the origin server that normally handles the request. Depending on how your site is configured, you might map all incoming host names to a single web server, or different host names might map to different web or application servers. When the

WebAccelerator receives a request that it cannot service from cache, it looks at the Host header information, compares that to the host map, and proxies the request to appropriate domain name.

Most configurations have two or more WebAccelerators. For a distributed configuration that can handle larger loads, you need multiple WebAccelerators. Two WebAccelerators are the minimum F5 recommends, ensuring your system has some fault tolerance.

With multiple WebAccelerators, you have one WebAccelerator and a Management Console together on an appliance, and the other WebAccelerators each as their own appliance. The Management Console controls all the WebAccelerators, providing information such as the host map to them. This multi-machine configuration requires a load balancer or some other mechanism, such as DNS round robin, for allocating incoming requests to the WebAccelerators.

In this figure, the Accelerators are grouped into one cluster, which handles the requests for all the origin servers:

Figure 2 Basic Multiple WebAccelerator Configuration

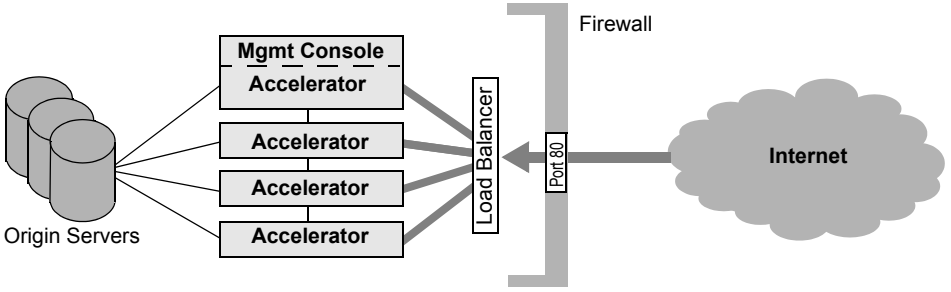

Clusters And Applications

Accelerators are grouped into clusters, and the clusters are assigned to your web applications to handle the requests directed to them. Clusters make it easier to manage multiple Accelerators. All the Accelerators handling the requests for one application can be managed as a group, with the same characteristics and configuration, instead of configuring each Accelerator individually. An application can only have a single cluster assigned to it, but a cluster can contain as many Accelerators as needed.

For each application that you want to manage individually, create a profile specifying:

- what policies are used to handle requests to that application
- the host map to be used for the application
- which cluster is assigned to it

For example, you might have three applications that you want the WebAccelerator to accelerate: the main external web server, an app server, and an internal web server. By setting up a separate application profile for each, and defining three clusters, each can have their own set of policies. This is useful if the characteristics of each or the traffic they get are very different and require different policies for optimal performance. This is also useful if the different servers are managed by different groups. This configuration allows each group to manage its policies independently.

Figure 3 Multiple Applications and Clusters

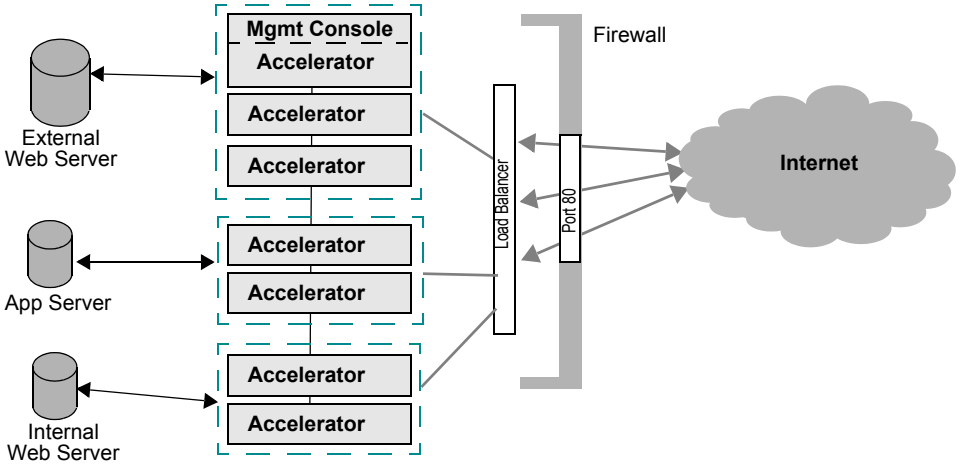

When an HTTP request arrives, the load balancer directs it to an Accelerator in the appropriate cluster, based on the address you specify for the cluster and the host name found on the request. The Accelerator looks at the host name, compares that to the host maps defined for the applications to identify which application profile to use, and applies the appropriate policy set to the request.

F5 WebAccelerator Remote

The F5 WebAccelerator Remote is designed to work in a remote location as an WebAccelerator that you locate on the receiving end of a response, near the clients. It works together with a standard WebAccelerator, and is configured through the Management Console of the WebAccelerator just like any other WebAccelerator, as part of a cluster. The WebAccelerator Remote has several benefits:

- because it is located near your remote users, responses to requests serviced from cache travel only a short distance, usually on the local network
- when a request cannot be serviced from cache, the WebAccelerator Remote proxies to a primary WebAccelerator. Because the primary WebAccelerator is

sending the response to another WebAccelerator instead of to a browser client, the sending WebAccelerator can perform additional processing and compression on the response stream. The WebAccelerator Remote translates the response back into something the browser understands. The additional processing and compression makes more efficient use of your network bandwidth and speeds up response times.

WebAccelerator Remote is ideal for remote offices that are located far from internal corporate websites or on networks with limited bandwidth. It is especially useful for remote offices that need to receive large amounts of data from your internal web applications.

In a basic WebAccelerator Remote configuration, a WebAccelerator Remote server is placed at the same location as the remote user group, which might be a geographically distant sales office or an offshore development group. The WebAccelerator Remote is configured to communicate with the primary WebAccelerator located closer to the data sources, for example, at corporate headquarters or the data center.

Figure 4 Basic WebAccelerator Remote Configuration

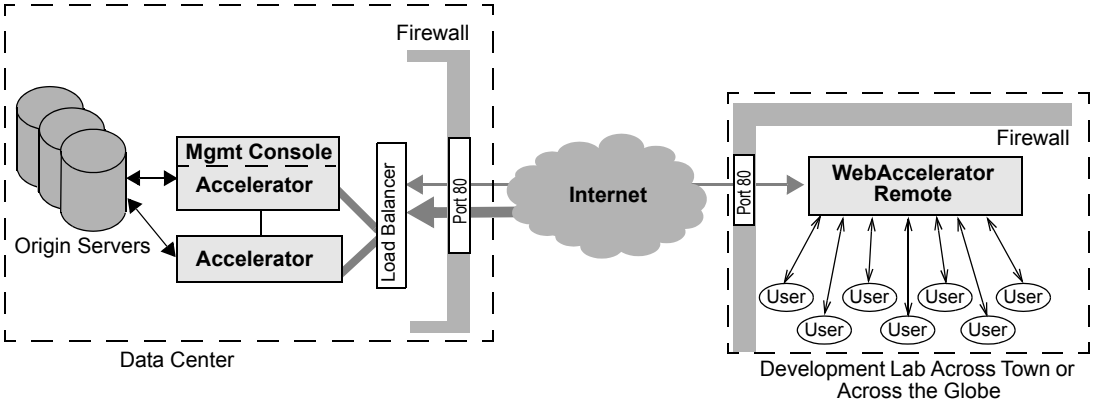

When a user at the remote site sends a request, it is routed to the WebAccelerator Remote, which either services the request from its cache or proxies it to the primary WebAccelerator. In turn, the primary WebAccelerator either services the request from its cache or proxies to the origin servers for content. The response is sent from the primary WebAccelerator to the WebAccelerator Remote, where it can be cached, and the WebAccelerator Remote sends the response on to the clients.

Installation Requirements

You only need to read this section if you purchased the WebAccelerator as a software license. Your site and the machines on which you plan to install WebAccelerator software must meet the requirements described in this section.

Disk Partitioning Requirements

You must have your disks partitioned so that there is enough space for WebAccelerator's files. The most basic partitioning, `/disk1` and `/disk2`, should allow for:

- on the first disk:
  - `/` 3.0 GB
  - `swap` 2.0 GB
  - `/var` 1.0 GB
  - `/disk1` all remaining disk space
- on the second disk: all disk space

Software Requirements

Your system must meet these requirements:

- Redhat Linux Enterprise ES or AS version 3 must already be installed on your system

  For recommendations on installing RedHat Linux to optimally support the WebAccelerator software, see Appendix A, Linux-Specific Information.
- Perl 5.6.0 or later, which can be installed as part of the WebAccelerator installation
- NTP 4.1.0, which can be installed as part of the WebAccelerator installation
- Sendmail 8.11.6, which can be installed as part of the WebAccelerator installation

Your WebAccelerator installation media includes an installer that places the Perl, NTP, and Sendmail packages on your system for you.

Internet Explorer 5.5 or higher is required to access all portions of the Administration Tool, also called the Admin Tool. Netscape 6.0 or higher can be used to access everything except the Policy Editor portion of the Admin Tool.

Network Requirements

Placing your WebAccelerator into production requires that you place it on your network between your users and the server or servers that run your web application. You must reconfigure DNS and your firewalls to support this new configuration.

This reconfiguration should occur after you have installed your WebAccelerator software, or after your SL400 appliance is up and running. See Chapter 8, Configuring Your Network for more information.

Changes and Enhancements

This section describes the changes and enhancements new to this release.

Version 5.1

Improved User Interface

The user interface for the WebAccelerator, the Admin Tool, has been reorganized and improved to make the task of managing the WebAccelerator easier. For this release, there is only one login to the Admin Tool, the administrator account. Many tasks, such as publishing policies and clearing the cache, have been simplified.

Previously, you could only use Internet Explorer 5.5 or higher to access the Admin Tool. In this release, most portions of the Admin Tool, with the exception of the Policy Editor, can also be accessed with Netscape 6.0 or higher.

Clusters

Accelerators, whether they are local WebAccelerators or WebAccelerator Remotes, can now be managed as groups called clusters, to make it easier to manage multiple Accelerators in large enterprise settings. New to this release, a system monitor performs health checks on the Accelerators in each cluster and displays the results in the Admin Tool Cluster table. For more information, see "Clusters And Applications" on page 3.

Applications

This release makes the concept of organizations obsolete. Instead, the WebAccelerator is configured around applications. You define the applications you want the WebAccelerator to accelerate, and create a profile for each that defines the policies, host maps, and Accelerators that should handle requests to the application. Any user

with access to the administrator login can change the configuration for an application. You no longer need to log out and log back in to modify the policies for different applications. For more information, see "Clusters And Applications" on page 3.

Performance Reports

The Performance Monitor is replaced in this release with an easier to use Performance Reports feature. You have more options for creating and viewing reports on performance and caching. See "Monitoring the WebAccelerator" on page 40.

Enhanced Response Processing

The response processing has been enhanced to better classify responses, and to improve compression and transform capabilities with streaming. Response processing now includes response object classification and content identity based on the actual content of the response that can be used to normalize a URL.

Response object classification enables you to classify a response by object type and group based on the extension or MIME type of the response. There are pre-defined object types that the WebAccelerator uses to determine whether to apply compression or rewrite scripts. You can modify these, and also create your own. For more information, see the *Administration Guide*.

The WebAccelerator can normalize URLs for systems that use changing URLs or different URLs for the same content. The content of the response is analysed and assigned an object identifier. The WebAccelerator can perform a redirect using a normalized URL based on the WebAccelerator object identifier. This enables the browser or WebAccelerator Remote to locate the object in cache even though the requesting URL was randomized. This identifier also enables the WebAccelerator to recognize the object even if the response passes through gateways that re-encode the URL.You can modify the default settings to control this feature and to enable authorization and authentication. For more information, see the *Administration Guide*.

Application Matching Extended

Application matching is now performed twice during processing, once against the request and once against the response if the request was proxied. The match of a request to a node in the Request Type Hierarchy determines the content variation, connection, invalidation trigger, and proxy policies to apply to the request. The match of the response to a node determines the content assembly, lifetime, responses cached, and cache-related portions of the proxy policies to apply to the response.

When you specify application matching policies, there is a new parameter you can specify called Content Type. When a request or response is classified by object type, the group and object type are appended to generate a content type

(*group.objectType*). For a request, the file extension is used to determine object type and generate the content type. For a response, the WebAccelerator looks at several things to determine object type, including MIME type, and a response might end up with a different content type than its request had.

In matching rules, you can specify that a request or response must match or not match a certain value for Content Type. Because the content type can be different for a response than for the request that originated it, the matching rules based on Content Type are what can cause a response to match a different node than its request.

F5 WebAccelerator Remote

This release supports the newest version of the WebAccelerator Remote. The WebAccelerator Remote previously was a lightweight WebAccelerator and was managed separately. In this release, the WebAccelerator Remote provides the functionality of a remote Accelerator and is managed through the Management Console of the primary WebAccelerator just like any other Accelerator.

The WebAccelerator Remote can be used to improve performance of web applications to users at a remote site, and can also accelerate applications located at the remote site.

Simplified Installation with Appliances

All WebAccelerator components are available as appliances. If you purchase them as appliances, you simply plug them in and configure them for your network. These are the configurations that are available on the SL400 appliance:

- the Management Console and an WebAccelerator, or
- a single WebAccelerator, or
- a WebAccelerator Remote

All WebAccelerator components continue to be available for purchase as a software license.

Corporate Name Changes

The WebAccelerator was previously produced by a company called Pivia, Incorporated. As a result, what used to be called the Pivia Performance Server is now called the F5 WebAccelerator. References to the company and product names have been changed throughout the documentation, tools, and products. However, where `pivia` is part of the inner workings of the products, it has not been changed. You may still see it, for example, appearing as a parameter, part of a command, imbedded in an automatically generated filename, or used as a system identifier.

F5 WebAccelerator Product Documentation

The F5 WebAccelerator product documentation set consists of:

- this *Getting Started Guide*
- the *Administration Guide*, which provides information on starting and stopping processes, managing organizations and accounts, backing up and restoring the system, managing various log and configuration files, and using SNMP to monitor the WebAccelerator
- the *Policy Management Guide*, which provides information on creating and editing policies to tailor your WebAccelerator for optimal performance, and on using Performance Reports
- release notes, which provide information about known issues and workarounds and any information that was too late to be included in the product documentation
- the online help in the Administration Tool, which provides information about the Administration Tool user interface

You access the online help by clicking the Help button found on each Administration Tool screen. You can access the product documentation as HTML or PDF files either on the product CD or through a link in the online help. After the product is installed, you can also access the documentation from this directory on the machine on which you installed the WebAccelerator:
`/opt/pivia/dac/docs`

Chapter 2

Planning Your Installation

Installation Overview
Installation Planning Worksheets

Read this chapter only if you are planning to install the F5 WebAccelerator software on your own hardware.

F5 WebAccelerator installation goes smoothly if you do some research and planning first. You should identify what machines you want to install F5 software on, and ensure those machines meet both the hardware and software requirements described in "Installation Requirements" on page 6. You should define which IP addresses the F5 software on various machines will use. You should know what kinds of HTTP requests you want to proxy to which of your origin servers. These are all issues of topology and deployment. Basically, you should understand how you plan to deploy the WebAccelerator in your system before you begin installation.

Read "Network Changes" on page 56 for helpful information on configuring DNS, and how this affects the hostnames you assign to your WebAccelerators.

Installation Overview

Here is an overview of the general installation and configuration tasks needed to get the WebAccelerator up and running:

1. Fill out the appropriate worksheets in "Installation Planning Worksheets" on page 12.
2. Install the main WebAccelerator software, described in Chapter 3, Installing Your WebAccelerator. Use the worksheets you filled out in Step 1 to answer the install prompts correctly.
3. Perform the initial configuration of the WebAccelerator. You log into the Administration Tool (Admin Tool) and create your first cluster with the WebAccelerator you installed in Step 2 and create profiles for the applications you want to accelerate, including choosing a policy set and setting up host maps. This is described in Chapter 4, Configuring Your WebAccelerator.
4. Optionally, you can set up your Linux system to rotate log files automatically for you.
5. Minimally confirm that your WebAccelerator installation is working correctly. See Chapter 7, Confirming Your Installation for more information.
6. Optionally install other F5 software on other machines, such as additional WebAccelerator WebAccelerators, standby Management Consoles, or WebAccelerator Remote products. See the relevant installation information for the software you want to install.
7. Log into the Admin Tool and assign the new Accelerators and WebAccelerator Remotes to clusters and applications.
8. Reconfigure your network and firewalls to place your WebAccelerator software into production use. See Chapter 8, Configuring Your Network for more information.

Installation Planning Worksheets

To prepare for installation, we recommend you complete these worksheets before beginning installation. Complete the appropriate worksheet for each install you plan to do.

WebAccelerator Machine

In all WebAccelerator installations, there is a single machine on which the primary Management Console and an WebAccelerator are installed. There might be other

machines with F5 software installed, but those might contain only an WebAccelerator or an WebAccelerator and a standby Management Console. This worksheet is for the main WebAccelerator machine, which has the primary Management Console and an WebAccelerator:

Table 1 WebAccelerator Installation Planning

Question	Considerations	Value to Enter During Install
a) Hostname for Management Console	This is the hostname you configure in your DNS server to be recognized as the primary Management Console for this WebAccelerator. **Make a note of this hostname:** you need to it to configure networking for all WebAccelerators not installed on this machine, so they know where their Management Console is located. You also need it to access the Administration Tool for configuration, account management, or performance monitoring. If you are installing one or more standby Management Consoles, they also use this hostname when they switch from standby to primary during failover. We generally recommend using a meaningful hostname that is independent of the machine on which the Management Console is installed. You can use any appropriate hostname if you add it to your DNS server with the IP address of the machine on which you are installing this WebAccelerator. **Default**: local hostname of machine	
b) Directory for WebAccelerator software to be installed	This is the directory on the machine into which the WebAccelerator software is installed. The default value works for most installations. Ensure the directory you specify has enough space. **Default**: `/opt/pivia` **Min Space required**: 1500 MB	
c) Directory for log file storage	This is the directory that the WebAccelerator uses for storing log files. The default value works for most installations. Ensure the directory you specify has enough space. We recommend that it be a dedicated partition. **Default**: `/opt/pivia/log` **Min Space required**: 500 MB	

Table 1 WebAccelerator Installation Planning

Question	Considerations	Value to Enter During Install
d) Directory for Management Console depot	This is the directory that the WebAccelerator uses for storing Management Console files. We recommend that it be a dedicated partition. **Default**: `/opt/pivia/depot` **Min Space required**: 5000 MB	
e) Directory for database data file storage	This is the directory that the WebAccelerator uses for storing database data files. The default value works for most installations. Ensure the directory you specify has enough space. We recommend that it be a dedicated partition. **Default**: `/opt/pivia/dbdata` **Min Space required**: 3072 MB	
f) Directory for database control files and logs	This is the directory that the WebAccelerator uses for storing database control and log files. The default value works for most installations. Ensure the directory you specify has enough space. We recommend that it be on a different physical partition than where you placed your data files. **Default**: `/opt/pivia/log/oracle` **Min Space required**: 300 MB	
g) Management Console password	This is the password of the administrator account used for logging into the Administration tool. It is also the password for the database and sys/system accounts used for configuring and accessing the database. **Make a note of this password:** you need it to log into the Administration Tool and complete WebAccelerator configuration. You can change this password in the Administration Tool.	
h) Email notification	Do you want to be notified using email whenever a WebAccelerator process exits abnormally? Most installations choose yes. **Default**: `yes`	Circle: YES NO
i) Email address	Provide the email address you want the notifications to be sent to. This can be any valid email address. There is no default value.	

Table 1 WebAccelerator Installation Planning

Question	Considerations	Value to Enter During Install
j) Directory for cached data	This is the directory that the WebAccelerator uses for caching files. The default value works for most installations. Ensure the directory you specify has enough space. We recommend that it be a dedicated partition. **Default**: `/opt/pivia/hds` **Min Space required**: 500 MB	

Accelerator Machines

WebAccelerator is the name for the software that actually services requests for your site. You always have one main WebAccelerator machine that contains the primary Management Console and one WebAccelerator. You can install additional WebAccelerators if they are included in your license. You can also install WebAccelerator Remotes. WebAccelerators and WebAccelerator Remotes can only be installed one to a machine.

Use this worksheet for installation planning on machines where you plan to install either an WebAccelerator or a WebAccelerator Remote:

Table 2 WebAccelerator Installation Planning

Question	Considerations	Value to Enter During Install
i) Hostname for Management Console	This is the hostname for the Management Console that you defined when you installed the main WebAccelerator. Use the same hostname that you used in your answer for a) Hostname for Management Console in Worksheet 1.	
ii) Directory for WebAccelerator software to be installed	This is the directory on the machine into which the WebAccelerator software is installed. The default value works for most installations. Ensure the directory you specify has enough space. **Default**: `/opt/pivia` **Min Space required**: 1500 MB	

Table 2 WebAccelerator Installation Planning

Question	Considerations	Value to Enter During Install
iii) Directory for cached data	This is the directory that the WebAccelerator uses for caching files. The default value works for most installations. Ensure the directory you specify has enough space. We recommend that it be a dedicated partition. **Default**: `/opt/pivia/hds` **Min Space required**: 500 MB	
iv) Directory for log file storage	This is the directory that the WebAccelerator uses for storing log files. The default value works for most installations. Ensure the directory you specify has enough space. We recommend that it be a dedicated partition. **Default**: `/opt/pivia/log` **Min Space required**: 500 MB	
v) Email notification	Do you want to be notified using email whenever an WebAccelerator process exits abnormally? Most installations choose yes. **Default**: `yes`	Circle: YES NO
vi) Email address	Provide the email address you want the notifications to be sent to. This can be any valid email address. There is no default value.	

Standby Management Console Machines

If you purchased the standby Management Console feature, you can choose to install a standby Management Console on one or more of the machines you chose to install an WebAccelerator on. The standby Management Console replicates the database from the primary Management Console, and can become the primary Management Console if the primary Management Console fails. Because the standby Management Console can become the primary Management Console at any time, it has the same installation requirements as a regular WebAccelerator machine. Use exactly the same values that you filled in for the main WebAccelerator on Worksheet 1, especially for hostnames.

Chapter 3

Installing Your WebAccelerator

Installation Process
Post-Installation Steps
Meeting Minimum Software Requirements
Uninstalling WebAccelerator Software
Installing a Standby Management Console

Read this chapter only if you are installing the WebAccelerator software on your own hardware.

This chapter guides you through the installation script that places a new WebAccelerator on the machine you choose as your WebAccelerator Management Console machine. The install script installs both a Management Console and an WebAccelerator on the machine you choose.

If you are upgrading your WebAccelerator from a previous version, go to Chapter 6, Upgrading Your WebAccelerator and follow those steps.

After installation, if you want to restore your server machine to its original state for any reason, you can uninstall the WebAccelerator software.

Installation Process

To install the WebAccelerator software:

1. Log in as root to the RedHat Linux machine on which you want to install the WebAccelerator. There should be no WebAccelerator or Performance Server software currently running on this machine. If there is, either:
   - uninstall it following the steps in "Uninstalling WebAccelerator Software" on page 22, or
   - choose a different machine
   - if the software is a Performance Server from a prior release, you can upgrade it following the steps in Chapter 6, Upgrading Your WebAccelerator
2. Place the WebAccelerator installation CD into the CD-ROM drive and mount that drive. For most RedHat Linux installations, you can mount the CD-ROM using:
   `# mount /mnt/cdrom`

   If this does not work, see "Managing CD-ROMs" on page 68.
3. Go to the directory on the CD-ROM:
   `# cd /mnt/cdrom`
4. Run the installation script by typing:
   `# ./install`

   The script begins running. The prompts follow these conventions:
   - If a prompt has a limited set of valid replies, such as `yes`, `no`, `1`, `2`, or `q` for quit, the reply options are either listed in the menu above the prompt or indicated in parentheses ( ) on the prompt line. For example, this prompt accepts only `y` or `n` as a valid response:
     `Create this directory? (y/n) [y]:`
   - The option value displayed in brackets next to any prompt is the default. This default value is used if you simply press Enter in response to a prompt. For example, the default value for this prompt is `yes`:
     `Answer yes or no (yes/no) [yes]:`
   - Use these special characters in response to a prompt to:

`!`	Launch `/bin/bash`
`q!`	Return to main menu
`#`	Print filesystem disk space usage
`?`	Print this help information
5. You see the initial dialog on your screen that includes F5's technical support email address and phone number.

The installation script checks to see if your machine meets the required software dependencies. See "Installation Requirements" on page 6 for a list of these dependencies. See "Meeting Minimum Software Requirements" on page 21 for information on what to do if your system does not meet the minimum software requirements.

6. You are shown your installation options. For the first install, you should install the full F5 WebAccelerator complete with Management Console. Enter **1**.
7. Because you are installing a Management Console, the next prompt asks if you want to install a Primary Management Console or a Standby Management Console. For this main F5 WebAccelerator install, you want to install a primary Management Console. Enter **1**.

   After you complete this install and configure your WebAccelerator, you might want to install a standby Management Console on a different machine, possibly one you had chosen to install an additional WebAccelerator on.
8. As the installer begins to install the WebAccelerator software, it prompts you for information specific to your system. Get your completed worksheet from Worksheet 1 WebAccelerator Installation Planning on page 13 and use it as you answer this next series of prompts. The default answers work for most WebAccelerator systems.
9. Enter the hostname to use for the Management Console. See the value you selected for question a) Hostname for Management Console on your worksheet. The worksheet explains the considerations involved in choosing this hostname.
10. You are shown the disk partitions for the machine. Enter the directory where you want the WebAccelerator to be installed. The default is `/opt/pivia`. See the value you selected for question b) Directory for WebAccelerator software to be installed on your worksheet.

    Note: If any directory you specify does not exist, the installer can create it for you. In most installations, you can simply use the default directory names.
11. Enter the directory where the WebAccelerator can store the log files. The default is `/opt/pivia/log`. See the value you selected for question c) Directory for log file storage on your worksheet.
12. Enter the directory where the WebAccelerator can store the Management Console depot. The default is `/opt/pivia/depot`. See the value you selected for question d) Directory for Management Console depot on your worksheet.
13. Enter the directory where the WebAccelerator can store the database files. The default is `/opt/pivia/dbdata`. See the value you selected for question e) Directory for database data file storage on your worksheet.
14. Enter the directory where the WebAccelerator can store the database control and log files. The default is `/opt/pivia/log/dblog`. See the value you selected for question f) Directory for database control files and logs on your worksheet.

15. You are prompted for the password for the Management Console. This password is used for the Administration Tool administrator account and also for the database and sys/system account. First enter the password, then enter it again for verification. See the value you selected for question g) Management Console password on your worksheet.
16. You are asked if you want to be notified by email whenever a WebAccelerator process exits abnormally. If you answer yes, you are then prompted to enter the email address to which these notifications should be sent. See the value you selected for question i) Email address on your worksheet.
17. The Management Console is installed, including the WebAccelerator database. The database is large, so this step can take a few minutes.
18. The WebAccelerator is installed next. Enter the location where the WebAccelerator can store cached data. See the value you selected for question j) Directory for cached data on your worksheet.
19. The installation is complete. You are returned to the main menu, where you must quit the installation. You cannot install a second Management Console or WebAccelerator on a machine on which one already resides.
20. You are shown the installation summary. Note any information you might want to reference in future. The installation summary information is stored in a log file called `pvinstall.log` in your root `/tmp` directory. We recommend that you save this file as a record of your selections during installation.
21. After the summary, unless there was an error during the install, you are asked if you want to start the WebAccelerator processes. Enter **`y`** for yes, unless there is some reason you want to postpone starting them.
22. Eject the installation CD using the command:
    `eject /mnt/cdrom`

    If this command does not work, see "Unmounting CD-ROMs" on page 69 for additional instructions.

Your installation is complete.

Post-Installation Steps

After your installation of the WebAccelerator is completed, you perform some other steps before you have a complete, running system:

1. You must do some basic configuration of the WebAccelerator before it can begin accelerating your applications. See Chapter 4, Configuring Your WebAccelerator. Generally, you do this basic configuration first, and then install any additional WebAccelerator software to other machines in your system.

2. Optionally, you can add WebAccelerators and WebAccelerator Remotes on other machines. See Chapter 5, Installing Accelerators. First you install them, then go to the Admin Tool and configure them.
3. Optionally, you can also add standby Management Consoles to provide failover support if your primary Management Console fails. See "Installing a Standby Management Console" on page 23.
4. You confirm your installation following the steps in Chapter 7, Confirming Your Installation.
5. You must reconfigure your network to use your WebAccelerator. This includes defining the hostname you selected in Step 9 to your DNS server. See Chapter 8, Configuring Your Network for more information.

Meeting Minimum Software Requirements

When you run the installation script, it checks to see if your machine meets the minimum requirements. If your machine does not meet the minimal software installation requirements, the installation script tells you this and provides you with an opportunity to exit the install script.

Run the special installer that F5 provides on the distribution CD that installs the necessary packages to meet the minimal software requirements. For example:

1. Early in the installation script, you see informational messages about checking dependencies. If a dependency check fails, you need to exit the installation and install the required software that is currently missing.

   Enter `n` at the next prompt to exit the installation script:

   ```
   Do you want to continue with the installation? (y/n) [y] : n
   ```
2. Go to the `extra` directory on your WebAccelerator distribution CD and run the installer:

   ```
   # cd /mnt/cdrom/dist/Linux/extra/install
   ```
3. The script begins by showing you what it will do. Press Enter when asked to do so.
4. Next the script steps through and installs each package as is necessary.
5. When the script reaches the NTP installation portion, you are asked for the location of the time server on your network regardless of whether NTP is already installed. If you do not have an NTP time server on your network, you can enter a publicly available one here. For more information, see:
   http://ntp.isc.org/bin/view/Servers/WebHome

   Here is an example:

   ```
   NTP:
   ```

```
Enter the host name of the reference NTP
server for your network: ntpserver.company.com
ntp - Version 4.1.0 is already installed.
This machine has an offset of 0.000972 seconds
from the reference NTP Server time.company.com.
Do you want to synchronize time with the NTP server [y/n]?
y
Time synchronized successfully with the NTP server.
        Starting NTP daemon...
        Done.
Done.
OpenSSH:
openssh - Version 3.1p1 is already installed.
openssh-askpass - Version 3.1p1 is already installed.
openssh-server - Version 3.1p1 is already installed.
openssh-clients - Version 3.1p1 is already installed.
        Starting sshd...
        Done.
Done.
```

When you get to a package that needs to be installed, there is a slight pause while the package is installed on your system.

You are done updating your machine to meet WebAccelerator's minimal software requirements. You can now perform the software installation as described in "Installing Your WebAccelerator" on page 17.

Uninstalling WebAccelerator Software

You uninstall any WebAccelerator software using the `uninstall` script:

1. Log in as root to the machine where you want to uninstall WebAccelerator software.
2. Enter:
   `/opt/pivia/uninstall`
3. The `uninstall` script stops the locally running WebAccelerator processes for you. If the WebAccelerator was not running, you see a message that says an error occurred while stopping the process. This error can be ignored.

   Enter **`y`** in response to continue with the uninstall.
4. You are asked if you want to delete the various directories where the WebAccelerator software data and log files are stored. Depending on your reasons for uninstalling the software, you might want to save the contents of these directories.

If you allow the uninstall script to delete a directory, and that directory happens to be a mount point for a disk partition, the script issues a warning:

```
Delete /opt/pivia/hds and its contents (y/n) [n]: y
Can't remove directory /opt/pivia/hds: Device or resource busy at
uninstall line 305
```

Under these circumstances, the directory is not deleted, but the data within the directory is deleted.

5. When the `uninstall` script is done, it displays a message with the status of the uninstall.

Installing a Standby Management Console

The standby Management Console is a second Management Console that you can bring online when your primary Management Console fails, due to a machine crash or network failure. Generally, standby Management Consoles are installed on machines chosen to run additional WebAccelerators. To make installation simpler, WebAccelerator combines the standby Management Console install with an WebAccelerator install.

Because the standby Management Console should be a mirror of your primary Management Console, it picks up much of its configuration from the primary Management Console. For any other information you need to enter during the install, refer to the same WebAccelerator worksheet you used to install the primary Management Console.

The installation process is similar to the installation of your primary WebAccelerator Management Console and WebAccelerator.

Installation Process

1. Log in as root to the Red Hat Linux machine on which you want to install the standby Management Console and an WebAccelerator. This machine must not have any WebAccelerator software currently installed on it. If the machine is already running an WebAccelerator, stop the WebAccelerator and uninstall it, following the steps in "Uninstalling WebAccelerator Software" on page 22.
2. Follow the same installation process that you used when you installed the main WebAccelerator, described in "Installation Process" on page 18. Begin with Step 2 on page 18 and continue on.
3. When you get to Step 7 on page 19, you are asked if you want to install a Primary Management Console or a Standby Management Console. Enter 2 for Standby Management Console. That is the only difference in the installation.

4. Continue with the rest of the installation steps, Step 8 on page 19 through Step 22 on page 20. You do not need to perform any configuration of the standby Management Console.

Chapter 7, Confirming Your Installation provides information on how you confirm that your standby Management Console is working correctly. This is the only post-installation step for standby Management Consoles.

Chapter 4

Configuring Your WebAccelerator

The Admin Tool
Host Mapping
Configuration Process
Changing Your Configuration
Changing Your Preferences
Clearing Your Cache
Monitoring the WebAccelerator
Log File Rotation
Finding the MAC Address
Advanced Configuration

Most configuration and administration tasks are performed in the Admin Tool. This chapter describes the basic configuration which sets up your F5 WebAccelerator to manage your web applications. You must perform these tasks regardless of whether you installed the WebAccelerator as software or whether you purchased it as an appliance.

Apart from these tasks that you perform using the Admin Tool, you can set up Log File Rotation directly in Linux. This eliminates the need for you to manually clean out your log files when they become too large.

For information on detailed configuration of the F5 WebAccelerator, see the *Policy Management Guide*.

The Admin Tool

All basic administration is performed using the Admin Tool. It is a web-based interface, and you access it by pointing your web browser to the host on which your WebAccelerator Management Console resides. For example, entering a URL such as this, which uses the hostname you specified for the Management Console machine, takes you to the default login for the Admin Tool:
`https://mynwa.mydomain.com:8443`

where you enter your Management Console subdomain and domain instead of `mynwa` and `mydomain`. If you have an unusual installation where SSL is not enabled, use `http` and port `8080` instead of `https` and port `8443`.

The first time you log into the Admin Tool, add the URL to Favorites in your browser to make it easy to get to the login screen in future. Perform these tasks to configure your WebAccelerator initially:

- set up your WebAccelerator clusters
- create your application profiles
- change any preferences as needed

After the initial configuration, you can use the Admin Tool to:

- add new WebAccelerators or WebAccelerator Remotes to your clusters
- update any pre-defined policy sets, if a more recent version has been created
- create new custom policy sets
- update your host mapping and change any option settings
- clear your cache when needed
- change your password or time zone preferences
- monitor WebAccelerator performance

In addition, there are a few standard administrative procedures you should perform occasionally. These include periodically backing up the WebAccelerator's data files and cleaning out your log files if you do not want to set up automatic log file rotation. These procedures are described in the *Administration Guide*.

Host Mapping

As part of the configuration process, you map domain names (Requested Hosts) to the actual locations where your web servers reside (Destination Host). This mapping must include any domain names that you expect your Accelerators to see on incoming

HTTP and HTTPS requests. The WebAccelerator is also capable of managing requests for unmapped domains. These are called Unmapped Requests.

When you set up your host mapping, you are telling the WebAccelerator where to go for content. It looks to the destination you specify for the initial content needed to respond to a request, for the content to refresh its cache, and as the location to proxy a request to when it cannot service the request from its cache.

Understand how host mapping works, and the options you can specify for each destination, before you set values in the Host Maps section of the Admin Tool.

You create a host map by identifying the domain as it appears on the HTTP HOST request header. These domains are called Requested Hosts. You then identify the hostname of your origin servers. These hostnames are called Destination Hosts. This hostname is the actual location where the WebAccelerator should go for content to service that request.

Valid Hostnames

The hostname of a domain is broken into subdomains, separated by periods, as in `www.somesite.com` or `catalog.sales.bigretailsite.com`. When you specify a hostname for a Destination host, you must specify the complete hostname, because the WebAccelerator must know exactly where to proxy for content. When you specify the hostname for the Requested Host, you can use an asterisk (*) as a wildcard for the first part of the domain. The asterisk must be the first character in the name and must be followed by a period. It can represent one or more subdomains. This allows you to map several subdomains to one origin web server in one step, a time-saver if your site has many subdomains.

Valid examples of a wildcard in the Requested Host name are:

- `*.sales.somesite.com`, which maps
  ```
  direct.sales.somesite.com
  marketing.sales.somesite.com
  marcom.marketing.sales.somesite.com
  ```
  all to the same Destination Host
- `*.somesite.com`, which maps
  ```
  www.somesite.com
  engineering.somesite.com
  direct.sales.somesite.com
  marketing.sales.somesite.com
  marcom.marketing.sales.somesite.com
  ```
  all to the same Destination Host
- `*.com`, which maps all incoming requests that end in `.com` to one Destination Host
- `*`, which maps all incoming requests to one Destination Host

These mappings have one exception: if several of the Requested Host hostnames could match a request, the most specific match is chosen. For example, if you define these Requested Host hostnames:

```
a.com
www.a.com
*.b.a.com
*.a.com
```

When requests are received that contain these URLs, they are mapped to the Requested Hosts as follows:

- a request to www.a.com is mapped to `www.a.com`, not to `*.a.com`
- a request to a.com maps to `a.com`
- requests to c.a.com and b.a.com both map to `*.a.com`
- a request to c.b.a.com maps to `*.b.a.com`

This allows you to define several specific Requested Host hostnames, and a much more general Requested Host hostname to map all other requests. For example, perhaps you want any requests directed to your internal HR and payroll applications to go to their own Destination Host, and all other requests to go to a different Destination Host. You could define these Requested Hosts and Destination Hosts:

- `*.hr.inside.anycompany.com > hr.anycompany.com`
- `*.payroll.inside.anycompany.com > payroll.anycompany.com`
- `*.anycompany.com > primary.anycompany.com`

Defining a Destination Host

When you add your origin servers to the host map, you can create a definition for each server. You can also set various connection values and timeouts in the definition. These properties allow the WebAccelerator to work as efficiently as possible with your origin servers. In many systems, the default values for these properties work well. Change them only if you need to tune your system for more efficient operation. You can change the values at any time after your initial configuration.

Protocol Security

This defines the protocol used for connections from the WebAccelerator and your origin servers. If you choose Same as Original Request, the protocol used for the client request is used for proxies to your origin servers. If you choose HTTP Only, then only HTTP is used for proxies even if HTTPS is used for the original request. Select the HTTP Only option only if your WebAccelerator and your origin servers are on the same protected network and you want to avoid the overhead of HTTPS connections.

Timeout Properties

- Connection Timeout

  Specifies the time in seconds your WebAccelerator waits for a response from your origin servers when proxying for content.

- Read/Write Timeout

  Specifies the time in seconds your WebAccelerator waits for a read or write operation to your origin servers to complete before giving up.

- Retries

  Specifies the number of times your WebAccelerator attempts to connect to your origin servers in the event of a connection failure.

- Lifetime

  Specifies the time your WebAccelerator holds a persistent connection open to your origin servers. You specify a number and select the unit of time to be used.

  Each time an WebAccelerator first proxies to an origin server, it forms a persistent connection to that server. All subsequent proxies are performed over that connection. However, some web servers and application servers timeout persistent connections regardless of the activity occurring on the connection. Therefore, before performing each proxy, the WebAccelerator checks to see if the connection has been opened for longer than Lifetime. If so, the WebAccelerator drops the connection and reconnects before attempting the proxy.

Express Connect Options

Express Connect options allow you to identify subdomains that the WebAccelerator automatically generates for qualifying links and embedded URLs. Most browsers create up to two persistent TCP connections for each domain from which they are requesting data. By modifying embedded URLs with unique subdomains, the WebAccelerator causes the browser to open more persistent connections (up to two per subdomain generated by the WebAccelerator). The browser downloads data much faster and the user perceives a much faster response time from your site.

To use Express Connect, first identify the number of subdomains you want the WebAccelerator to generate for each supported protocol (HTTP or HTTPS). This is the maximum number of subdomains that you want the WebAccelerator to generate for each page that it serves. Remember that for each additional subdomain you request, the browser probably opens two additional persistent connections over and above what it normally does. So if you request 1 subdomain for the HTTP protocol, a browser could open up to a total of 4 persistent connections to the WebAccelerator when requesting pages over the HTTP protocol. The WebAccelerator uses these additional subdomains only on embedded URLs or links that request images or scripts.

In addition to identifying the number of subdomains, you also assign a prefix to those subdomains. For example, suppose the Requested Host for this mapping is `www.somesite.com` and you request two additional subdomains for the HTTP protocol and assign a subdomain prefix of `pv`. The WebAccelerator changes the domain on qualifying embedded URLs and links to use these domains:
`pv1.www.somesite.com`
`pv2.www.somesite.com`

You must configure DNS with these entries and they must map to the same IP address as your base origin server (`www.somesite.com` in this example). Note that these domains are used only for requests/responses between the client and the WebAccelerator. Your origin servers never see a request that uses these domains.

Unmapped Requests

Any incoming HTTP requests that do not match a Requested Host are not mapped to a Destination Host. These are called unmapped requests. These unmapped requests generate an error message. You can define your Requested Hosts so that there is always a match for legitimate incoming requests.

Alternatively, for each Accelerator cluster, you can choose to have it process unmapped host requests. Check the "Process requests for unmapped hosts" box. When this option is enabled, unmapped requests do not cause an error. Instead, the requests with unmapped hosts are forwarded to whatever domain is on the request. This can be useful for initial testing of your system.

Security is a concern if you enable this option. If you check the unmapped hosts box and do not specify a proxy server, you enable the WebAccelerator to be used as a relay. Unless the WebAccelerator and your origin servers are private and heavily protected, we do not recommend this option. If you specify a proxy server to forward the requests to, and the proxy server is set up to deal with unwanted or unsanctioned requests, this might be an acceptable method to handle requests for unmapped hosts.

Configuration Process

Step 1 Logging In

1. To log in to the Admin Tool, use your web browser to navigate to the host on which it is installed. For example:
   `https://mynwa.mydomain.com:8443`

   This brings you to the login screen. You should see two text entry boxes, one for username and one for password.

2. To log in, use the administrator account. Enter the username and password:
   - Username: **`administrator`**
   - Password: the Management Console password you set during installation or appliance setup (see Step 15 on page 20)
3. You see your home page. The home page you see when you first log in shows some performance reports, which should be empty because the system is not running yet.

Step 2 Setting Up Clusters

You must first create Accelerator clusters. Every Accelerator must be assigned to a cluster. You can have as many clusters as you like, and a cluster can have one or more Accelerators assigned to it. An Accelerator can only be assigned to one cluster.

Clusters are arbitrary groupings of Accelerators. Set them up to organize your system in a way that makes the most sense. You might group them by physical location, or by application that they would accelerate. Keep in mind that an application can only have one cluster assigned to handle its requests. All the Accelerators that you want assigned to an application must be in the same cluster. An application treats any Accelerator not in its cluster as a WebAccelerator Remote. However, a cluster can support multiple applications.

An example of grouping by location: if you have several Accelerators in your data center, you could put them all in one cluster even though they support several different applications. However, if there is another Accelerator in a different office that also handles requests for one of those applications, it must also belong to the data center cluster, even though it is not located in the data center. In this case, you might group by application instead, and define a cluster for each application, with the appropriate Accelerators assigned to the appropriate clusters.

You create a cluster for this first Accelerator, and then assign the Accelerator to it:

1. Click on the Clusters menu item in the menu bar. You go to the Clusters screen:

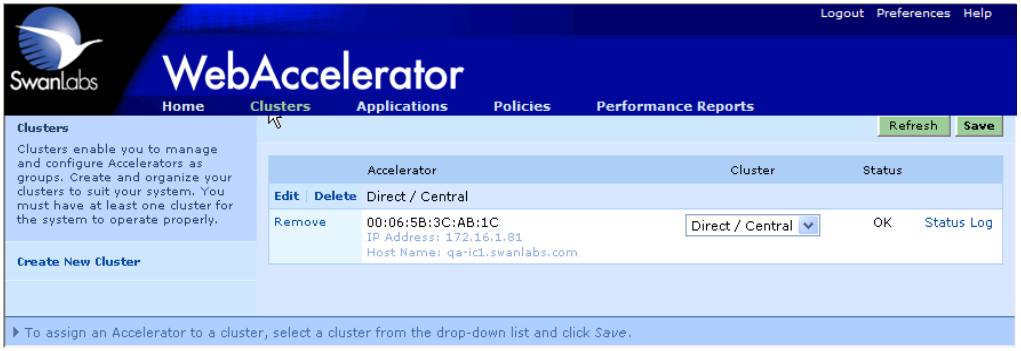

Recently installed Accelerators that have not been added to a cluster are shown at the beginning of the table under the New Accelerator heading, and show a status of Standby. New Accelerators are in standby mode until you assign them to a cluster. At this point, you should see the single Accelerator that was installed with the Management Console during your WebAccelerator installation.

Accelerators are shown on this table with as much information as the WebAccelerator knows about them. The MAC address is shown, in addition to the host name and IP address if possible. For information on how to look up the MAC address for a machine, see "Finding the MAC Address" on page 42.

2. Click on the Create New Cluster link at the left of your screen.
3. You see the cluster definition screen:

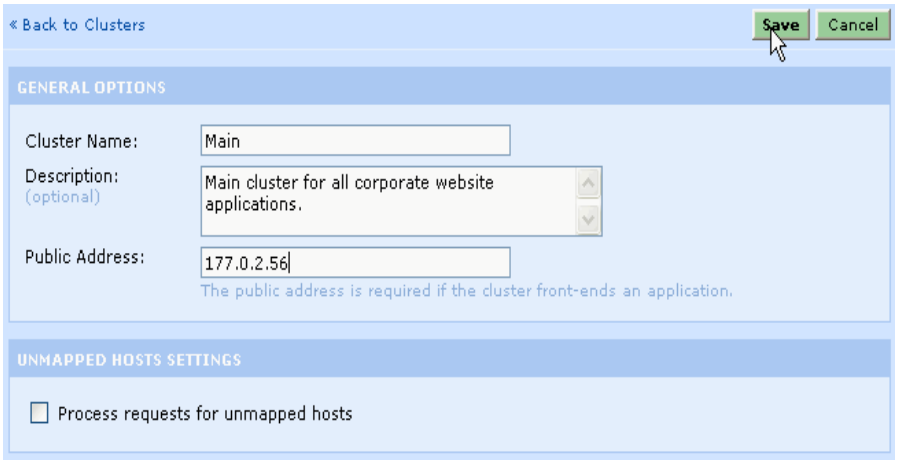

4. Enter a name for the cluster. Provide a descriptive name that is unique to this WebAccelerator installation.
5. Optionally, provide a description of this cluster. You might include its location, what applications it accelerates, or other relevant information.
6. Public Address: Provide the public address of the cluster, either IP address or hostname. This is required if the cluster is accelerating an application. The address you enter is used internally to direct remote Accelerators to this cluster when they send requests for this cluster's application. You might need to reconfigure your network, such as defining the address in DNS, so the address can be resolved to an IP address. See "Defining WebAccelerator Hostnames" on page 56.
7. Optionally, under Unmapped Hosts Settings, check the "Process requests for unmapped hosts" box.

Check this box if you want the cluster to process Unmapped Requests requests for a host that is not in the host map. Be sure you are aware of the security implications discussed in "Unmapped Requests" on page 30. If you check this box, provide the additional information on how to process these requests. If you do not check this box, skip ahead to Step 10 on page 34.

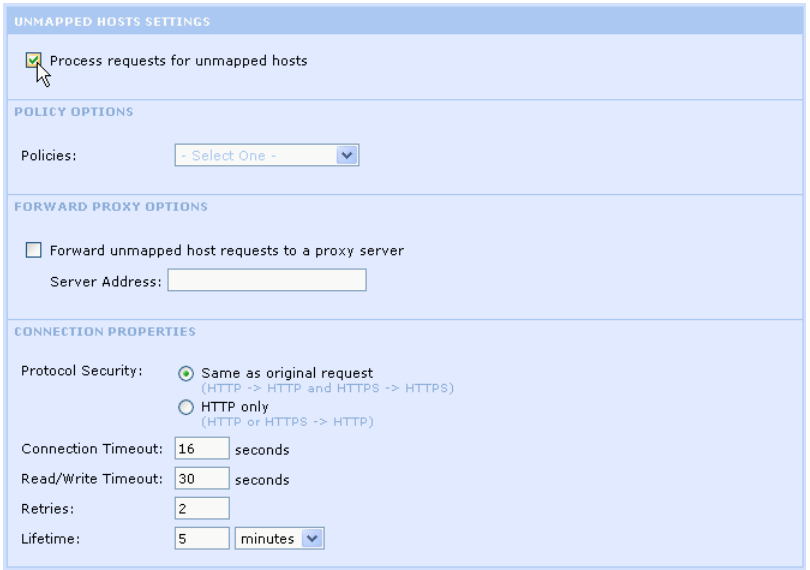

a. Use the drop-down list to select which policy set to apply to any requests for unmapped hosts.

b. Optionally, check the "Forward unmapped host requests to a proxy server" box if you want all requests for unmapped hosts forwarded to a proxy server. Provide the address, IP address or hostname, of the proxy server.

c. Use the default settings or change the values set for the connection properties that are used when proxying the unmapped host request. See "Defining a Destination Host" on page 28 for information on what the properties control. These are the same properties you set for any destination host:

   – Protocol Security
   – Connection Timeout
   – Read/Write Timeout
   – Retries
   – Lifetime

8. Click the Save button to save and activate the settings. This takes you back to the main Cluster screen.
9. Assign the Accelerator to the newly created cluster. In the scrolling box next to the Accelerator, select the cluster you want to assign it to.
10. Click the Save button.
11. Click the Refresh button to see the changes. It might take a moment.

Step 3 Creating Application Profiles

After you created your clusters, you create profiles for the applications you want the WebAccelerator to accelerate. The application profile provides key information to the WebAccelerator so that it can handle requests to your application appropriately.

This involves assigning an Accelerator cluster to the application and setting up a host map that tells the WebAccelerator where to proxy for content when handling requests for that application. You also choose a set of pre-defined policies to use for the application. Later, you can tailor those policies or create an entirely new policy set.

1. Click on the Applications menu item in the menu bar. You go to the Applications screen, where you should see an empty table because no applications are currently defined. This example shows several test applications:

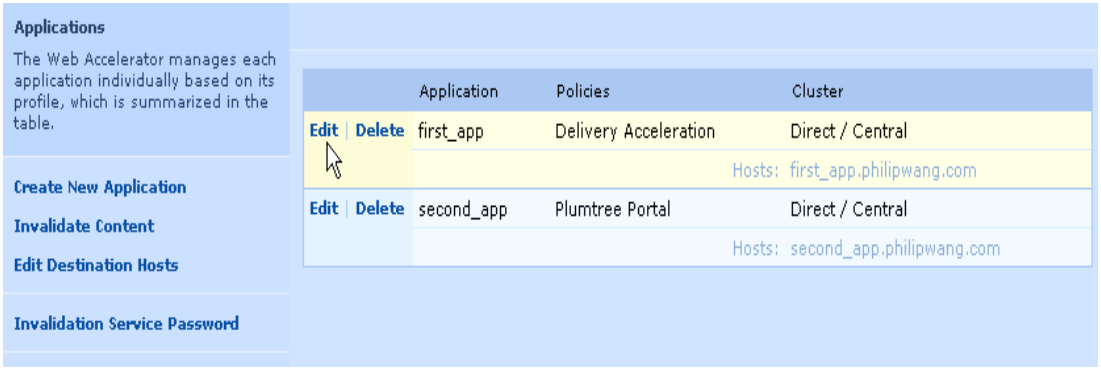

2. Click on the Create New Application link at the left of your screen.

3. You see the applications profile screen. After you have completed a profile, it should look like this:

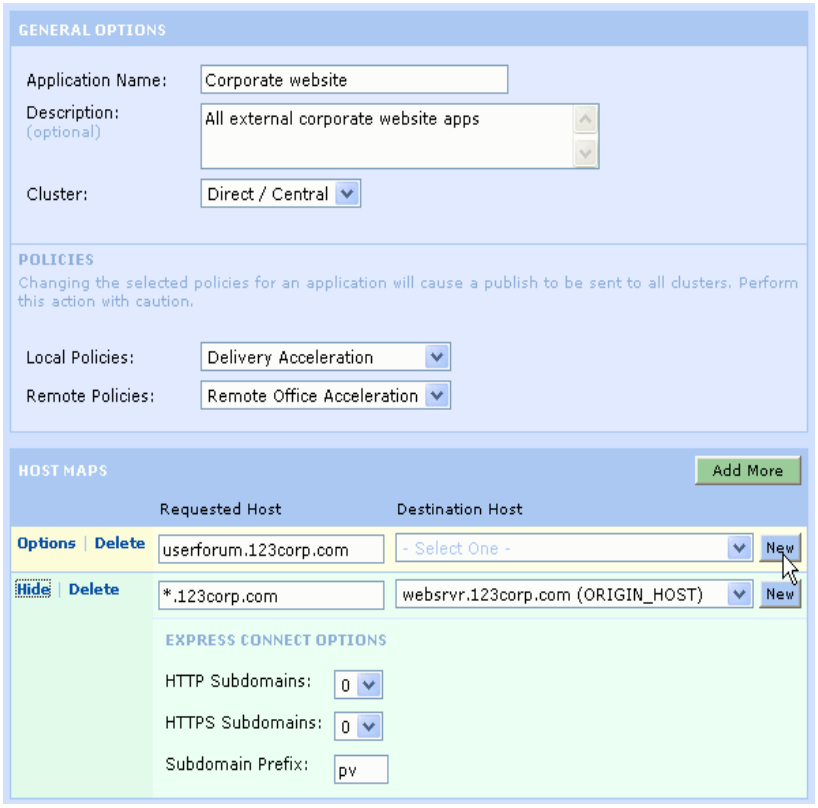

4. Enter a name for the application. Provide a descriptive and useful name for this application.
5. Optionally, provide a description of this application.
6. Select the appropriate Accelerator cluster from the drop-down list.
7. Under Policies, select the appropriate local policy set from the list. The list is divided into the pre-defined policy sets and any published custom sets created for your system.
8. Select the appropriate remote policy set to be used for requests that come from a remote Accelerator. A request that comes from any Accelerator not in the cluster is treated as coming from a WebAccelerator Remote. In most cases, use the default value.
9. Create a host map for the application.

Here is an example of a host map: There are three domains expected on incoming requests that must each be mapped to a different origin web server in order to be serviced properly. The fourth Requested Host uses a wildcard so that all other incoming requests directed to the company's web servers are mapped to `websrvr.123corp.com.`:

Requested Host	Destination Host
`www.123corp.com`	`websrvr1.123corp.com`
`ww2.123corp.com`	`sales.123corp.com`
`userforum.123corp.com`	`msgboard.123corp.com`
`*.123corp.com`	`websrvr.123corp.com`

Start by entering a Requested Host. Enter each domain name (hostname) that might appear in HTTP requests for your web application. These are the same hostnames that are mapped in DNS to the server machine on which your WebAccelerator is running.

When you specify the hostname for the requested host, you can use an asterisk (*) as a wildcard for the first part of the domain name. See "Valid Hostnames" on page 27. You can use wildcards to map a group or range of requested hostnames to a single Destination Host.

10. Next, enter the matching Destination Host. Click the New button to specify a new destination host that has not yet been defined.

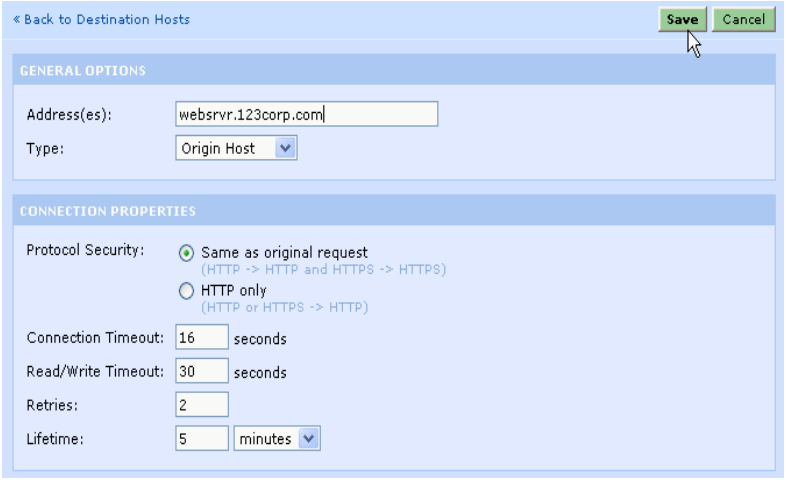

a. In Address(es), specify the DNS name (IP address or host name) of the destination host, or if this is a group of hosts, specify all the IP addresses or host names separated by a comma:
`12.1.100.2, 12.1.100.3, 12.1.100.4, 12.1.100.5`

If you are specifying the name for an origin server, you can include the origin server ports for HTTP and HTTPS:
`hostname``:`*`insecure`*`:`*`secure`*

where *`hostname`* is the address for the origin server, such as `main.123corp.com`, *`insecure`* is the port number for the HTTP port, and *`secure`* is the port number for the HTTPS port. The HTTP port is always shown first. You do not need to include the HTTPS port. If you want to specify an HTTPS port, you must include the HTTP port. These are valid host addresses:
```
main.123corp.com:8000:8443
main.123corp.com:8000
123corp.com:80:443
123corp.com
```

The last two host addresses are equivalent.

To implement load balancing, you can enter a range of servers here, delimited by a slash (/), or a list of servers delimited by a comma (,), or some combination of both. If you specify a range, it must take the form of names with a numerical value at the end of a subdomain name. For example, these are legal examples of ranges:
```
ows4.f5.com/ows20.f5.com
ows.appserver2.f5.com/ows.appserver50.f5.com
```

If a particular hostname is unavailable, the WebAccelerator does not use that hostname in its load balancing until that host is available.

By default the WebAccelerator sets a cookie to ensure a given client's connections are always sent to the same origin server. This capability can be managed using the appropriate `pvsystem.conf` parameters. See the *Administration Guide* for more information.

b. Select the Type:

Origin Host – the destination host is an actual web or application server. This is the default.

Proxy Server – the destination host is a proxy server. If this option is selected, the WebAccelerator ensures that all proxy requests to the destination host are performed in a way that is compatible with a proxy server.

For example, the GET field on the request might change from:
```
GET /index.html HTTP/1.1
```
to
```
GET http://www.somesite.com/index.html HTTP/1.1
```

c. Use the default settings or change the values set for the connection properties that are used when proxying a request to this host. See "Defining a Destination Host" on page 28 for information on what the properties control:
   - Protocol Security: Same as Original Request or HTTP Only
   - Connection Timeout
   - Read/Write Timeout
   - Retries
   - Lifetime

d. Click the Create button to save and activate the settings.

e. Choose the new destination host from the drop-down Destination Host list.

11. Optionally, click the Options link to open a box where you can enter Express Connect options for each host (see "Express Connect Options" on page 29)
    - HTTP Subdomains: enter the number of subdomains for HTTP requests.
    - HTTPS Subdomains: enter the number of subdomains for HTTPS requests.
    - Subdomain Prefix: select the prefix that you want used for your subdomains.
12. Click the Add More button to add another line if you need to add another Requested Host and Destination Host pair.
13. Click the Save button to save this application profile or Cancel to ignore these settings.

Your system should be up and running now. The required configuration steps are complete. You can logout now or do a few optional configuration steps.

Changing Your Configuration

You can change your configuration at any time by logging back into the Admin Tool.

You can change, add, or delete Accelerators and clusters. You can change, add, or delete any application profile. You can import or create new policy sets and assign them to an application.

However, always clear your cache when you change your policy sets or change your destination hosts. This ensures new requests are serviced with fresh content, instead of content cached under previous policies or from the original destination hosts. To clear your cache, use the Invalidate Content link on the Applications screen.

Changing Your Preferences

The password for the administrator account was set during installation. It is easy to change the password at any time in the Admin Tool:

1. Log in to the Admin Tool by navigating to the host on which it is installed. For example:
   `https://mynwa.mydomain.com:8443`

   This brings you to the login screen.
2. Log in as administrator by entering `administrator` as the username and entering the current password for the account.
3. Click on the Preferences link in the upper right of the screen. This brings you to a dialog box where you can change:
   - your time zone, and you can check a box to indicate if you want to enable Daylight Savings Time for your zone
   - your administrator password

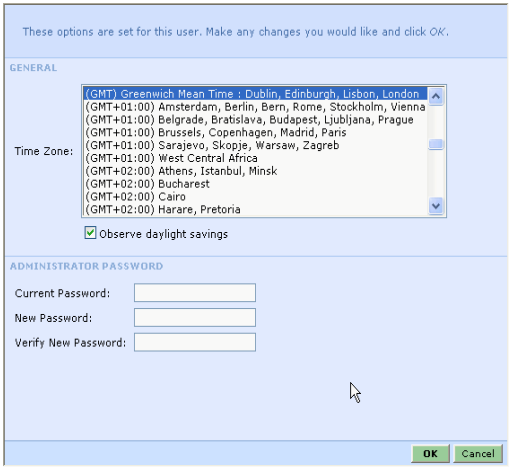

4. When you are done, click the OK button in the lower right of your screen.

Clearing Your Cache

At any time, you can clear the cache by logging in to the Admin Tool and going to the Applications screen. Click on the Invalidate Cache link to go to a screen where you can clear your cache.

You do not need to clear your cache often. If you want to ensure your WebAccelerators refresh their content, you can do this by clearing the cache. If you feel the wrong pages are being served from cache, you can try clearing the cache so that new content is fetched. It is good practice to clear only the content from cache that you feel is incorrect or needs refreshing. You can specify options that limit what content is cleared from cache, and this reduces how many requests the WebAccelerator must proxy to your origin servers for fresh content.

Always clear your cache when you change your policy sets or change your destination hosts. This ensures new requests are serviced with fresh content, instead of content cached under previous policies or from the original destination hosts.

Monitoring the WebAccelerator

Use the Performance Reports tool to view statistics and generate charts and reports on the WebAccelerator's performance and traffic load.

You can change the type of report by choosing a new report from the list on the left. For example, you can change from Hits by Application to Load Time by Content Size.

Click on the Edit Filter button to change the filter criteria, including the time period covered by the report. The filter criteria stay in effect until you change them, even if you switch to a different type of report. Click the Update button in the Edit Filter popup window to update the values, or click the Update and Close button to update and also close the window.

The first time you use Peformance Reports, you should edit the filter to ensure the reports make sense for your applications. It is often helpful to use the Transaction Type selection box to choose an application or nodes for an application as part of the filter. Then all the data you see is for that application.

You can change the time period by selecting a time period from the drop-down list. If you select `Custom...` from the drop-down list, a calendar lets you select the dates for the time period you want. Any invalid dates are greyed out in the calendar. Invalid dates are dates for which there is no historic data, such as any time in the future or dates in the past for which data is no longer saved. Because the WebAccelerator collects so much system data, it cannot store historic data for more than a few weeks.

You can change which data is displayed by selecting certain filter criteria:

- Cluster, where you choose which clusters to show data for
- Content Type, where you choose which responses to show data for, based on how they were classified by type of file. Content type represents the object types defined in the global fragment configuration file. For more information, see the *Administration Guide*.

- Content Size, where you choose which responses to show data for, based on the size of the response
- Transaction Type, where you choose which responses to show data for, based on the node they matched against in the Request Type Hierarchy during response matching

Log File Rotation

Each WebAccelerator keeps several types of log files. They can contain large amounts of data and should be rotated or else the disk partition can fill up. If the disk partition fills up, the WebAccelerator can crash. By default, the WebAccelerator software does not rotate log files, but the WebAccelerator appliance is set up to rotate log files automatically. For detailed information on log files, see the *Administration Guide*.

If you installed the WebAccelerator software on your own machines, you can rotate the log files manually or you can use the Linux logrotation facility to cause the operating system to rotate your log files for you. For more information, see the `logrotate(8)` manual page.

For example, in order to rotate your WebAccelerator's access logs, create a file called `/etc/logrotate.d/webaccl` containing this code:

```
/opt/pivia/log/access_log*log
    {compress
     missingok
     rotate 7
     daily
     size 1000M
     postrotate
     /etc/init.d/webaccl hup_pvac
     endscript

}
```

By default, log rotation occurs daily. For sites with very high traffic, you might want to change your log rotation time table to hourly. Copy:

`/etc/cron.daily/logrotate`

to

`/etc/cron.hourly/logrotate`

Finding the MAC Address

Your Accelerators are identified on the Clusters screen with as much information as the WebAccelerator knows about them. To be sure you can always identify your Accelerators by a unique ID, the MAC address is always shown. The MAC (media access control) address is a hardware address that uniquely identifies every node on a network. If you are uncertain which MAC address is for which physical machine, you can always log into a machine and look up its MAC address.

On Linux systems, the ethernet device is typically called `eth0`. In order to find the MAC address of the ethernet device, you must first switch to `root`, using `su`. Then, use the `ifconfig -a` command to look up the relevant info. For example:

```
# ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:60:08:C4:99:AA
          inet addr:131.225.84.67  Bcast:131.225.87.255  Mask:255.255.248.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:15647904 errors:0 dropped:0 overruns:0
          TX packets:69559 errors:0 dropped:0 overruns:0
          Interrupt:10 Base address:0x300
```

The MAC address is the `HWaddr` listed on the first line. In this example, it is `00:60:08:C4:99:AA`.

Advanced Configuration

Your WebAccelerator should be up and running after this basic configuration. There are a few more steps you need to do, such as changing DNS and doing a few checks to confirm that your system is running correctly. There are also optional steps such as installing additional WebAccelerators. All these steps are described in the following chapters of this guide.

For advanced administration tasks, such as setting options in the `pvsystem.conf` configuration file and global configuration fragment file, or backing up and restoring your system, see the *Administration Guide*.

For information on how to tailor the WebAccelerator policies or create new ones, see the *Policy Management Guide*. That guide describes how requests and responses are processed, and how to create new policies, and change or delete existing policies.

Chapter 5

Installing Accelerators

Installation Process
Uninstalling Accelerators

Read this chapter only if you are installing Accelerator software on your own hardware.

The WebAccelerator is the software that is actually responsible for improving your site's performance. When you install a WebAccelerator as described in Chapter 3, Installing Your WebAccelerator, you install both WebAccelerator software and the software that manages them, called the Management Console.

Depending on your requirements and your software license, you can add any number of machines to your system dedicated to running just the WebAccelerator software. Doing this can improve your site's availability, responsiveness, and increase the total number of connections serviceable by your site.

The WebAccelerator Remote product, if you purchase it as a software license, is installed just like any other Accelerator.

This chapter guides you through the process of installing and defining additional WebAccelerators or WebAccelerator Remotes to your WebAccelerator installation. Do not install them until you install the Management Console with which they need to communicate.

Installation Process

To install an WebAccelerator or WebAccelerator Remote on its own machine:

1. Log in as root to the RedHat Linux machine on which you want to install the WebAccelerator. There should be no WebAccelerator software currently running on this machine. If there is any WebAccelerator software running on this machine, either:
   - choose a different machine for the install, or
   - uninstall the software following the steps in "Uninstalling WebAccelerator Software" on page 22, or
   - if the software is an WebAccelerator from a prior release, upgrade the WebAccelerator by following the steps in Chapter 6, Upgrading Your WebAccelerator
2. Place the WebAccelerator installation CD into the CD-ROM drive and mount that drive. For most RedHat Linux installations, you can mount the CD-ROM using:
   `# mount /mnt/cdrom`

   If this does not work, see "Managing CD-ROMs" on page 68.
3. Go to the directory on the CD-ROM:
   `# cd /mnt/cdrom`
4. Run the installation script by typing:
   `# ./install`

   The script begins running. The installation is similar to the main WebAccelerator installation you performed.
5. You see the initial dialog on your screen that includes F5's technical support e-mail address and phone number.

   The installation script checks to see if your machine meets the required software dependencies. See "Installation Requirements" on page 6 for a list of these dependencies. See "Meeting Minimum Software Requirements" on page 21 for information on what to do if your system does not meet the minimum software requirements.
6. You are shown your installation options. An WebAccelerator and a WebAccelerator Remote are the same install, the WebAccelerator alone, without Management Console. Enter **2**.
7. As the installer begins to install the WebAccelerator software, it prompts you for information specific to your system. Get your completed worksheet from Worksheet 2 WebAccelerator Installation Planning on page 15 and use it as you answer this next series of prompts. The default answers work for most WebAccelerator systems.

8. Enter the hostname for the primary Management Console with which the WebAccelerator or WebAccelerator Remote must communicate. You defined the hostname for the primary Management Console in Step 9 on page 19 when you installed your WebAccelerator or configured your appliance. See the value you selected for question i) Hostname for Management Console on your worksheet. The worksheet contains additional explanation.
9. You are shown the disk partitions for the machine. Enter the directory where you want the WebAccelerator to be installed. The default is `/opt/pivia`. See the value you selected for question ii) Directory for WebAccelerator software to be installed on your worksheet.

   Note: If any directory you specify does not exist, the installer can create it for you. In most installations, you can simply use the default directory names.
10. Enter the directory where the WebAccelerator can store the cached data. The default is `/opt/pivia/hds`. See the value you selected for question iii) Directory for cached data on your worksheet.
11. Enter the directory where the WebAccelerator can store the log files. The default is `/opt/pivia/log`. See the value you selected for question iv) Directory for log file storage on your worksheet.
12. You are asked if you want to be notified by email whenever a WebAccelerator process, such as the WebAccelerator, exits abnormally. If you answer yes, you are then prompted to enter the email address to which these notifications should be sent. See the value you selected for question vi) Email address on your worksheet.
13. The installation is complete. You are returned to the main menu, where you must quit the installation. You cannot install a second WebAccelerator on a machine on which one already resides.
14. You are shown the installation summary. Make note of any information you might want to reference in future. The installation summary information is stored in a log file called `pvinstall.log` in your root `/tmp` directory. We recommend that you save this file as a record of your selections during installation.
15. After the summary, unless there was an error during the install, you are asked if you want to start the WebAccelerator processes. Enter **`y`** for yes, unless there is some reason you want to postpone starting them.
16. Eject the installation CD using the command:
    `eject /mnt/cdrom`

    If this command does not work, see "Unmounting CD-ROMs" on page 69 for additional instructions.

Your installation is complete.

Uninstalling Accelerators

To uninstall an Accelerator or WebAccelerator Remote for any reason, use the `uninstall` script found in this directory on the machine on which you installed the Accelerator:
`/opt/pivia/uninstall`

For more information on using this script, see "Uninstalling WebAccelerator Software" on page 22.

Chapter 6

Upgrading Your WebAccelerator

Upgrade Restrictions
Upgrade Process

This chapter describes how to upgrade your existing WebAccelerator installation to the latest release of the software. Use the upgrade process only when you are upgrading a machine that already has WebAccelerator software on it. If a machine does not have a Management Console or an Accelerator already installed, use the installation process described in Chapter 3, Installing Your WebAccelerator or Chapter 5, Installing Accelerators. Do not upgrade a F5 appliance yourself.

If you are upgrading to add a standby Management Console to a machine that already has an Accelerator, F5 recommends you first uninstall the Accelerator, following the instructions in "Uninstalling WebAccelerator Software" on page 22, and then do a new install of your standby Management Console and Accelerator, following the instructions in "Installing a Standby Management Console" on page 23.

Upgrade Restrictions

Before upgrading your WebAccelerator, note these restrictions:

- you must upgrade the machine running the Management Console before you upgrade the machines running your other WebAccelerators
- after you upgrade your Management Console, you cannot publish policies or perform cache invalidation until you have completed upgrading your other Accelerators
- you must manually merge in any custom configuration you made to `/opt/pivia/dac/conf/pvsystem.conf`

  The upgrade script only partially preserves changes you made to `pvsystem.conf`. Information such as host and port information that is captured in `/etc/pivia.manifest` is reflected in your upgraded system. For other information, such as maximum thread counts or compile queue settings, default values are used. Your customized values are not set in the new `pvsystem.conf`.

  The upgrade script makes a copy of the `pvsystem.conf` file before it overwrites it. However, you might want to make your own copy of the `pvsystem.conf` file that you can refer to after the upgrade script is finished.

Upgrade Process

The upgrade process is similar to the installation process described in Chapter 3, Installing Your WebAccelerator. To upgrade, you run the same installation script, but you choose the Upgrade option from the Main Menu.

To begin the upgrade:

1. Log in as root to the RedHat Linux machine that runs the Management Console for the F5 WebAccelerator you want to upgrade. You must upgrade this machine first.
2. Place the WebAccelerator installation CD into the CD-ROM drive and mount that drive. For most RedHat Linux installations, you can mount the CD-ROM using:
   `# mount /mnt/cdrom`

   If this does not work, see "Managing CD-ROMs" on page 68.
3. Go to the directory on the CD-ROM:
   `# cd /mnt/cdrom`
4. Run the installation script by typing:
   `# ./install`

   The script begins running. The prompts follow these conventions:

- If a prompt has a limited set of valid replies, such as `yes`, `no`, `1`, `2`, or `q` for quit, the reply options are either listed in the menu above the prompt or indicated in parentheses ( ) on the prompt line. For example, this prompt accepts only `y` or `n` as a valid response:
  ```
  Create this directory? (y/n) [y]:
  ```
- The option value displayed in brackets next to any prompt is the default. This default value is used if you simply press Enter in response to a prompt. For example, the default value for this prompt is `yes`:
  ```
  Answer yes or no (yes/no) [yes]:
  ```
- Use these special characters in response to a prompt to:

`!`	Launch `/bin/bash`
`q!`	Return to main menu
`#`	Print filesystem disk space usage
`?`	Print this help information

5. You see the initial dialog on your screen that includes the F5 Customer Confidence e-mail address and phone number.

   The installation script checks to see if your machine meets the required software dependencies. See "Installation Requirements" on page 6 for a list of these dependencies. See "Meeting Minimum Software Requirements" on page 21 for information on what to do if your system does not meet the minimum software requirements.
6. This same installation script is used for several different WebAccelerator product installations. You are shown a list of products and are prompted for which product you want to install now. Enter **`1`** for your F5 WebAccelerator.
7. You are shown your installation options. Enter **`3`** to indicate you want to upgrade.
8. The upgrade script prompts you for the information it needs to upgrade the software.
9. You are asked if you want to stop your WebAccelerator processes. Enter **`y`** for yes because you must stop the processes before you can proceed with the upgrade.
10. You are asked if you want to delete backup versions of directories from the previous installation. You can safely allow these to be deleted if you backed up any custom configurations you might have made in those directories. The only configuration file that you should have customized is `/opt/pivia/dac/conf/pvsystem.conf`, and this file is backed up automatically by the upgrade script.

    If you do not allow these directories to be deleted, make sure you have enough disk space to accommodate the new installation and the backed up directories.

11. The next part of the onscreen dialog is important. It tells you the location of your previous `pvsystem.conf` file. You need this if you made any manual changes to it. Press Enter after you have written down the location of the file.
12. Next, there are several prompts asking you about deleting the old versions and directories. Generally, you should answer yes.
13. The installation is complete. You are returned to the main menu, where you should quit the installation.
14. You are shown the installation summary. Make note of any information you might want to reference in future.
15. After the summary, you are asked if you want to start the WebAccelerator processes. Enter **`y`** for yes, unless there is some reason you want to postpone starting them.
16. Eject the installation CD using the command:
    `eject /mnt/cdrom`

    If this command does not work, see "Unmounting CD-ROMs" on page 69 for additional instructions.
17. Merge any manual changes from your previous `pvsystem.conf` file to your new `pvsystem.conf` file and your upgrade of this machine is completed. Remember, your previous `pvsystem.conf` file is located in the path you noted in Step 11, and your new `pvsystem.conf` file is located in `/opt/pivia/dac/conf/`.
18. If you have any WebAccelerators running on machines other than this one, you must upgrade them also so that they can work with the upgraded Management Console.

    Upgrade each of those WebAccelerators now, by running the upgrade script on each machine that has an WebAccelerator.

    Follow the steps, just as you did for the main upgrade, beginning with Step 1 on page 48, except that you log in to the WebAccelerator machine you want to upgrade, instead of the Management Console machine.

After you have upgraded all your WebAccelerators, your F5 WebAccelerator installation upgrade is complete. Go to "Configuring Your WebAccelerator" on page 25 for information on configuring your upgraded WebAccelerator.

Chapter 7

Confirming Your Installation

Check Processes
Check Proxy

Once you have installed and configured your WebAccelerator machines, you can minimally confirm that the software is functioning correctly. This chapter shows you how to check each machine.

Your initial configuration, logging into the Admin Tool and setting up policies and host mapping, confirms that a large portion of your WebAccelerator is operating correctly. This additional confirmation checks to ensure that your WebAccelerator software is operating correctly and that your initial policies and host maps were selected and entered without error.

Check Processes

Check to make sure that the WebAccelerator processes are running. Log in as root to the machine on which you have installed your WebAccelerator and issue the command:

```
# /opt/pivia/dac/rc/watchdog_client
```

You should see that these processes are up:

- ccm_pri
- ii
- tomcat
- log_collector
- log_pusher
- pvac
- hds_prune

If you are checking that a standby Management Console is running, log in as root to the machine on which you installed the standby Management Console and issue the command:

```
# /opt/pivia/dac/rc/watchdog_client
```

You should see that these processes are up:

- imon
- log_pusher
- hds_prune

If you installed a F5 appliance, you can check the processes by connecting a console to the appliance, starting a terminal emulator, and selecting the View Swan Labs WebAccelerator Processes option from the WebAccelerator menu that appears. This runs the `watchdog_client` script for you.

Check Proxy

Make sure your WebAccelerator is capable of proxying to your origin servers for fresh content. Because you have not yet reconfigured DNS for this installation, this step involves manually changing `/etc/hosts` files to override the information in DNS. Perform these steps:

1. First, choose one of the applications for which you defined a profile during configuration ("Creating Application Profiles" on page 34). For this example, let's say you choose your main web site application, located at `www.somesite.com`.

Look at its profile and see which cluster you assigned to it. Choose an WebAccelerator from this cluster and note where it is located and what its IP address is.

You must access files on this WebAccelerator, whether it is an appliance or a machine on which you installed the software. You can either connect a terminal or console to the machine, or remotely log in as root using `ssh`. If the WebAccelerator is an appliance and you connected a terminal to it, use the Run Interactive Shell menu option to access files on the appliance.

2. On a machine where you run a browser, edit the `/etc/hosts` file and change it such that the hostname used by your users for your web site application points to the IP address of that WebAccelerator machine.

   For example, your web site is accessed at the domain `www.somesite.com` and the WebAccelerator is on a machine with IP address 11.1.11.3. You change the line in your `/etc/hosts` file to:

   ```
   11.1.11.3   www.somesite.com
   ```

   All network traffic from this machine for `www.somesite.com` subsequently goes to the machine on which the WebAccelerator is installed.

   Under Windows 2000 and Windows XP, the `hosts` file can be found at:
   `c:\WINNT\system32\drivers\etc\hosts`

3. Go to the WebAccelerator machine to edit its `/etc/hosts` file. When you performed host mapping in your application profile for your web site, you identified at least one Destination Host address. Provide an entry for this Destination Host address.

   For example, if your web server for `www.somesite.com` is located at IP address 11.1.11.2, and your host map includes a Destination Host address of `ows.somesite.com`, enter this in the `/etc/hosts` file on the WebAccelerator machine:

   ```
   11.1.11.2   ows.somesite.com
   ```

   Your WebAccelerator should now proxy to the correct location for fresh content when it gets a request with domain `www.somesite.com`.

4. Back on the web browser machine in Step 2, request a page from `www.somesite.com`. If everything is working correctly, you see the page you would have received if your browser had accessed your origin servers directly.

   If you receive an "Access denied by intermediary" error response, check `/etc/hosts` and the host map for the application. Be sure you used a domain in your request that matches to a Requested Host in the host map which maps to the Destination Host you entered in the WebAccelerator's `/etc/hosts` file.

   If your browser times out the request, either the WebAccelerator software is not running on the machine, or your firewall is blocking access to port 80 on your WebAccelerator machine.

5. Repeat Step 2 through Step 4 for each WebAccelerator that is part of the cluster for that application.

   If you are going to ask for the same page for each WebAccelerator test, remember to clear your browser cache before requesting the page. See your browser documentation for information on how to clear its cache.

6. If you have more than one cluster, choose a new application that uses a different cluster, and repeat all the steps until every machine with an WebAccelerator has been tested.

7. Once you have confirmed that your installation is working to this level of minimal correctness, remove all the entries that you changed or added to the `/etc/hosts` files on your web browser and WebAccelerator machines.

You can now confidently reconfigure your network to place your WebAccelerator installation into your production environment. You create DNS entries, or use some other method depending on your topology, to point requests for `www.somesite.com` to the WebAccelerators, and redefine your origin servers so that they are no longer accessed with `www.somesite.com` but instead with `ows.somesite.com`. See Chapter 8, Configuring Your Network for more information.

Chapter 8

Configuring Your Network

Network Changes
Configuring Your Firewalls

After you have confirmed your WebAccelerator software installation, you must reconfigure your network to place your WebAccelerator into production use. Without this critical step, the WebAccelerator cannot find your origin web servers, and requests are not directed correctly to the WebAccelerator.

Network Changes

To change your network to support the WebAccelerator, you must:

- define the hostnames used by the Management Console and clusters so that they are resolvable to an IP address
- reroute network traffic to your WebAccelerator instead of the machine on which your web application is running
- define new hostnames for that your origin servers can use, now that network traffic for their former hostnames is being routed to the WebAccelerator
- if you have WebAccelerator Remotes installed, configure the local DNS servers at those locations so that HTTP requests originating there are directed to the local WebAccelerator Remote and not to your WebAccelerator
- if you have WebAccelerator Remotes installed, configure the global DNS server so that HTTP requests outside the WebAccelerator Remote locations are directed to the WebAccelerator

For specific information on how to configure your DNS servers, see the documentation that came with your DNS software.

Defining WebAccelerator Hostnames

When you installed the WebAccelerator, you defined the hostname of the Management Console for the WebAccelerator as part of the installation. During your configuration steps, in Step 6 on page 32, you specified public addresses for each of your clusters. You must be sure that each of these addresses, for the clusters and the Management Console, is resolvable to an IP address so that requests and other communication can be routed to them. This might involve changing DNS, reconfiguring load balancers, updating your VIP (virtual IP) system, or some other network change depending on your network topology.

The default value for the Management Console hostname is the local hostname of its machine. This local hostname might already be defined in DNS, so if you chose the default, you might not need to do any additional configuration.

The Management Console hostname is used by WebAccelerators and WebAccelerator Remotes to find the Management Console. It is also used to access the Administration Tool over the web so you can perform configuration. You might have chosen a logical, machine-independent value for the hostname, especially if you installed a primary and several standby Management Consoles. You might not want administrators to use a machine-dependent name to access the Administration Tool, when they could use `admin.SL400.yourcompany.com`, for example.

If you selected a logical hostname for the Management Console, you must ensure the hostname can be resolved to an IP address. For example, you might create an entry in

your DNS server that defines the hostname and the IP address of the Management Console to DNS.

You must perform similar steps, depending on your network setup, to ensure the public addresses you specified for your clusters are also resolvable to an IP address. These addresses are used by WebAccelerator Remotes to forward requests to the correct cluster handling the requested application.

Routing Requests to the WebAccelerator

You must reconfigure your network to direct requests for your web applications to the WebAccelerator clusters handling the application. This might involve changing DNS, or changing your load balancer, or updating your VIP (virtual IP) system.

For example, suppose your web site is normally accessed using:
`http://www.somesite.com`

You configured the WebAccelerator with a cluster called Somesite that contained three WebAccelerators, which are behind a load balancer. In the application profile for your web site, you picked Somesite as the cluster.

In this case you want to reroute traffice so that requests to `www.somesite.com` go to the load balancer in front of Somesite and update the load balancer to balance the requests between the three WebAccelerators.

You also have an internal web-based financial application that is defined to the WebAccelerator in the application profile Expense Reports. This profile specifies cluster ER, composed of one WebAccelerator, as handling all the requests for this application. You can change DNS to route all requests for this application to the IP address of the WebAccelerator in cluster ER.

Accessing your Origin Servers

You must also create new entries in DNS to allow your origin web servers to be accessed. For example, now that `www.somesite.com` is defined as pointing to the load balancer in front of the Somesite cluster, you need to create a new hostname for your origin web server. You might use this hostname:
`http://ows.somesite.com`

for your origin web server, and use the IP address of its machine when defining the hostname. This hostname should also appear in the host maps of the WebAccelerator, so the WebAccelerator can use this hostname when proxying requests to your origin web server.

Routing Requests to the WebAccelerator Remote

You must reconfigure the local DNS to direct requests for your website at the remote office to your WebAccelerator Remote installation. In turn, your WebAccelerator Remote directs requests to your WebAccelerator, using the host map and cluster information to send requests to the appropriate cluster.

For example, suppose your web site is normally accessed using:
`http://www.somesite.com`

and you use a local DNS server at the remote office to handle routing. In this case you must change that local DNS server by adding an entry with the `www.somesite.com` domain and the IP address of the WebAccelerator Remote machine, so that requests to `www.somesite.com` go to the WebAccelerator Remote server.

You do not make this change to the global DNS server, because you want only local requests to go through the WebAccelerator Remote. The global DNS server should still be directing `www.somesite.com` requests to the appropriate WebAccelerator cluster.

For the global DNS server, ensure that there is also an entry for the hostname of the Management Console or some way to resolve the hostname to an IP address. You have probably already done this when you changed your network to support the WebAccelerator (see "Defining WebAccelerator Hostnames" on page 56). This hostname should be the hostname you specified during your WebAccelerator Remote install. This hostname is used for peer-to-peer communication between the WebAccelerator Remote and the Management Console.

Requests are routed as follows:

1. A user in the remote office sends a request to `www.somesite.com` that requires fresh content.
2. The local DNS server resolves it to the WebAccelerator Remote machine, and directs the request there.
3. The WebAccelerator Remote determines which application the request is for, and that the application is being handled by cluster Somesite. It applies the remote policies to the request and sends it to the public address for cluster Somesite.
4. The local DNS server does not have an entry for this hostname, so it is passed to the global DNS server, which directs the request to the cluster, where it is passed to an WebAccelerator within the cluster.
5. The WebAccelerator handles the request and if needed, proxies it to the appropriate origin web server, based on its own host mapping.

123 Corporation Usage Scenario

A company, 123 Corporation, has its headquarters in the United States, with a manufacturing division in China. All divisions of the company need to access certain applications and documents through the internal corporate web server located at headquarters. The main internal corporate website is `inside.123corp.com`. There is a WebAccelerator system installed at headquarters to accelerate this application, called Inside. In that system, an WebAccelerator cluster called 123HQ has been defined to handle all the requests for the Inside application. This includes the WebAccelerator on the Management Console machine. Requests to the WebAccelerators in the cluster are routed through a load balancer that has a public address of `10.1.10.3`. When the WebAccelerator was configured, the cluster's public address was specified as `10.1.10.3`. There is also a WebAccelerator Remote installed at the company's offices in China.

The network has been reconfigured as follows:

- the internal corporate web server at headquarters has a new DNS record so that it is now `apps.inside.123corp.com` instead of `inside.123corp.com`.
- the 123HQ cluster is defined in the global DNS server as `inside.123corp.com` and resolves to the IP address of the load balancer.
- there is also a hostname for the Management Console that was defined when it was installed, `webacmc.123corp.com`. This hostname also has an entry in the global DNS server, resolving to the IP address of the Management Console machine.

When a user in a sales office in the U.S. sends an HTTP request to `inside.123corp.com`, the global DNS server needs to route that request to cluster 123HQ at headquarters. However, users in China also send their HTTP requests for corporate information to `inside.123corp.com`. These requests need to be directed by their local DNS server to their local WebAccelerator Remote, not to headquarters. The local DNS server, at the China office, needs an entry for hostname `inside.123corp.com` that points to the IP address for the local WebAccelerator Remote machine. This local entry needs to override the entry for `inside.123corp.com` in the global DNS server.

The WebAccelerator Remote has received all the policies, application profiles, cluster information, and host mappings from the Management Console. When it handles requests, it can use this information to determine where to proxy to handle a request.

The diagram shows how a request is serviced for clients in the China office and for other clients outside of headquarters:

Figure 1 Example of DNS Configuration with WebAccelerator Remote

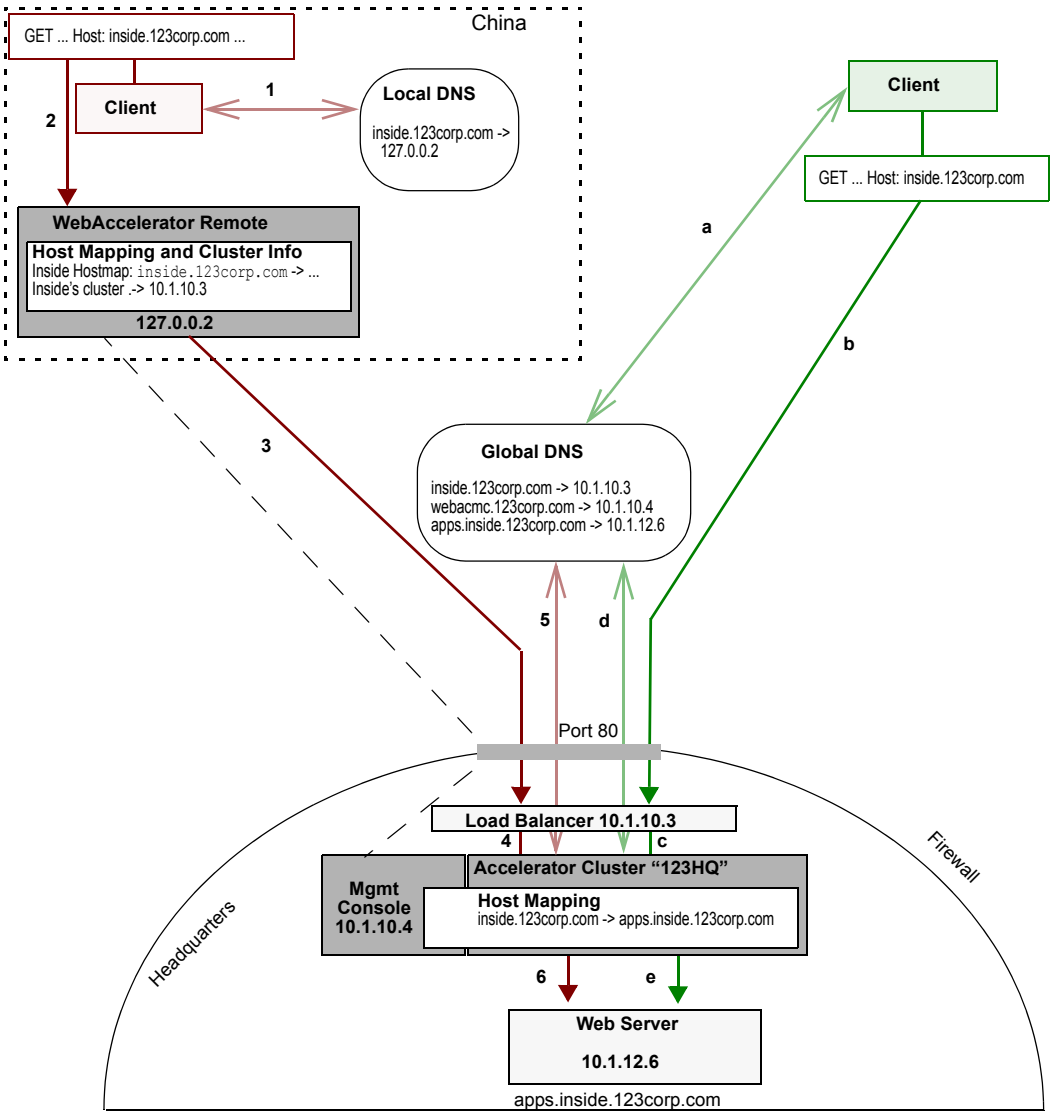

Request from Remote Location Client

1. A client browser in China creates an HTTP request for a new document from the corporate web server. The browser queries the local DNS server for the hostname. The local DNS server has an entry for `inside.123corp.com` and points the

browser to the WebAccelerator Remote machine. This local DNS entry must override any entry for `inside.123corp.com` in the global DNS server.

2. The client browser sends the request to WebAccelerator Remote.
3. WebAccelerator Remote processes the request, looks through all its host maps and finds `inside.123corp.com` as a Requested Host in the host map for the Inside application.

   The WebAccelerator Remote determines it must proxy for content. In its application profiles and cluster information, it finds that the Inside application is assigned to cluster 123HQ, which shows a public address of 10.1.10.3. It proxies the request to that IP address.
4. The load balancer receives the request and sends it to one of the WebAccelerators in the cluster.
5. The WebAccelerator processes the request, looks through all its host maps and finds `inside.123corp.com` as a Requested Host in the host map for the Inside application. The WebAccelerator knows it is assigned to handle requests for the Inside application.

   The WebAccelerator determines it must proxy for content, and queries the global DNS server for the destination hostname that maps to `inside.123corp.com` in its host map, which is `apps.inside.123corp.com`. The global DNS server provides the IP address for the origin web server.
6. The WebAccelerator proxies to the Web Server for content, and the response follows a similar path back to the client. Note that if the document is cacheable, the next time a client in the China office requests the document, the WebAccelerator Remote can provide the document directly to the client.

Request from Client Outside Headquarters

a) A client browser at a U.S. sales office creates an HTTP request for a new sales document from the corporate web server. The browser queries the global DNS server for the hostname. The global DNS server has an entry for `inside.123corp.com` and points the browser to the load balancer at headquarters.
b) The client browser sends the request to the load balancer.
c) The load balancer receives the request and sends it to one of the WebAccelerators in the cluster.
d) The WebAccelerator processes the request, determines it must proxy for content, and queries the global DNS server for the destination hostname that maps to `inside.123corp.com` in its host map, which is `apps.inside.123corp.com`. The global DNS server provides the IP address for the origin web server.
e) The WebAccelerator proxies to the Web Server for content, and responds to the client. Note that if the document is cacheable, the next time any client requests the

document, the WebAccelerator can provide the document directly to the client, without proxying to the Web Server.

Configuring Your Firewalls

Your network firewalls must be reconfigured to allow WebAccelerator communications and traffic. After installing the WebAccelerator, there are some ports you must open. There are also network requirements related to the administration and monitoring of your WebAccelerator installation.

Note: Your origin servers must continue to be accessible for any web traffic that the WebAccelerator can proxy or redirect to them.

WebAccelerator Ports

You need to open ports for the Management Console and any Accelerator that is on the same machine as the Management Console. This table summarizes the WebAccelerator's traffic requirements and the default ports that it uses to satisfy those requirements.

Port	Purpose
8080 or 8443	Needed to access the Admin Tool. The actual port depends on whether you want to access the UI using HTTP (8080) or HTTPS (8443). You can cause the UI to be available on a nonstandard port if you choose, in which case these ports might not be needed. Access to this port is only required by the personnel responsible for administering the WebAccelerator installation.
80 / 443	Needed for client requests to your web application.
12400, 12402, 12403, 12406, 12407, 12408, 12409	Default ports used for SNMP monitoring of the various WebAccelerator processes. These ports must be open only if you are using SNMP. Access is required only for the IP addresses where your network monitoring software is running.
161/udp 162/udp	Default SNMP ports used for only for monitoring the operating system.

See the documentation that came with your firewall for information on how to open ports.

Note: Using SNMP to monitor your WebAccelerator is an optional activity. You only need to open SNMP ports if you are going to perform SNMP monitoring. For information on how to use SNMP with your WebAccelerator, see the *Administration Guide*.

If you have a standby Management Console, you need to configure the ports the same as for any other WebAccelerator. Use the table above.

WebAccelerator Ports

If your installation has Accelerators on different machines than the Management Console, there are additional ports you must open. They are:

Port	Purpose
80 / 443	Needed to service client HTTP/S requests.
12202 or 12006	Used for communications between the Management Console and WebAccelerator machines. By default, this traffic is encrypted (SSL), and by default the encrypted traffic uses ports 12202. If you choose to use unencrypted traffic for internal WebAccelerator communications, then by default port 12006 is used.
22	Needed for ssh traffic from the Management Console to the WebAccelerator machine. This is used to transport configuration information to the individual WebAccelerator machines.
12406, 12407, 12408, 12409	Default ports used for SNMP monitoring of the various WebAccelerator processes. These ports must be open only if you are using SNMP. Also, access is only required for the IP addresses where your network monitoring software is running.
161/udp 162/udp	Default SNMP ports. Used for only for monitoring the operating system.
123/udp	Required for NTP traffic between the Management Console and the WebAccelerators.

Appendix A

Linux-Specific Information

Linux Installation Planning
Operating System Installation
Setting up Red Hat up2date
Managing CD-ROMs

The F5 WebAccelerator must be installed on RedHat Linux system, which means you must install Redhat Linux before you install the F5 WebAccelerator.

RedHat Linux's operating system installer is friendly but flexible. You are presented with several installation choices. Your choices can affect how well the F5 WebAccelerator works. This chapter guides you through the key decisions that impact the subsequent performance of your F5 WebAccelerator.

Linux Installation Planning

Before performing the installation, gather this information:

Interface:	Eth0 / Eth1
Hostname:	
IP Address:	
Net Mask:	
Gateway:	
DNS 1:	
DNS 2:	
DNS 3:	

Operating System Installation

You install Red Hat Enterprise ES using CDs 1 and 2 from the CD set. Place CD 1 in the CD-ROM drive and boot the machine.

F5 recommends you make these choices during the installation process:

1. When you are asked to select an Installation Type, select 'Custom'. This allows you to customize the packages that are placed on your machine for this installation.
2. For Partition Strategy, select a manual partition strategy. This allows you to strategically separate data created by your WebAccelerator onto individual disk partitions. Red Hat and F5 recommend that you select "Manually partition with Disk Druid" for custom disk partition installations such as this.
3. When you are asked about a boot loader, choose lilo.
4. When asked about a firewall, choose 'No Firewall'. Firewall support should be provided by some other device in your infrastructure.
5. For time zones, ensure all of the machines participating this WebAccelerator installation are using the same time zone.
6. For user account configuration, create only those user accounts that you need for the administration of this machine. The WebAccelerator installation software creates the user and group required to run the WebAccelerator.
7. For Package Selection:

a. Turn off everything except for:
   - "XWindow System"
   - "Network Support"
   - "Server"

b. Select "Individual Packages".

c. Select "Flat view".

d. Make sure that "rhn_register" and "up2date" are selected.

e. Select 'kernel-smp' and deselect 'kernel'.

f. Make sure your preferred editor is on the system. By default, emacs is not installed and vi is installed.

When you are done, examine the packages the installer wants to put on the machine in order to resolve package dependencies, if any. If the installer is putting anything on the machine that you do not want there, use the back button to go back and deselect the packages that are forcing the unwanted software onto the machine. When you are done, allow the installer to resolve package dependencies if necessary.

8. When you see the 'Customize X' configuration screen, allow the installer to configure your X server only if you need it to help administer the machine. This step is not required for the WebAccelerator.

   If you do choose to configure your X server, make sure that the installer does not cause the X server to automatically start when the system boots.

When you are done, the installer installs the operating system and all the packages you selected. Upon completion of this process, you must reboot your machine.

You can now proceed with your WebAccelerator installation.

Setting up Red Hat up2date

up2date is a service that Red Hat provides to allow you to easily keep your machine current with the latest patches. To set up up2date:

1. Visit this URL and register your copy of Red Hat Linux Enterprise Workstation: http://www.redhat.com/apps/activate/

   If you are not a registered user of Red Hat Network, you are asked to create an account (username and password) at this time.

2. Update the up2date and rhn_register certificate (the certificate initially installed on your system expired on August 8th, 2003).

3. Register your system with Red Hat Network. Run this command as root on the system you have just installed:
   `[root@localhost] rhn_register`

   Note that rhn_register prompts you for the username and password from Step 1.

Once you have registered your system, there are several ways that you can use the Red Hat Network. The mechanism that F5 recommends is that you allow Red Hat Network to automatically update your system. There is a daemon (rhnsd) that pulls and applies updated packages for you, and this is running by default. However, you first have to log into RHN and indicate that you want updates applied automatically to your system.

1. Log into the Red Hat Network site:
   http://rhn.redhat.com
2. Click on the `systems` tab.
3. In the System Overview screen, locate the system that you just registered and click on its name.
4. Under `System Properties`, click on the `"Edit these properties"` link.
5. Select `"Auto Errata Update"`.
6. Click the `Update Properties` button at the bottom of the page.

Your system now automatically updates itself with any relevant packages that may be made available. Note that your system only downloads and installs updates to packages that already exist on the machine.

Managing CD-ROMs

Depending on how your RedHat Linux system is put together, there are several methods that you might use to mount a CD-ROM.

Mounting CD-ROMs: Method 1

The most common method is to simply use:
`mount /dev/cdrom`

If you see something like this in your `/etc/fstab` file, then this method works:
`/dev/cdrom /mnt/cdrom iso9660 noauto,owner,kudzu,ro 0 0`

Under these circumstances, after mounting, the CD-ROM is available under `/dev/cdrom`.

Mounting CD-ROMs: Method 2

If your CD-ROM drive is not identified in your `/etc/fstab` file, you can manually mount your CD-ROM using this command. Note that before this can work, the mount location (`/mnt/cdrom` in this example) must already exist:

```
mount -t iso9660 -o ro /dev/cdrom /mnt/cdrom
```

Mounting CD-ROMs: Method 3

In some rare cases, your RedHat system might have been installed such that the `/dev/cdrom` symbolic link does not exist. In this case, you must mount your CD-ROM using the device identifier assigned to it when your system boots. This device identifier is stable across reboots (unless you change your hardware configuration in some way), so once you have found your CD-ROM device identifier you should not have to look for it again.

Your CD-ROM device identifier is dependent on whether your CD-ROM drive is connected as a SCSI or IDE device. It is further dependent on the CD-ROM drive's position on the device bus or chain. For example, a CD-ROM drive installed as the master device on the first IDE bus in your system has the device identifier:

```
/dev/hda
```

A CD-ROM drive installed on the second position of a SCSI chain has the device identifier:

```
/dev/sdb
```

You can find your CD-ROM device identifier by examining your system's boot log. Use the `dmesg` command to examine this information. Generally, a CD-ROM drive identifies itself as a CD-ROM or DVD/ROM device. Try using `grep` to search for the word 'ROM' in the `dmesg` output for the device:

```
dmesg | grep ROM
```

The result of this command might be:

```
hda: TEAC CD-ROM CD-224E, ATAPI CD/DVD-ROM drive
hda: ATAPI 24x CD-ROM drive, 128kB Cache, UDMA(33)
```

The device is on `/dev/hda`, so mount the drive using:

```
mount -t iso9660 -o ro /dev/hda1 /mnt/cdrom
```

Notice the `1` appended to the device identifier. This indicates that you want to mount the first partition on the device. Typically, CD-ROMs only have one partition on them, so you usually append `1` to the device identifier.

Unmounting CD-ROMs

You use either the `umount` or `eject` command to unmount a CD-ROM. The `eject` command, if it is installed, should both unmount and physically eject the CD-ROM from the drive.

Use the directory that you mounted your CD-ROM to as the argument to the `umount` or `eject` commands. For example:

```
mount -t iso9660 -o ro /dev/cdrom /mnt/cdrom
eject /mnt/cdrom
```

Note that if the attempt to unmount the CD-ROM results in a 'device is busy' warning, ensure there are no applications accessing anything under the mount directory before attempting the unmount again.

Index

A

abnormal exits
  email address 14
  notification 20
Accelerator
  clusters 3
  configuration with multiple 3
  identifying 42
  installation 43
  introduction 2
  planning 15
  refreshing content 39
  uninstall 22
accessing the Admin Tool 30
accounts 7
address
  cluster 32
  destination host 37
Admin Tool
  advanced use 42
  basics 26
administrator account
  changing password 39
  logging in 31
appliance, WebAccelerator Remote 9
application matching 8
applications 3
  initial steps 34
  profile 34

B

browsers
  Internet Explorer requirements 6
  Netscape 7

C

cache, clearing 39
CD-ROM
  mounting 68
  unmounting 69
changing
  password 39
  time zone 39
clusters 3
  address 32
  assigning to applications 34
  initial steps 31
  options 32
connection timeout, destination host field 29
content type 8
content, refreshing 39
content-based object identity 8
conventions in prompts 18

D

de-install 22
deployment planning 12
destination host
  connection timeout field 29
  description 27
  lifetime field 29
  properties 28
  read write field 29
  retries field 29
  security field 28
  valid hostname 27
directing incoming requests 26
directories
  partitions 6

selecting for install 19
disk partitioning 6
DNS
configuring for WebAccelerator 56
defining hostnames 56
directing requests to WebAccelerator 57, 58
Express Connect configuration 30
domain mapping 26

E

Express Connect 29, 38
number of subdomains 29
subdomain prefix 30, 38

F

figures
DNS with WebAccelerator Remote 60
multiple Accelerators 3
multiple clusters 4
single WebAccelerator 2
single WebAccelerator Remote Edition 5
files
editing /etc/hosts 53
install log 20
log, rotating 41
pvsystem.conf upgrade considerations 48
upgrade backups 49
firewalls
Accelerator ports 63
configuring 62
Management Console ports 62

H

handling unmapped hosts 32
host map 2
creating 35
host mapping
description 26
verifying 52
hostname
specifying 27
used by Accelerator 45
used for Management Console 13
verify mapping 53
wildcards 27
hosts, options for unmapped 32
HTTP ports 62

I

installation
conventions 18
overview 12
prerequisite software 21
solo Accelerators 43
standby Management Console 23
steps to follow 18
uninstall 22
upgrading 47
WebAccelerator Remote 43
worksheets 12
Internet Explorer, required version 6

L

lifetime, destination host field 29
Linux
installing 66
logrotation facility 41
MAC Address 42
requirements 6
load balancer, basic configuration 3
log files
pvinstall 20
rotating 41

M

MAC address 42
Management Console
defining hostname to DNS 56
install overview 12
introduction 2
primary
installing 19
running 52
setting hostname 13

standby
installation 23
planning 16
running 52
uninstall 22
mapping
domains and hosts 26
verification 52
monitoring Accelerators 7
monitoring performance 40
mounting CD-ROMs 68

N

Netscape support 7
normalized URLs 8
NTP ports 63

O

object classification 8
object IDs 8
options
destination host 28
Express Connect 29
Linux install 66
organizations 7
OS requirements 6

P

partitioning disks 6
partitions full 41
password, changing 39
Performance Reports 40
pivia, defined 9
planning for installation 12
policies
assigning to applications 34
changing set 38
ports, firewall 62
preferences, setting 39
processes, confirm status 52
profile, application 34
profiles 3
proxy
checking 52
specifying hostname 26
proxy server
forwarding unmapped host requests 32
options 32
proxying for content 2

R

randomized URLs 8
read write, destination host field 29
remote Accelerator policies 35
removing WebAccelerator software 22
requested host
description 27
wildcards in hostname 27
requests, routing 57, 58
requirements
partitioning 6
planning 12
software 6
responses, enhanced 8
restrictions on upgrade 48
retries, destination host field 29
rotating log files 41

S

security
destination host field 28
normalized URLs 8
SNMP ports
Accelerator 63
Management Console 62
specifying hostnames using wildcards 27
system monitor health checks 7

T

time zone, changing 39

U

uninstall 22

unmapped requests
  described 30
  setting options 32
unmounting CD-ROMs 69
up2date 67
upgrade considerations 48
URL
  normalized 8
  randomized 8
  re-encoding 8
URLs, mapping to a server 26

W

watchdog_client
  confirm process status 52
  confirm standby Management Console 52
WebAccelerator
  DNS configuration 57, 58
  monitoring 40
  upgrade 48
WebAccelerator Remote
  basics 4
  installation 43
  setting policies 35
wildcards in hostname 27
worksheet
  installing additional Accelerators 15
  installing WebAccelerator 12