HP WebQoS Administration Guide

Edition 4

E0801

U.S.A.

Legal Notices

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

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

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

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

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

Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19(c)(1,2).

Contents

Contents

Contents

Printing History

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

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

First Edition: August 2001

1 Understanding HP WebQoS

This chapter explains the roles and benefits of HP WebQoS and how it fits into your environment. It also describes fundamental WebQoS concepts to help you understand how WebQoS technology works.

The Role and Benefits of HP WebQoS

HP WebQoS is an enhancement to your operating environment that stabilizes, optimizes, and prioritizes Internet-related applications and transactions that run on Windows 2000 and NT servers (Windows) and SPARC[1] servers (Solaris[1] operating environment). HP WebQoS works with web-enabled applications built on top of the iPlanet Web Server[1] version 4.1 (UNIX®) and Microsoft® IIS[2] (Internet Information Server) version 4.0 (Windows NT) and version 5.0 (Windows 2000).

HP WebQoS enables you to use your resources efficiently to deliver predictable and differentiated service levels for your web-based applications, based on the type of customer or transaction. It allows you to determine these service levels based on business policies.

Although the Internet represents significant opportunities for expanding your business, the following risks may hamper your ability to deliver web-based services to customers and partners reliably:

- Unpredictable and possibly fluctuating demand for services.

  Because the Internet offers instantaneous access to your site, promotional, seasonal, or current events may instigate sudden and immediate interest in your services. Although the additional interest and business may be welcome, your server may not be able to handle the load if significant numbers of clients come all at one time. Or if you are a service provider hosting multiple sites on one server, you should make sure that the demands of one site do not restrict use by other sites on the same server.

- Heavier than expected system processing per user request.

  The actual workloads exhibited by each visitor to your site may not align with the workload models you generated during the capacity planning stage. Perhaps users are performing far more searches than you anticipated or are spending a lot more time performing compute intensive operations while at your site.

1. SPARC, Solaris, and the iPlanet Web Server are products of Sun Microsystems, Inc.
2. Microsoft IIS is a product of Microsoft Corporation.

- All requests are treated the same.

  During overload conditions, all user requests and all transactions are impacted. Even your most important customers who are trying to access their accounts to make purchases, or sales representatives who are trying to make a sale at the end of the month, are at the mercy of Web surfers.

Each of the above conditions can cause increased load at your site or may impede your ability to deliver your services at optimum levels.

By configuring WebQoS based on the type of services offered and traffic experienced at your site, you ensure that important web-based transactions are handled optimally in order to maximize revenue and customer satisfaction.

Key capabilities and benefits of WebQoS allow you to do the following:

- Stabilize the system during server overload situations.

  When the server is at capacity, WebQoS works to alleviate the load and improve the performance. For more information, refer to “Capacity Protection” on page 15.

- Prioritize client requests to the server system.

  WebQoS allows you to provide differentiating service to different types of client requests by prioritizing requests submitted to the site. For more information, refer to “Request Classifications” on page 14.

- Establish rules for the classification of service requests from users.

  WebQoS fulfills high priority service requests with optimum performance while lower priority service requests are handled in accordance with your instructions. For more information, refer to “Request Classifications” on page 14.

WebQoS Concepts

This section describes fundamental WebQoS concepts and capabilities. Understanding these concepts is necessary before you can understand how WebQoS technology works.

Service

A service refers to something of value that a business is offering to other businesses or consumers. For example, a business may provide email, messaging, documentation publishing, cataloging, and Electronic Commerce services. A WebQoS service is a way to logically group web sites related to a business application that you offer to customers or partners.

Request Classifications

Request classifications determine access priority for requests submitted to a web site. They enable you to give preferential treatment to your most important customers or transactions, allowing you to meet formal or informal service level agreements.

WebQoS request classifications differentiate requests as they enter the server system by dividing requests into categories based on application, client or destination IP addresses, destination port number, and URL document paths. These request classifications can be assigned different priority levels.

WebQoS supports three request classification priority levels: high, medium, and low. During periods of heavy system load high priority requests can be given preferred access to the server and low priority requests can be redirected or rejected. If a request is accepted, it is scheduled based on its request classification priority. Based on your configured policies, the request might be immediately processed, or it may wait while higher priority requests are processed first.

Sessions

A session is composed of one or more requests to a web site from the same user. These requests may arrive over one, or over several connections. Once a session has been granted access to a web site, a user's remaining requests are guaranteed to be forwarded to the web server, unless the session times out.

Under periods of heavy system load, new sessions may not be granted.

This is controlled by the WebQoS policies. When new sessions are rejected, redirected, or deferred, existing sessions continue uninterrupted. User request prioritization and session management are the primary tools used by WebQoS to provide capacity protection. The network manager is responsible for updating the timers that define a session, and for establishing the policies that control if the session is admitted.

Capacity Protection

Capacity Protection prevents system overload with prioritization and admission control. This minimizes the impact of unexpected surges in demand while maximizing the volume of completed transactions. Capacity Protection makes sure that performance levels for active customers and their currently running transactions are not compromised. New user requests are not admitted to a site unless their transactions can be completed quickly.

Capacity is assessed by service level objectives and thresholds. Protection is provided by corrective actions.

SLOs, Thresholds, and Corrective Actions

To ensure that your business policies are reflected in your Web applications, they need to be translated into service level objectives (SLOs) and thresholds. SLOs are business-oriented policies and thresholds are operations-oriented policies. When an SLO or threshold is violated, a configured corrective action is performed.

SLOs WebQoS has the following types of business-oriented SLOs:

- Capacity - Maximum number of concurrent sessions supportable with acceptable performance. Typically coupled with a response time SLO or performance threshold.
- Response times - A key component of customer satisfaction. Measured from the time a request enters the server to the time a response is sent back.
- External measurement - Measurements collected from external sources can be used to determine corrective actions.

Thresholds WebQoS uses an operations-oriented load threshold to keep your system operating properly during times of peak load. System load is assessed as a function of CPU (central processing unit) usage.

Corrective Actions Corrective actions help bring the violated SLO or threshold into compliance. You configure and prioritize corrective actions, which are executed when a violation occurs.

There are three types of corrective actions:

- Redirect requests to another system.
- Defer requests for a few moments until current sessions are completed.
- Reject low priority customers in extreme overload conditions.

How WebQoS Fits Into Your Environment

WebQoS works with your web server and application servers for all of your web-enabled applications. Refer to Figure 1-1, "WebQoS in Your Environment".

Figure 1-1 WebQoS in Your Environment

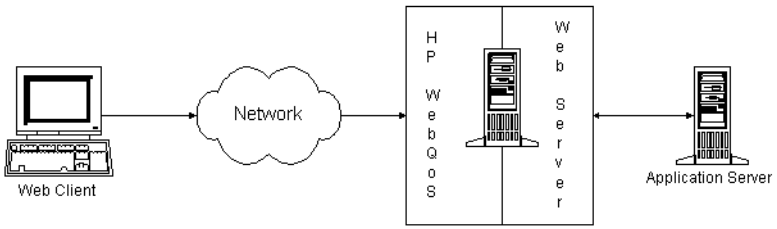

To ensure acceptable web-based interactions between you and your customers, all parts of your web environment (client, network, web server, and application server which includes an application) must work together efficiently and effectively. The network alone cannot deliver acceptable service levels if the server is creating delays due to excessive demand or component failures. The server alone cannot deliver acceptable service levels if the network is exhibiting bottlenecks.

The following describes briefly the function of each part of the environment:

- The client identifies itself and initiates requests.
- The network distinguishes class priorities and routes priority packets efficiently while delivering best effort service to lower priority packets.
- The web server distinguishes WebQoS request priorities and resource management policies to ensure high priority service requests are processed quickly.
- The application server contains an application with the processing rules that enables service requests to be classified and handled by the server and network.

WebQoS Components

WebQoS comprises a number of components that perform priority-based resource management and service request handling (see Table 1-2, WebQoS Components). In addition, management components allow you to configure WebQoS for the systems on which it is installed as well as define the SLOs and threshold policies that are important to your business.

Note that not all components are required on every WebQoS-enabled system but at least one of each component must exist in a WebQoS domain. Refer to the HP WebQoS release note for more detailed information.

Figure 1-2 WebQoS Components

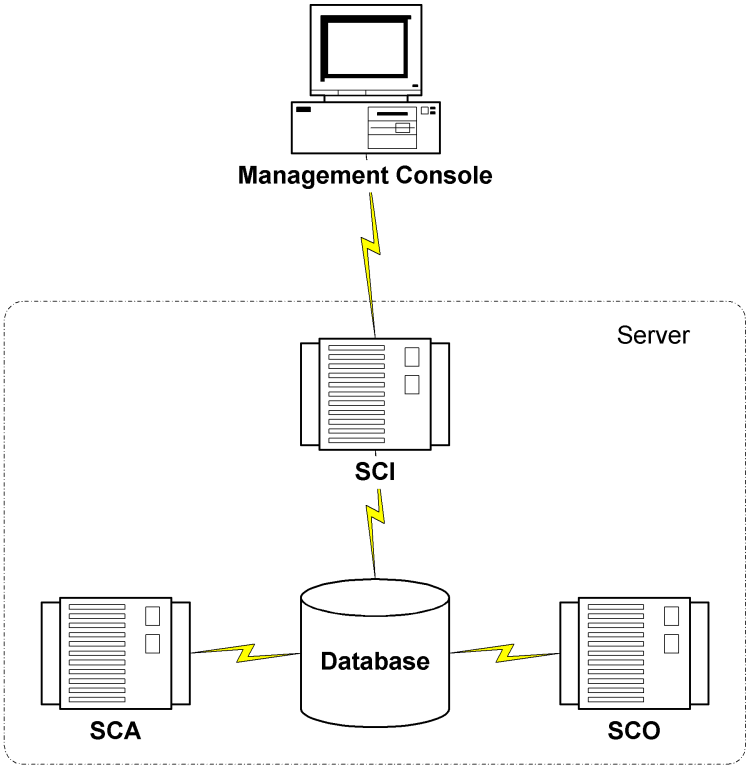

Component	Installed On	Description
Management Console	Any Windows NT/95/98/00 system. Can be installed on multiple PCs.	Allows you to define performance objectives for the components of the service and to monitor performance levels and SLOs.
Database	Any system*.	Stores configuration and event information, statistics, and component status. Resides on only one system in your WebQoS domain.
SCA (Service Control Agent) or **Web Services**	Any system* running a web server that is managed by WebQoS.	Monitors site level SLOs, classifies user requests, collects performance information, and performs corrective actions. Resides on at least one system in your WebQoS domain.
SCI (Service Control Interface) or **Communications**	Any system*.	Receives configuration information from the management console and API and monitors the SCO and SCA(s). Resides on at least one system in your WebQoS domain.
SCO (Service Control Operator) or **Management Server**	Any system*.	Monitors service level SLOs. Resides on only one system in your WebQoS domain.

*If you are running WebQoS on Windows 2000 or NT 4.0, any system is any Windows 2000 or NT 4.0 server or any Advanced Server system.

If you are running WebQoS on the Solaris operating environment, any system is any SPARC server running Solaris 8.

2 Setup

This chapter covers the initial steps you must complete before configuring policies and rules using WebQoS. Complete the following steps:

1. Install the WebQoS software.
2. Modify the WebQoS configuration files.
3. Run the WebQoS `setup` script and start your web server(s) (UNIX).

WebQoS Software Installation

Information about installing, upgrading, and removing WebQoS and its management console can be found in the HP WebQoS release note.

On a Windows system, the Microsoft IIS software should be installed, configured, and running on the web server system before WebQoS is installed.

See “Microsoft IIS” on page 81 for more information about using Microsoft IIS with WebQoS.

NOTE In this manual, the location where the WebQoS program files are installed for Windows is referred to as *<install-directory>*. The default installation directory is `C:\Program Files\Hewlett-Packard\HPWebQoS`.

WebQoS Configuration Files

The following configuration file can be modified before starting WebQoS:

- `qos.conf` - WebQoS global configuration file.

qos.conf

qos.conf is the WebQoS global configuration file. In UNIX, this file is located in the /etc/opt/webqos directory. In Windows, this file is located in *<install-directory>*.

On a UNIX system, configure qos.conf before running the WebQoS setup script. If you update qos.conf after running the setup script, individual components (for example, the web server, SCA, SCI, and/or SCO) must be restarted for any changes to take effect (see "Restarting WebQoS Components (UNIX)" on page 30).

On a Windows system, if you update qos.conf, individual components (for example, the web server, SCA, SCI, and/or SCO) must be restarted for any changes to take effect (see "Restarting WebQoS Components (Windows)" on page 30).

The following parameters can be set:

SCA and Web Server Parameter

The following parameter affects the SCA and web server. If you change this parameter, you must restart the SCA and web server.

- MeasurementInterval (seconds)

  Default: *30*
  Range: 5 - 43200 (12 hours)

  How often, in seconds, WebQoS collects performance measurements. The MeasurementInterval must be an even multiple of the MonitorInterval.

  If Site Statistics Logging is enabled from the management console, this parameter also specifies the interval at which the SCA generates site statistics.

Web Server Parameters

The following parameters affect the web server. If you change any of these parameters, you must restart the web server.

- EXCLUDE_FILE_EXTNS

  Default: *gif jpeg jpg*
  Range: Up to 10 extensions can be specified. Each extension can be up to 10 characters.

Extensions of file types that are typically embedded in a requested page and should not be considered new sessions. Any policies or rules should not be applied to them. For example, gif, jpeg, and jpg are graphic files that are typically embedded in a requested page. A separate request is generated for each of these files but these requests should not be considered new sessions. Also, any policies and/or rules configured are applied only to the requested page, not to these embedded files.

- `HTTPVersion`

  Default: *HTTP/1.1*

- `KeepDeferringAfterMaxDeferTime`

  Default: *TRUE*
  Range: TRUE or FALSE

  Based on the calculated defer time and the maximum defer time configured in the `Defer Corrective Actions` screen (default maximum defer time is 120 seconds). If the calculated defer time is greater than the maximum defer time and this parameter is set to TRUE, new sessions are deferred. If the parameter is set to FALSE, new sessions are re-evaluated (in accordance with the configured corrective actions).

- `MonitorCPUWeight` (%)

  Default: *20*
  Range: 1 - 100

  Percentage of the current CPU utilization (`pstat` is used to determine current CPU utilization) used to calculate the weighted CPU utilization. The weighted CPU utilization is based on the current and previous (last measured) CPU utilization:

  weighted CPU% = (`MonitorCPUWeight`% * current CPU) + [(100 - `MonitorCPUWeight`)% * previous CPU]

  For example, if you set the MonitorCPUWeight to 20, the current CPU utilization is 60, and the previous CPU utilization is 50, the weighted CPU utilization is (20% * 60) + [(100 - 20)% * 50] or 52%.

  If you want to monitor the current CPU utilization only, set this parameter to 100.

  The `MonitorInterval` parameter determines how often the weighted CPU utilization is calculated (default is every 5 seconds).

- `MonitorInterval` (seconds)

  Default: *5*
  Range: 1 - 43200 (12 hours)

  How often, in seconds, threshold policies are monitored.

- `MonitorResponseTimeWeight` (%)

  Default: *50*
  Range: 1 - 100

  Percentage of the current response time used to calculate the weighted response time. The weighted response time is based on the current and previous (last measured) response time:

  weighted response = (`MonitorResponseTimeWeight`% * current response time) + [(100 - `MonitorResponseTimeWeight`)% * previous response time]

  For example, if you set the `MonitorResponseTimeWeight` to 20, the current response time is 3 seconds, and the previous response time is 2 seconds, the weighted response time is (20% * 3) + [(100 - 20)% * 2] or 2.2 seconds.

  If you want to monitor the current response time only, set this parameter to 100.

  The `MonitorInterval` parameter determines how often the weighted response time is calculated (default is every 5 seconds).

- `NumOfDeferredSessionsPerSecond`

  Default: *1*

  The number of deferred sessions that have the same deferral time allocated to them. The default minimum deferral time is 10 seconds and the default maximum deferral time is 120 seconds.

  For example, if you set this parameter to **1** and three requests arrive at the same time, the first deferred session is allocated a deferral time of 10 seconds, the second deferred session is allocated a deferral time of 11 seconds, and the third deferred session is allocated a deferral time of 12 seconds. If you set the parameter to **2** and three requests arrive at the same time, the first and second deferred sessions are allocated a deferral time of 10 seconds and the third deferred session is allocated a deferral time of 11 seconds.

  The default minimum and maximum deferral times (also called wait

times) can be configured for a template by selecting the Advanced tab and then Corrective Actions Policy Details.

SCA Parameters

The following parameters affect the SCA only. If you change any of these parameters, you must restart the SCA.

These parameters impact measurement log files. In UNIX, the measurement log files are located in the /var/opt/webqos/measures directory. In Windows, these files are located in the *<install-directory>*\measures directory.

- LogInterval (minutes)

  Default: *30*
  Range: >=1

  How often, in minutes, WebQoS logs measurements to a log file. The LogOn parameter must be enabled (set to 1) before WebQoS logs measurements to the log file.

- LogOn

  Default: *0* (disabled)
  Range: 0 (disabled) or 1 (enabled)

  Enable or disable measurement logging. To enable measurement logging, set this parameter to 1. To disable measurement logging, set this parameter to 0.

  WebQoS creates a new measurement log file daily for each WebQoS site. The measurement log file name is the site name you entered in the management console followed by the date. The log file contains performance measurements collected during the day for the site including:

  — Average response time for each request class

  — Number of new, redirected, deferred, and rejected sessions for each request class

- SaveFiles

  Default: *0* (disabled)
  Range: 0 (disabled) or 1 (enabled)

  Enable or disable saving measurement log files. If enabled, all measurement log files are saved. If disabled, measurement log files

more than two days old are deleted.

SCO Parameters

The following parameters affect the SCO only. If you change any of these parameters, you must restart the SCO.

- `HistoryEventDays` (days)

  Default: *7*
  Range: 1 - 365

  How many days of history events to save. At midnight, history events older than the specified age are deleted.

- `SCOLogExternMeas`

  Default: *0* (disabled)
  Range: 0 (disabled) or 1 (enabled)

  Enable or disable external measurement logging.

  The `SCOLogInterval` parameter determines how often the SCO writes external measurement data to a log file.

  In UNIX, the external measurement log files are located in the `/var/opt/webqos/stats` directory. In Windows, these files are located in the *<install-directory>*`\stats` directory.

- `SCOLogInterval` (minutes)

  Default: *5*
  Range: 1 - 1440 (24 hours)

  How often, in minutes, WebQoS logs aggregate service statistics to a service statistics log file. Service statistics logging must be enabled using the management console before aggregate service statistics are logged. See “Service Statistics” on page 50 for a list of the logged aggregate statistics.

  In UNIX, the service statistics log file is located in the `/var/opt/webqos/stats` directory. In Windows, these files are located in the *<install-directory>*`\stats` directory.

- `ServiceSLOInterval` (seconds)

  Default: *30*
  Range: 5 - 43200 (12 hours)

  How often, in seconds, the SCO monitors the service level SLOs.

Logging and Tracing Parameters (UNIX Only)

Refer to “Changing Logging and Tracing Levels” on page 116 for information on how to change these levels on your Windows server.

The following parameters affect the amount of information written to the logging and tracing files. Logging and tracing is dynamically updated for the web server (information written to `qoslog`, `qoslog.child`, `qostrace`, and `qostrace.child`). However, you must restart the SCA, SCI, and SCO for changes to affect logging and tracing for WebQoS components. You can control logging and tracing levels by component by restarting specific components. For example, if you change the `LOGLEVEL` from 3 to 4 but only restart the SCO, detailed logging only occurs for the SCO. If you do not restart the SCA and SCI, the logging level remains at level 3.

- `LOGLEVEL`

  Default: *3* (informative logging)
  Range: 0 - 4

  The amount of logging information that is written to the `qoslog`, `qoslog.child`, `scalog`, `scilog`, and `scolog` log files. Logging values are:

  0 - None
  1 - Error
  2 - Warning
  3 - Informative
  4 - Detailed

  The larger the value, the more logging information is generated.

  The log files are located in the `/var/opt/webqos/logs` directory.

- `TRACELEVEL`

  Default: *0* (no tracing)
  Range: 0 - 3

  The amount of tracing information that is written to the `qostrace`, `qostrace.child`, `scatrace`, `scitrace`, and `scotrace` files. Tracing values are:

  0 - None
  1 - Data flow
  2 - Detailed
  3 - Procedural

The larger the value, the more tracing information is logged.

The trace files are located in the `/var/opt/webqos/logs` directory.

URL Encoding Parameters

The following parameters allow you to support users who have disabled cookies on their browser. These parameters are useful only if you are using an application server such as the BroadVision One-To-One Server[1]. If you change any of these parameters, you must restart the SCA.

- `APP_SESS_ENABLE_URL_ENCODING`

  Default: *0* (disabled)
  Range: 0 (disabled), 1 (enabled), or 3 (dynamic enable)

  Enable or disable URL encoding. If disabled, cookies are used to pass state information and users who have disabled cookies are not supported. If enabled, URLs are encoded to pass state information, users who have disabled cookies are supported, but the aggregate session count SLO is not useful. If set to 3, the application server you are using allows dynamic switching between cookies and URL encoding.

- `APP_SESSION_ID_STR`

  No default.

  The application server session identifier. How the application server uniquely identifies its sessions. For example, BroadVision uses the identifier `BV_SessionID`, so you would enter **`APP_SESSIONS_ID_STR BV_SessionID`** in the `qos.conf` file.

  You must enable `APP_SESS_ENABLE_URL_ENCODING` to use this parameter.

1. The BroadVision One-To-One Server is a product of BroadVision, Inc.

Restarting Individual WebQoS Components

In UNIX, if you have configured parameters in the `qos.conf` file after you have run the WebQoS `setup` script, you must restart (by stopping and starting) specific components (the SCA, SCI, and/or SCO).

In Windows, if you have configured parameters in the `qos.conf` file, you must restart specific components (the SCA, SCI, and/or SCO).

Restarting WebQoS Components (UNIX)

To restart the SCA, type the following:

```
/opt/webqos/bin/sca_stop
/opt/webqos/bin/sca_start
```

To restart the SCI, type the following:

```
/opt/webqos/bin/sci_stop
/opt/webqos/bin/sci_start
```

To restart the SCO, type the following:

```
/opt/webqos/bin/sco_stop
/opt/webqos/bin/sco_start
```

Restarting WebQoS Components (Windows)

On a Windows NT system, to restart a WebQoS component, go to the Control Panel and select **`Services`**. Highlight the component you want to restart (WebQoS SCA, WebQoS SCI, or WebQoS SCO), click on the Stop button and then the Start button.

On a Windows 2000 system, to restart a WebQoS component, click on **`Start -> Programs -> Administrative Tools -> Services`**. Double-Click on the component you want to restart (WebQoS SCA, WebQoS SCI, or WebQoS SCO).

WebQoS setup Script (UNIX)

This utility enables and disables the WebQoS daemons and “qosifies” your web server. You run it after your web server is installed and after WebQoS is installed on systems running any WebQoS component or if

you add a new instance of a web server to your “qosified” web server.

Run the following script on any UNIX server you are using as an application tier or a web server tier for WebQoS:

```
/opt/webqos/bin/setup
```

Single Server Setup

In a single server setup, you are using only one server in your WebQoS domain.

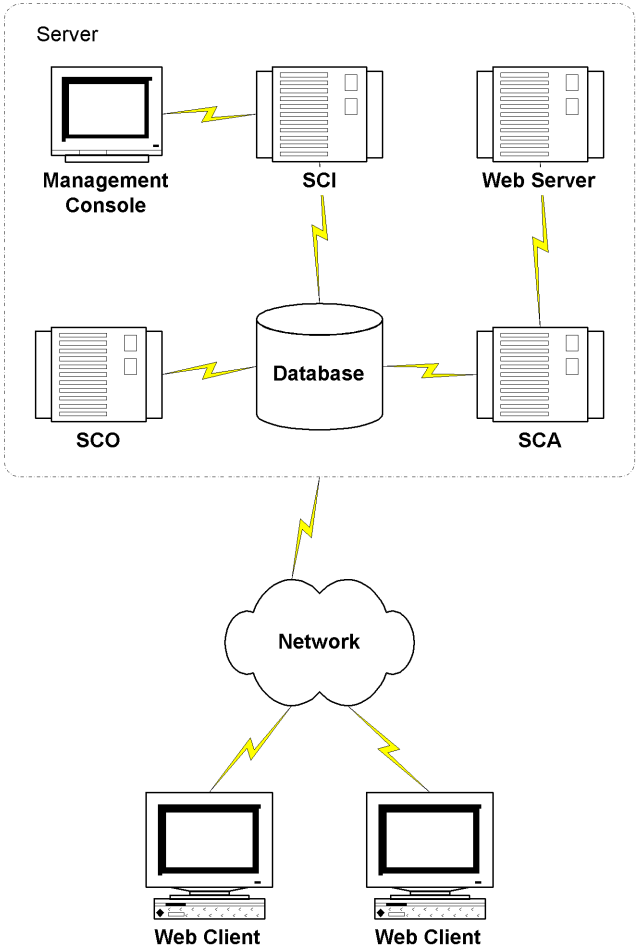

When you run the `setup` script, provide the following information:

- `Continue?`

  Enter **y** or hit **Enter** to continue. If you do not want to run WebQoS or read the software license terms, enter **n** to exit the script now.

- `Do you accept the HP Software License Terms?`

  If you accept the software license terms and want to continue running the `setup` script, enter **y**. Otherwise, enter **n** to exit the script.

- `Would you like to create the WebQoS database on this system?`

  You must create a database in your WebQoS domain in order to configure and manage your WebQoS-enabled web servers.

  Enter **y** or hit **Enter** to create the WebQoS database.

- `Enter a database password:`
  `Re-enter the password:`

  Enter a password for the WebQoS database. This password is not recoverable. If you forget this password, you must recreate the database.

- `Enter a password:`
  `Re-enter the password:`

  Enter a password for the WebQoS management console. There is one password for all management consoles in a WebQoS domain. The management console is used to configure policies and rules.

- `Would you like to enable the SCO?`

  You must enable one SCO in your WebQoS domain in order to monitor service level SLOs.

  Enter **y** or hit **Enter** to enable the SCO.

- `Would you like to enable the SCI?`

  You must enable one SCI in your WebQoS domain in order to run the management console and the API.

  Enter **y** or hit **Enter** to enable the SCI.

- `Root pathname of Netscape server.`

  This is the path to the web server you want to manage on this system. The setup script installs WebQoS for all the web servers operating

under the directory you specify.

After the script finishes, start your web server(s).

Multiple Server Setup

In a multiple server setup, you are using more than one server in your WebQoS domain.

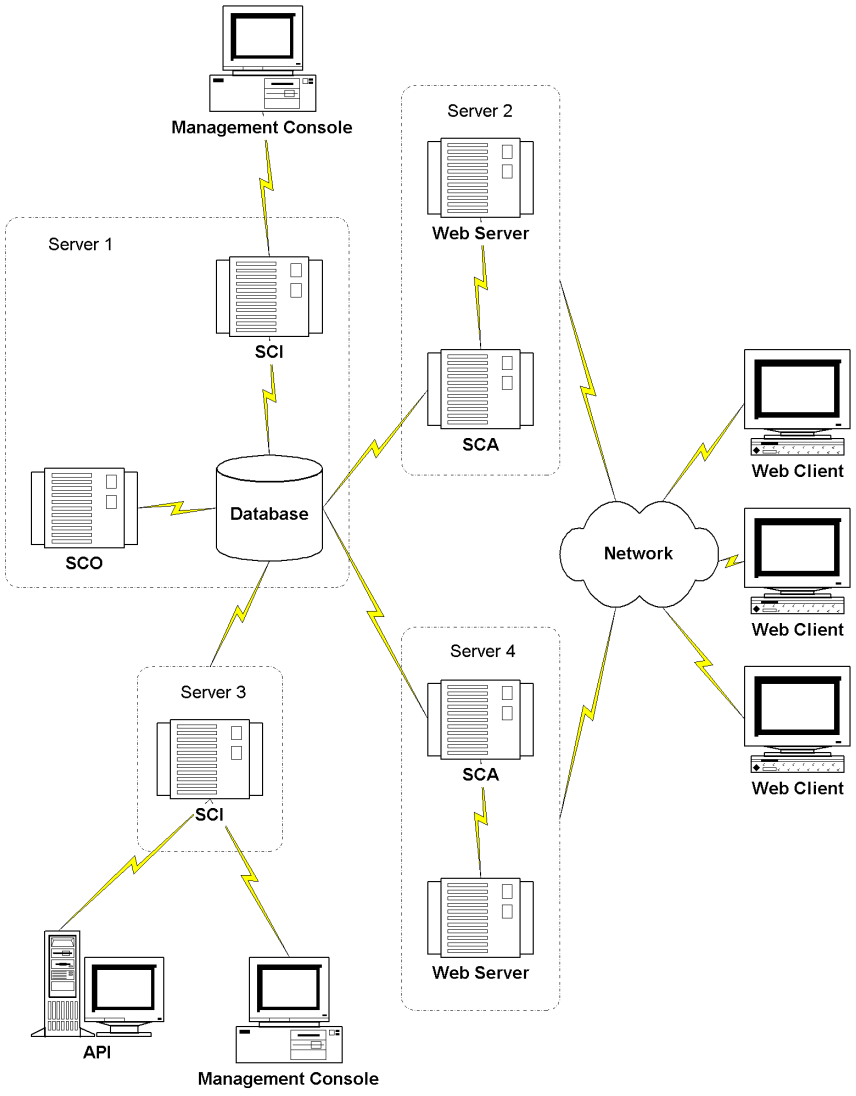

Server 1 Setup

On server 1, you are creating the database and enabling the SCO and SCI. When you run the `setup` script on server 1, you must provide the following information:

- `Continue?`

  Enter **y** or hit **Enter** to continue. If you do not want to run WebQoS or read the software license terms, enter **n** to exit the script now.

- `Do you accept the HP Software License Terms?`

  If you accept the software license terms and want to continue running the `setup` script, enter **y**. Otherwise, enter **n** to exit the script.

- `Would you like to create the WebQoS database on this system?`

  You must create one database in your WebQoS domain in order to configure and manage your WebQoS-enabled web servers. The database is created on this server.

  Enter **y** or hit **Enter** to create the WebQoS database.

- `Enter a database password:`
  `Re-enter the password:`

  Enter a password for the WebQoS database. This password is not recoverable. If you forget this password, you must recreate the database.

- `Enter a password:`
  `Re-enter the password:`

  Enter a password for the WebQoS management console. There is one password for all management consoles in a WebQoS domain. The management console is used to configure policies and rules.

- `Would you like to enable the SCO?`

  You must enable one SCO in your WebQoS domain in order to monitor service level SLOs. The SCO is running on this server.

  Enter **y** or hit **Enter** to enable the SCO.

- `Would you like to enable the SCI?`

  You must enable one SCI in your WebQoS domain in order to run the management console and the API. You are running the SCI on this

server.

Enter **y** or hit **Enter** to enable the SCI.

- `Root pathname of Netscape server.`

  Enter **NONE**. You do not want to configure any web servers.

Servers 2 and 4 Setup

On servers 2 and 4, you are running the web servers. You must identify server 1 as the system on which the database is running. When you run the `setup` script on servers 2 and 4, you must provide the following information:

- `Continue?`

  Enter **y** or hit **Enter** to continue. If you do not want to run WebQoS or read the software license terms, enter **n** to exit the script now.

- `Do you accept the HP Software License Terms?`

  If you accept the software license terms and want to continue running the `setup` script, enter **y**. Otherwise, enter **n** to exit the script.

- `Would you like to create the WebQoS database on this system?`

  Enter **n**. The database has been created on server 1.

- `Please configure the name of the system where the database runs.`

  Enter the name of server 1.

- `Enter the database password:`

  Enter the password for the WebQoS database on server 1. If you have forgotten the password, you must recreate the database.

- `Would you like to enable the SCO?`

  Enter **n**. You are not running the SCO on this server.

- `Would you like to enable the SCI?`

  Enter **n**. You are not running the SCI on this server.

- `Root pathname of Netscape server.`

  Enter the path to the web server you are managing on this system. The setup script installs WebQoS for all the web servers operating

under the directory you specify.

After the script finishes, start your web server(s).

Server 3 Setup

On server3, you are running the SCI. You must identify server 1 as the system on which the database is running. When you run the `setup` script on server 3, you must provide the following information:

- `Continue?`

  Enter **y** or hit **Enter** to continue. If you do not want to run WebQoS or read the software license terms, enter **n** to exit the script now.

- `Do you accept the HP Software License Terms?`

  If you accept the software license terms and want to continue running the `setup` script, enter **y**. Otherwise, enter **n** to exit the script.

- `Would you like to create the WebQoS database on this system?`

  Enter **n**. The database has been created on server 1.

- `Please configure the name of the system where the database runs.`

  Enter the name of server 1.

- `Enter the database password:`

  Enter the password for the WebQoS database on server 1. If you have forgotten the password, you must recreate the database.

- `Would you like to enable the SCO?`

  Enter **n**. You are not running the SCO on this server.

- `Would you like to enable the SCI?`

  Enter **y**. You are running the SCI on this server.

- `Root pathname of Netscape server.`

  Enter **NONE**. You do not want to configure any web servers.

Note that you were not asked to enter a management console password. The password was configured when you ran the `setup` script on server 1 and stored in the database.

Removing WebQoS

If you want to "unqosify" (not use WebQoS with) your web server, run the following command.

```
/opt/webqos/install/setup -r
```

WebQoS configuration data is preserved in case you want to run WebQoS again later.

When you run `setup -r`, you must provide the following information:

- `Enter the root pathname of Netscape server.`

  This is the path to the web server you want to unqosify. The `setup -r` script unqosifies all the web servers operating under the directory you specify.

After the script finishes, start your web server(s).

3 The Management Console

This chapter describes how to log in to the management console and its basic functionality.

Management Console Login

To log in to the WebQoS management console, do the following:

1. You should have already installed the management console software. Refer to the HP WebQoS release note for information on installing this software.
2. On your Solaris system, type **`/opt/webqos/bin/console_start`**.

   On a PC, double-click on the WebQoS icon on the Windows desktop. The WebQoS login dialog box appear (Figure 3-1, “WebQoS Login Dialog Box”).

Figure 3-1 WebQoS Login Dialog Box

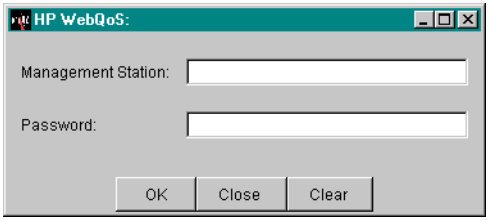

3. Enter the system name where the SCI is installed in the Management Station field.
4. Enter the management console password. This is the management console password you configured when you ran the `setup` script (UNIX) or the installation file (Windows). There is one password for all management consoles in a WebQoS domain.

   You can change the management console password. See “Password” on page 44 for more information.
5. Choose `[OK]`.

Console Main Screen

When you successfully log in, the management console's main screen appears. Your systems and sites are automatically discovered.

The following figure and table describe the management console's main screen. For more information, refer to the online help.

Figure 3-2 The Management Console Main Screen

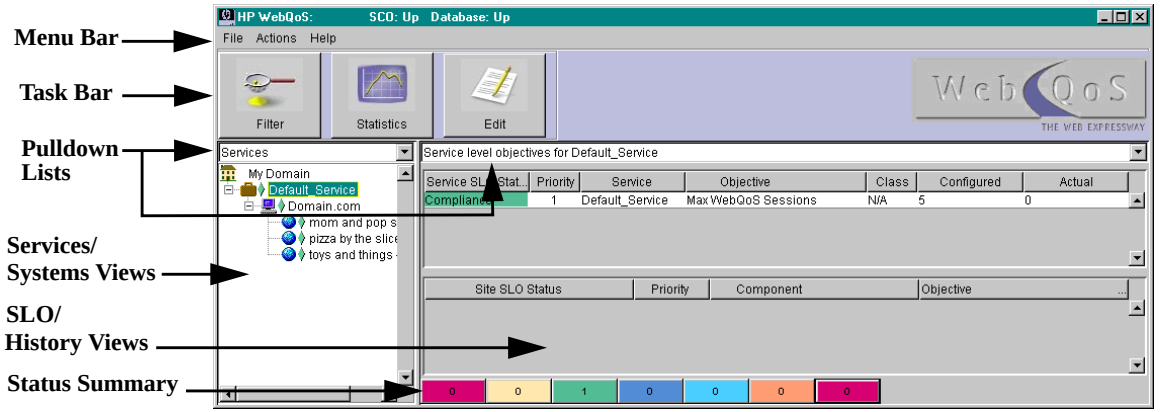

Window Area	Description
Menu Bar	The `File` menu allows you to change the administrator's password, select a management station to which to connect, or exit the management console. The `Actions` menu allows you to add, move, edit, or delete services, systems, and/or sites. You can also filter SLOs and the History log, set the default template, and view statistics. The `Help` menu provides overview online help and product information. It also provides a link to HP's support web site.
Task Bar	Each button provides quick access to an action in the `Actions` menu.

Window Area	Description
Pulldown Lists	The Services/Systems pulldown list allows you to select either the services view or the systems view. The Service Level Objective/History log pulldown list allows you to select either the SLO view or the History log view.
Services/Systems Views	These views display a tree containing your configured services, systems, and sites (your domain topology). The Services view displays your domain by services. The systems and sites that provide each service are displayed below the service. The Systems view displays your domain by systems. The services and sites available on each system are displayed below the system. For an explanation of the symbols see “The Services/Systems View” on page 43. When editing or deleting a service or site, you must first select it in this view area.
SLO /History Views	Displays the SLOs or history log for the service, system, or site that you selected in the Services/Systems view area. For more information about these views, refer to Chapter 7, “Logs and Filters.”.
Status Summary	The buttons at the bottom report the number of SLOs being violated, at risk, in compliance, and inactive. For an explanation of the color codes, see “The Service Level Objectives View” on page 96.

Auto Discovery

WebQoS uses a database to store information about your domain. WebQoS uses this information to automatically display the system(s) and site(s). The systems displayed are those that are using the same database to store configuration information. The site(s) displayed are those that are configured for each system.

For more information about adding and removing systems, refer to “Systems” on page 82.

For more information about sites, refer to “Sites” on page 80.

The Services/Systems View

The services/systems view shows the services, systems and sites configured in WebQoS. A WebQoS domain is a set of sites residing on one or more systems. Sites can also be grouped logically in services so that you can easily manage all the sites in that service.

By selecting the appropriate service, system, or site, you can determine which services, systems, and sites are meeting their SLO and threshold objectives.

You can navigate in this view by using the following navigation symbols:

- House - displays everything in your domain.
- Magnifying glass - click on the magnifying glass symbol to expand or reduce your view.
- Briefcase - displays a service.
- Monitor - displays a system name.
- Globe - displays a site name.
- Diamond - color coded to let you see which services, systems and sites are meeting their objectives. For a description of the color codes, see the section, “The Service Level Objectives View” on page 96.

In the services view, all configured services are displayed at the first level. Expanding the service displays the systems configured for that service. Expanding the system displays all the sites configured on the system under that service.

In the systems view, the systems are displayed at the first level. Expanding the system displays the services configured for that system.

Expanding the service displays all the sites configured for that service on that system.

Console Administration

Basic tasks that you can complete from the management console's main screen are changing the management console's password and managing another domain.

Password

To change the password that is used by the management console, do the following:

1. Select `File: Change Password`.
2. In the Change Password dialog box, type in the old password. You need to enter the new password twice.
3. Click on **OK**.

If you forget the management console's password, you can reset it. In UNIX, on any system running a WebQoS component, type:

```
/opt/webqos/install/setup -p
```

In Windows, go to the *<install-directory>* and type the following:

```
jre1.3\bin\java.exe -jar InstallPwd.jar <db_server_name> <db_password> <new_console_password>
```

Domain Management

To change the WebQoS domain that you are managing by connecting to another database, do the following:

1. Choose `File: Connect Management Station`.
2. In the WebQoS dialog box, type in the name of the system containing the database to which you want to connect and that WebQoS domain's management console password.
3. Click on **OK**.

4 Templates

This chapter describes the WebQoS templates.

Templates are used to configure your SLOs, corrective actions, thresholds, and/or request classification rules. If you have multiple

services or sites that use the same policies and rules, you only have to configure the policies and rules once in a template. You can then apply the template to each service or site.

Preconfigured and Default Templates

WebQoS comes with preconfigured service and site templates. These templates are also set as the default templates.

Preconfigured Templates

WebQoS comes with the following two preconfigured templates:

- Default_Service_Template
- Default_Site_Template

The templates do not have any policies or rules configured.

When WebQoS is started for the first time, any sites that are automatically discovered are assigned the default site template. Any systems and sites that are automatically discovered are assigned to a default service named `Default_Service`. `Default_Service` uses the default service template.

Because the preconfigured templates do not have any policies or rules configured, your sites will function as if WebQoS is not configured on your server.

Default Templates

The default template is used when you delete a template that is still assigned to a service or site. If you delete a template that is configured to a service, that service is reconfigured with the default service template.

Likewise, if you delete a template that is configured to a site, that site is reconfigured with the default site template.

You can set the default templates by selecting `Actions: Set Default Template`.

Service Template

Description

This screen displays the name of the service template.

Figure 4-1 Service Template Description

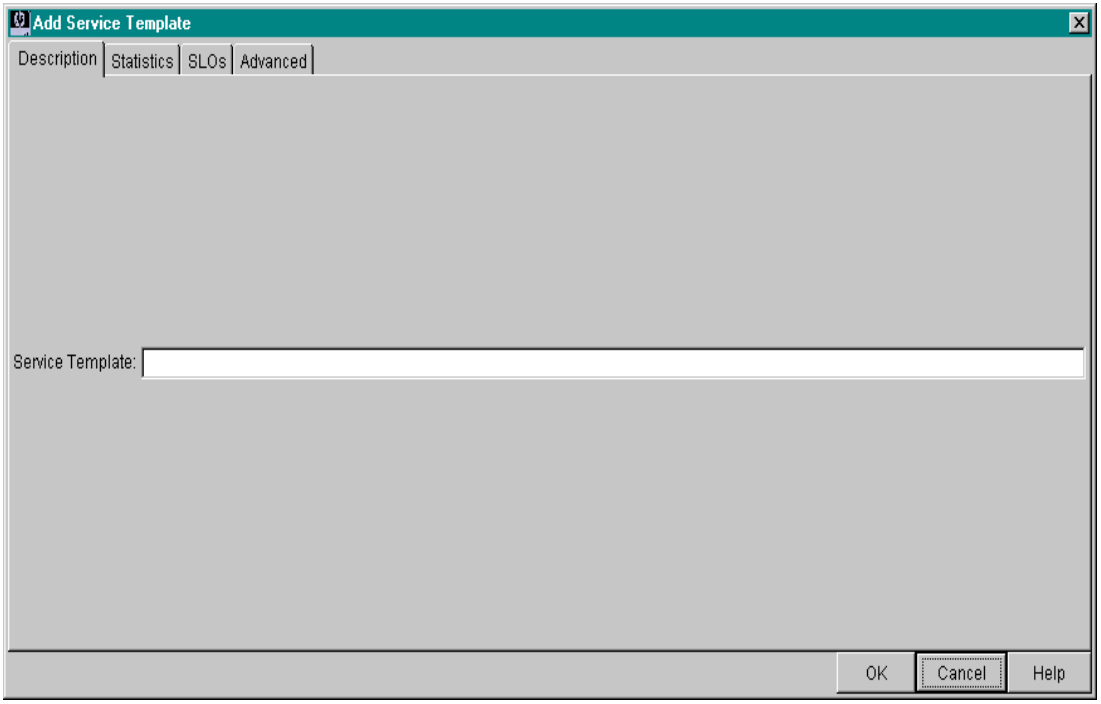

Enter a service template name. A service template name can be anything that is meaningful to you as an administrator of WebQoS. For example, a service template name could be “Big Company Template.” This template can be assigned to the “Big Company” service that might contain sites such as: `www.BigCompanyStore.com`, `www.BigCompanySupport.com` and `www.BigCompanyHR.com`. Spaces are allowed in the service template name.

If you wish to log service statistics logging, configure SLOs, or configure corrective actions, go to the next sections. Otherwise click `[OK]`.

Statistics

This screen displays whether or not service statistics logging is enabled for the service template.

Figure 4-2 Service Template Statistics

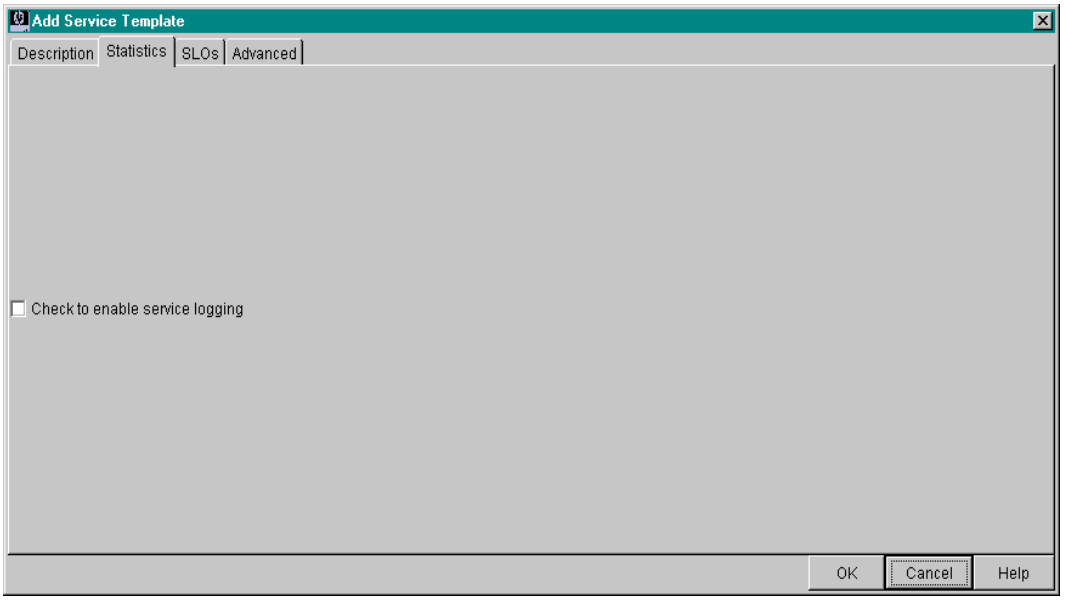

By default, service statistics logging is disabled.

If you enable service statistics logging, WebQoS logs service statistics to a log file on the management server (the server on which the SCO is installed). In UNIX, the log files are located in the `/var/opt/webqos/stats` directory. In Windows, the log files are located in the *<install-directory>*`\stats` directory.

You can specify how often the statistics are generated by editing the `SCOLogInterval` parameter in the `qos.conf` file. The default value is 5 minutes. For more information on the `SCOLogInterval` parameter, see “qos.conf” on page 23.

Service Statistics

If you enable service statistics logging, WebQoS logs service statistics to a log file named *<service_name>*`.stat` where *service_name* is the name of the service to which this template is assigned.

The log file is created daily until you disable service statistics logging. At midnight, the file is moved to <*service_name*>.stat.<*previous_date*>.

The following aggregate service statistics are logged:

Timestamp
Sessions
Admissions
Deferrals
Redirections
Rejections
Compliances (Number of sites in compliance)
N/A (This field is not used)
Violation (Number of sites in violation)
Other (Number of sites in states other than compliance or violation)
Response High
Response Medium
Response Low
Admissions High
Admissions Medium
Admissions Low
Deferrals High
Deferrals Medium
Deferrals Low
Redirections High
Redirections Medium
Redirections Low
Rejections High
Rejections Medium
Rejections Low
Deferred Outstanding High
Deferred Outstanding Medium
Deferred Outstanding Low
Deferral Time Average High
Deferral Time Average Medium
Deferral Time Average Low
Deferral Time Current High
Deferral Time Current Medium
Deferral Time Current Low

URL Group Statistics

If you enable service statistics logging and the SLO Ensure less than <NUMBER> millisecond response time for <CLASS> requests to <URL GROUP> is configured, URL group statistics are logged to the URL_<*service_name*>_<*URL GROUP*>_<*CLASS*>_<*NUMBER*>.stat file

where *service_name* is the name of the service to which this template is assigned, and *URL GROUP*, *CLASS*, and *NUMBER* correspond to the parameters defined in the SLO.

The following aggregate URL group statistics are logged:

Timestamp
URL Group Response Time

SLOs

This screen displays the SLOs configured for the service template.

Figure 4-3 Service Template SLOs

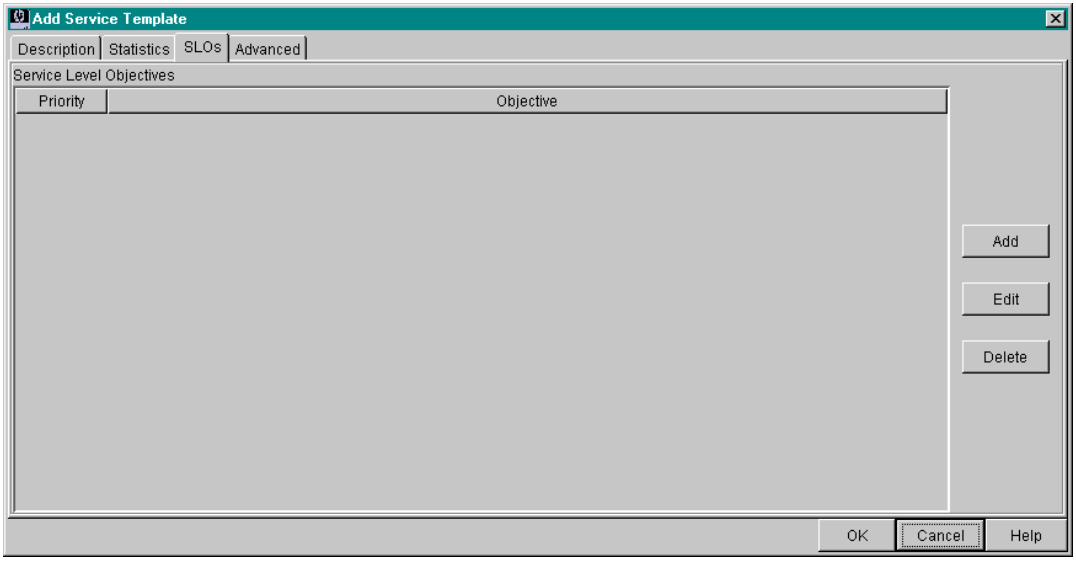

An SLO specifies the level of service that you expect to provide. They can be configured to set a minimum average response time, the maximum concurrent sessions supported, and to meet an externally specified measurement.

WebQoS monitors the SLOs to make sure they are met. If they are not met, WebQoS will execute any configured corrective actions.

If you have configured an SLO, the results can be viewed from the management console's main screen. For more information on statistics, see “Service Statistics” on page 86.

To configure an SLO, click [Add] or highlight an existing SLO and click [Edit].

Add/Edit SLOs

This screen displays the SLOs you can configure and the corrective actions configured to the SLOs for the service template.

Figure 4-4 Add/Edit Service Level Objectives

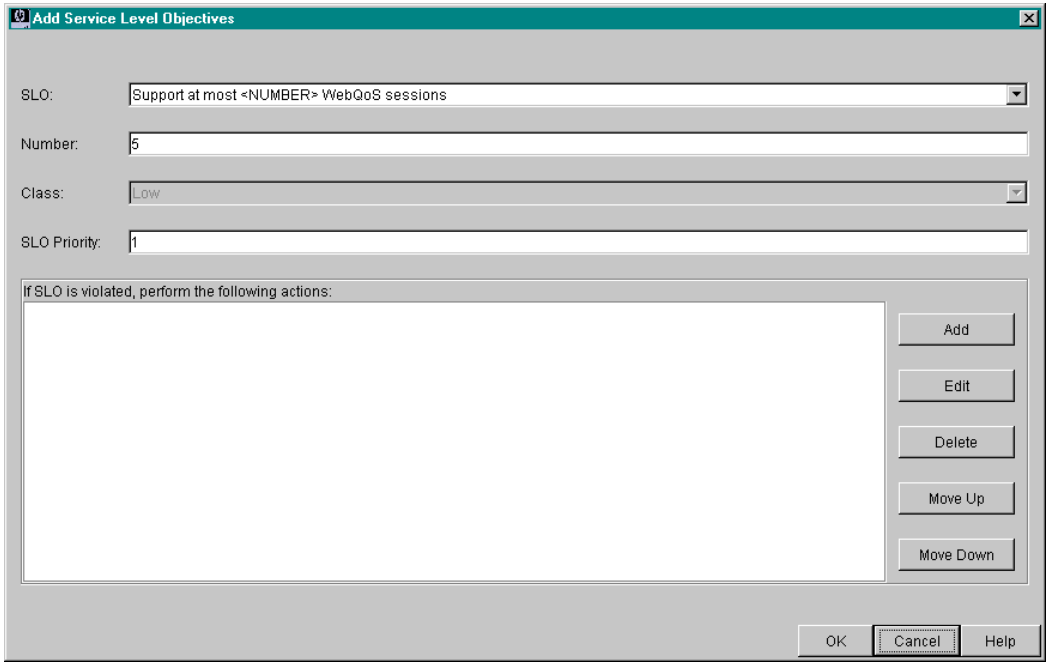

You can configure the following SLOs:

- Support at most <NUMBER> WebQoS sessions
- Ensure less than <NUMBER> millisecond response time for <CLASS> requests to <URL GROUP>
- Ensure <Measurement> is <CONDITION> <NUMBER>

If you modify an SLO, all of its events in the history log are deleted.

For more information about these SLOs, refer to the online help and “Service Level Objectives for the Service” on page 123.

If you wish to add corrective actions for when the SLO is violated, go to

the next section. Otherwise click [OK].

Add/Edit Corrective Actions This screen displays the corrective actions you can configure for the service template.

Figure 4-5 Add/Edit Corrective Actions

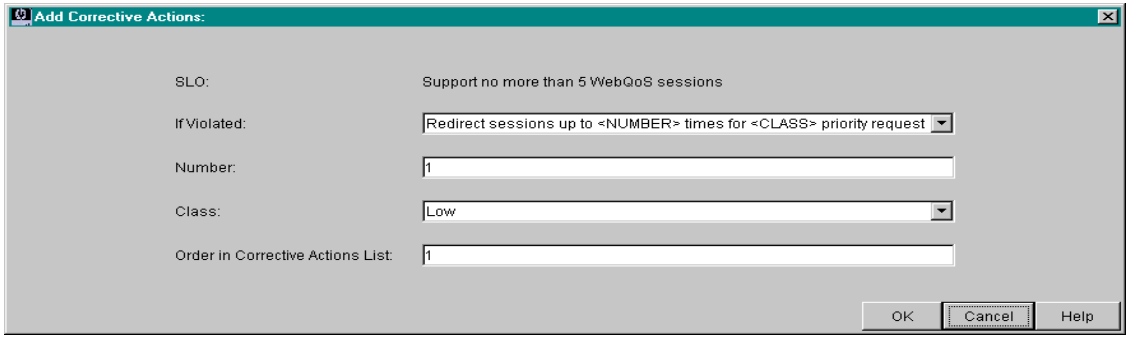

Corrective actions are those actions taken to bring an SLO into compliance. Corrective actions may involve making trade-offs between requests belonging to different request classifications, or making trade-offs between requests going to different sites or services running on the same system. You can add, edit or delete corrective actions by selecting the appropriate button.

You can configure the following corrective actions:

- Redirect sessions up to <NUMBER> times for <CLASS> priority requests
- Reject <CLASS> priority requests
- Defer <CLASS> priority requests

For more information about these corrective actions, refer to the online help and “Corrective Actions for SLO and Threshold Policy Violations” on page 124.

If more than one action is configured for one class, the action with the highest priority occurs first. Priority is based on the order assigned to the action, with order 1 having the highest priority.

For example, a set of corrective actions can be:

Order 1: Redirect low class sessions 3 times

Order 2: Reject low class sessions

A low class session is redirected up to three times. If a session has been redirected three times, it is rejected.

It is recommended that the reject action be configured as the last action (lowest priority) for the class.

You can change the order of the corrective actions by moving their order using [Move Up] and [Move Down].

Advanced

From this screen, you can configure how a session is deferred, redirected, or rejected. You must have configured a corrective action to a SLO for these policies to take effect for the service template.

For information on customizing a defer and reject web page, refer to “Customizing the Defer and Reject Web Pages” on page 125.

Figure 4-6 Service Template Advanced

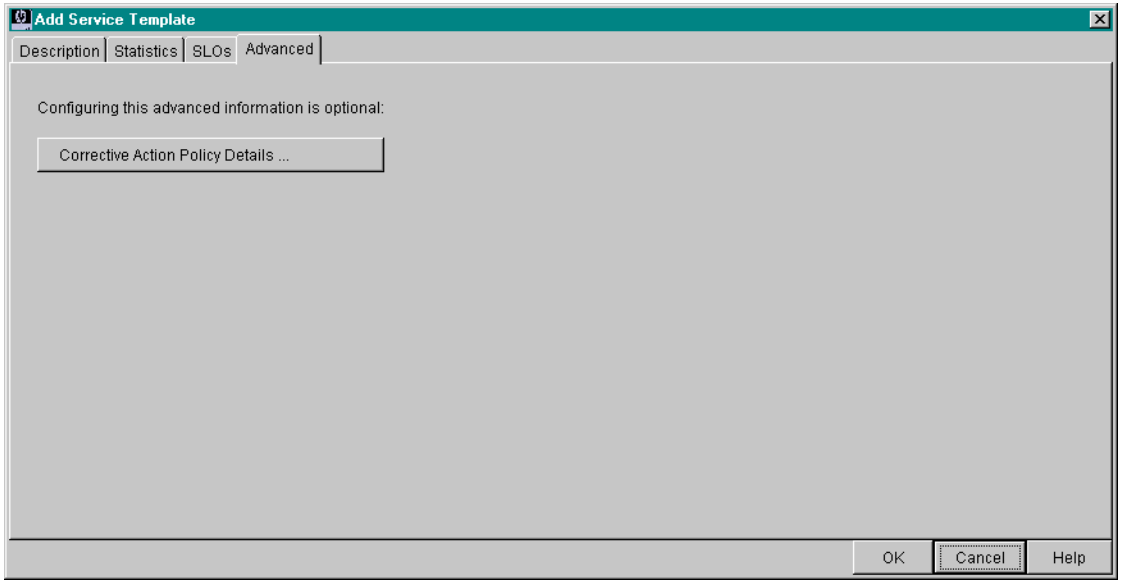

Defer

From this screen you can configure how a defer corrective action is executed for the service template.

Figure 4-7 Defer Corrective Action

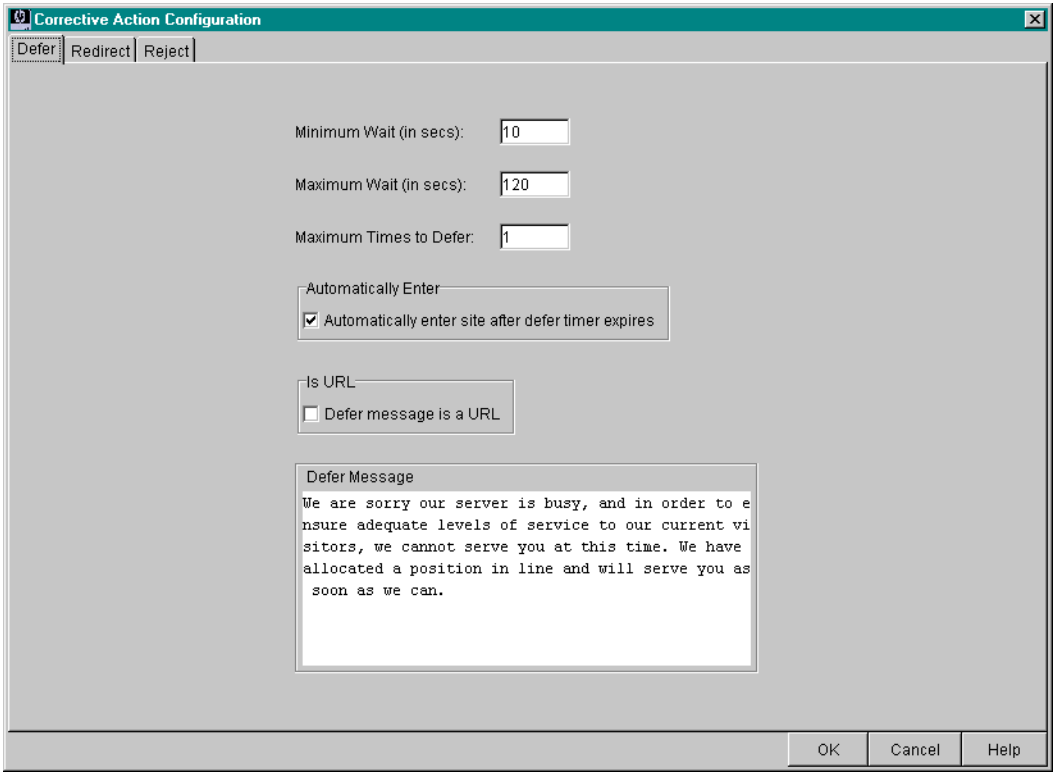

You can configure the minimum and maximum wait times of a deferred session, the maximum number of times a session is deferred, automatic entry after a defer, and the type of defer message. Refer to the online help for more information.

Redirect

From this screen you can configure how a redirect corrective action is executed for the service template.

Figure 4-8 Redirect Corrective Action

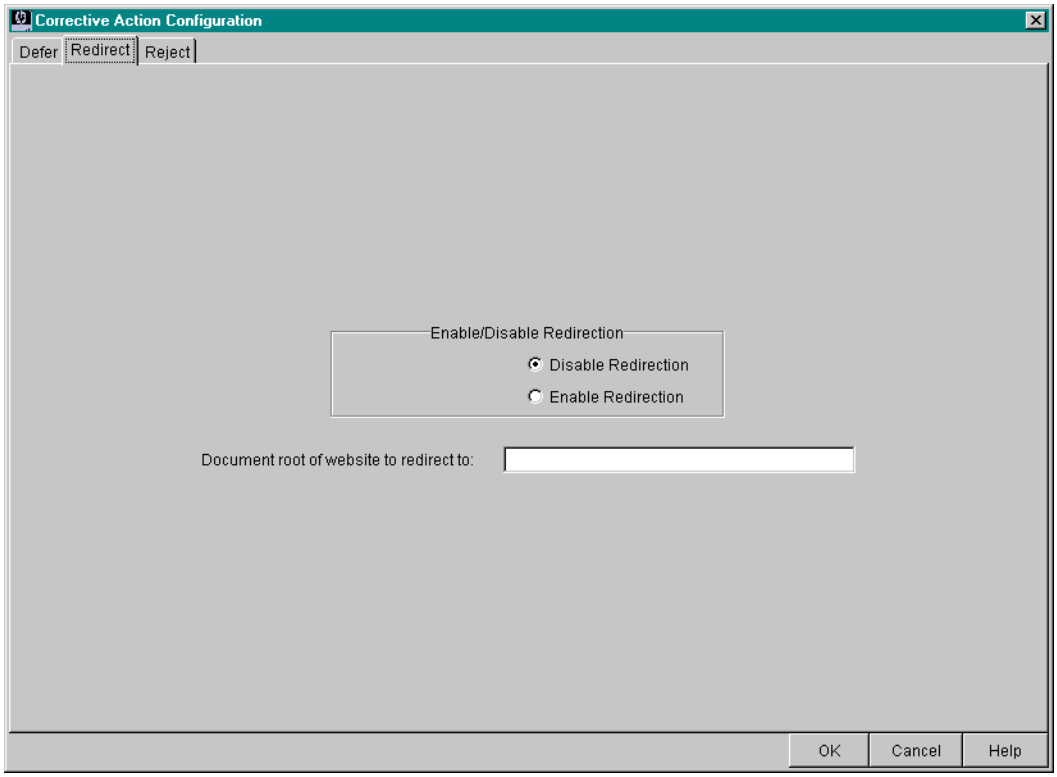

You can enable or disable redirection and configure the URL of a web site to which to redirect. Refer to the online help for more information.

If you have configured a redirect corrective action but disable redirection in this screen, all redirect corrective actions are ignored.

Reject

From this screen you can configure how a reject corrective action is executed for the service template.

Figure 4-9 Reject Corrective Action

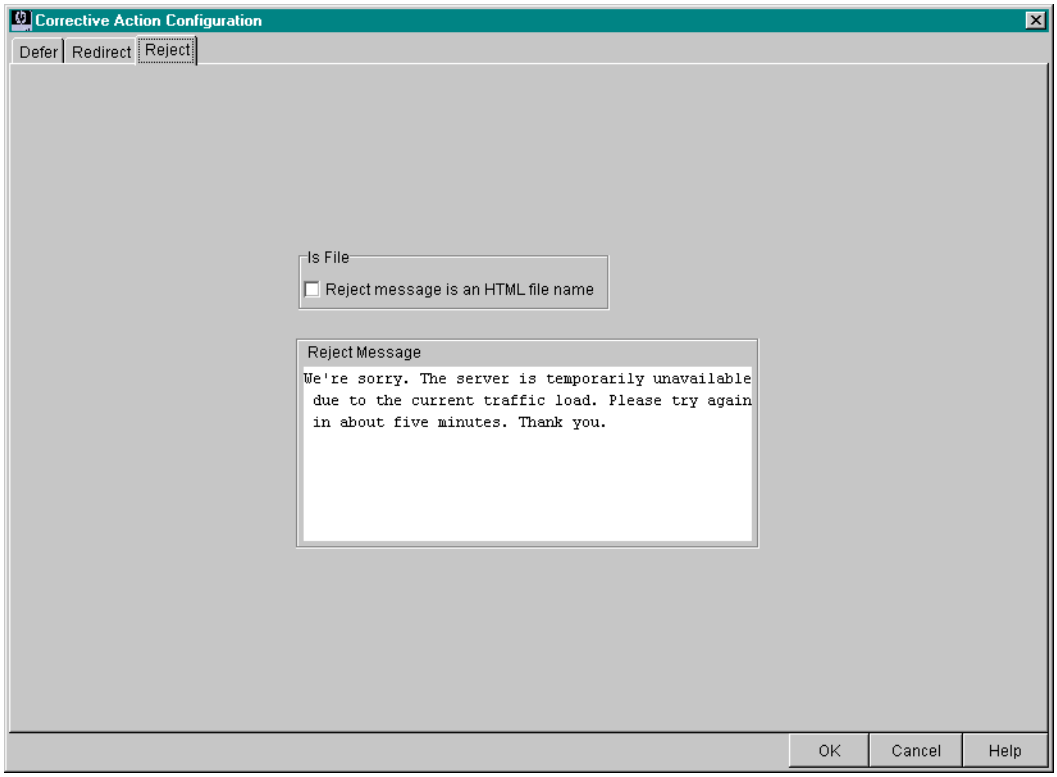

You can configure the message sent to a rejected session. Refer to the online help for more information.

Site Template

Description

This screen displays the name of the site template and the service to which this site template is associated.

Figure 4-10 Site Template Description

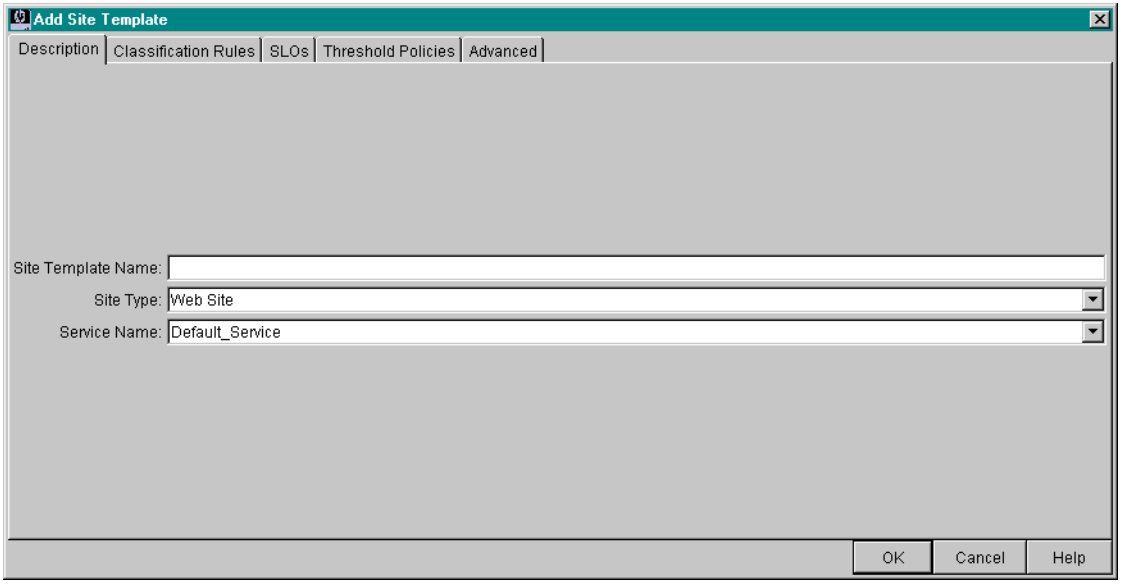

Enter a site template name. A site template name can be anything that is meaningful to you as an administrator of WebQoS. Select a service name from the pulldown box. The service names listed are those services that have already been configured.

Classification Rules

This screen displays the request classification rules configured for the site template.

Figure 4-11 Site Template Classification Rules

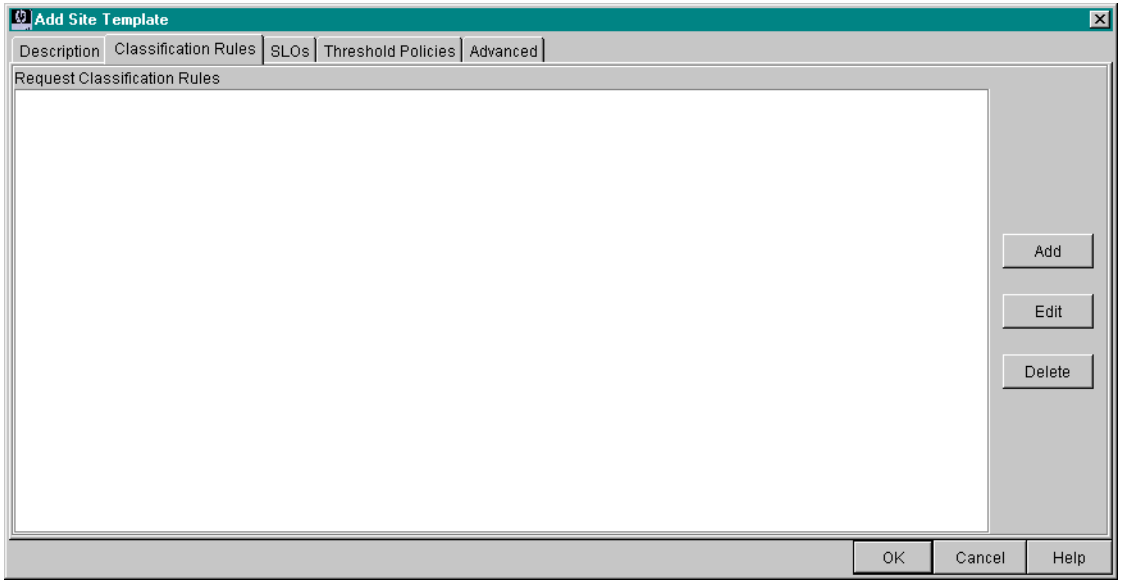

Request classification rules prioritize requests to the site. Requests can be differentiated by IP address, port number, URL document path, and/or virtual server name and assigned a class. These rules, when combined with the response time SLO, can determine how quickly a request is processed.

Adding request classification rules is optional. In Windows and Solaris, if you do not configure any request classification rules, all SSL *and* non-SSL requests are assigned the default (lowest) level of service.

Add/Edit Classification Rules

This screen displays the request classification rules you can configure for the site template.

Figure 4-12 Add/Edit Request Classifications Rules

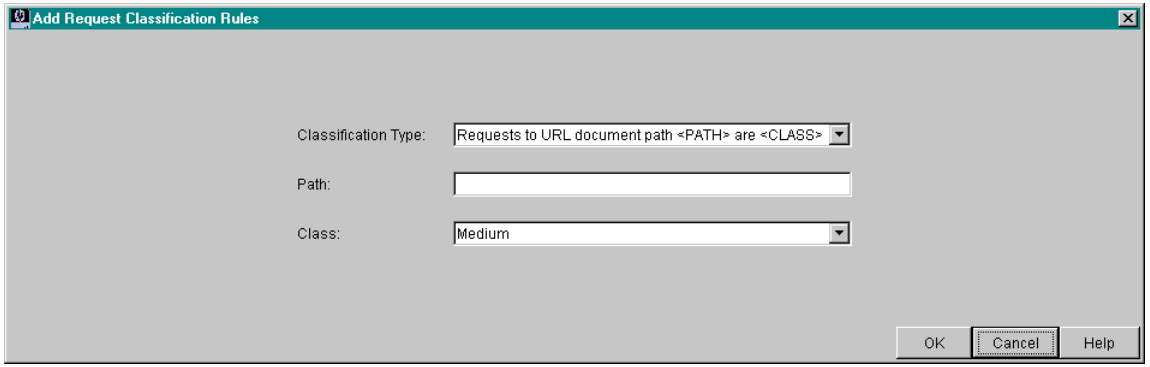

You can configure the following request classification rules:

- Requests to IP address <ADDRESS> are <CLASS>
- Requests from IP address <ADDRESS> are <CLASS>
- Requests to port number <NUMBER> are <CLASS>
- Requests to URL document path <PATH> are <CLASS>
- Requests to virtual server <NAME> are <CLASS>

For more information about these request classification rules, refer to the online help and “Request Classification Rules” on page 120.

SLOs

This screen displays the SLOs configured for the site template.

Figure 4-13 Site Template SLOs

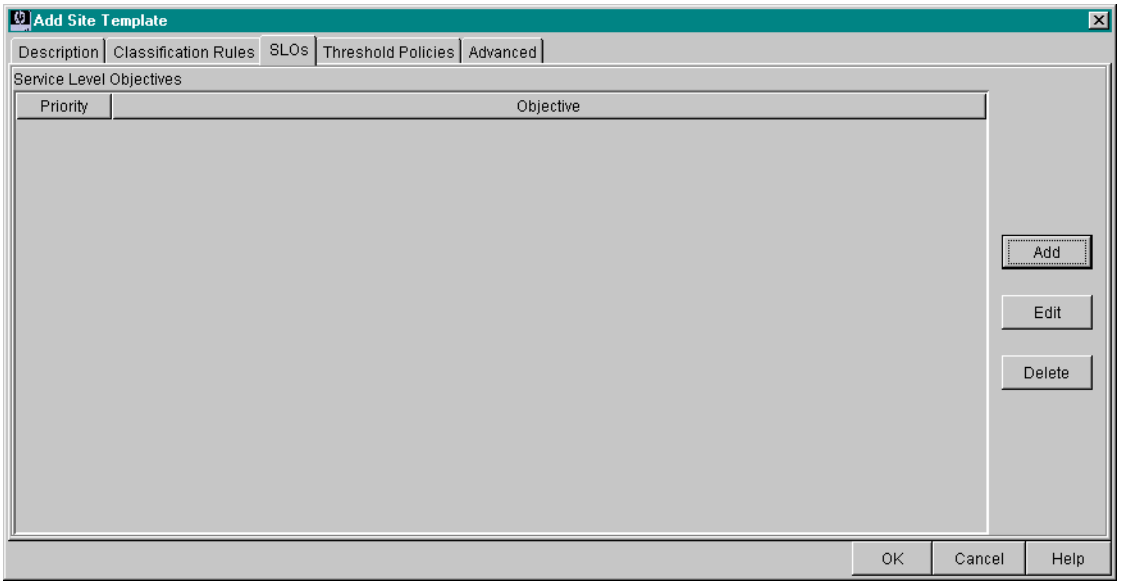

An SLO specifies the level of service that you expect to provide. They can be configured to set a minimum average response time and/or the minimum and maximum concurrent sessions supported.

WebQoS monitors the SLOs to make sure they are met. If they are not met, WebQoS will execute any configured corrective actions.

If you have configured an SLO, the results can be viewed from the management console's main screen. For more information on statistics, see "Site Statistics" on page 90.

Add/Edit SLOs

This screen displays the SLOs you can configure and the corrective actions configured to the SLOs for the site template.

Figure 4-14 Add/Edit SLOs

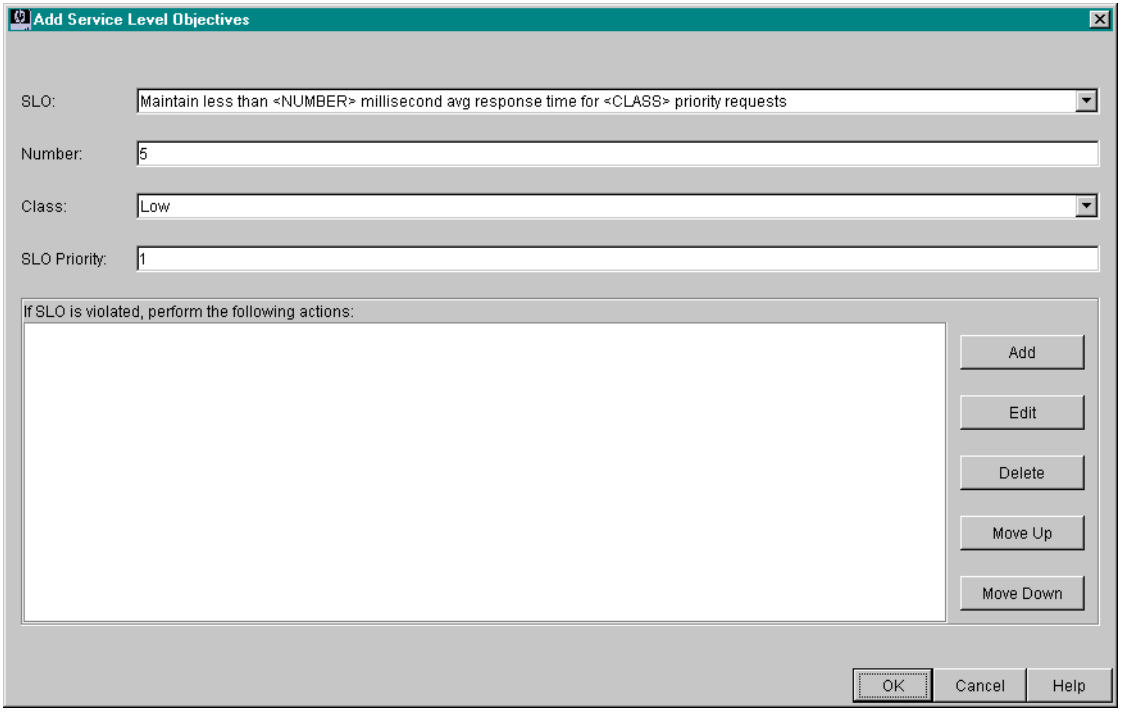

You can configure the following SLOs:

- Support at most <NUMBER> WebQoS sessions
- Maintain less than <NUMBER> millisecond avg response time for <CLASS> priority requests
- Support at least <NUMBER> WebQoS sessions

If you modify an SLO, all of its events in the history log are deleted.

For more information about these SLOs, refer to the online help and “Service Level Objectives for Site” on page 122.

If you wish to add corrective actions for when the SLO is violated, go to the next section. Otherwise click [OK].

Add/Edit Corrective Actions This screen displays the corrective actions you can configure for the site template.

Figure 4-15 Add/Edit Corrective Actions

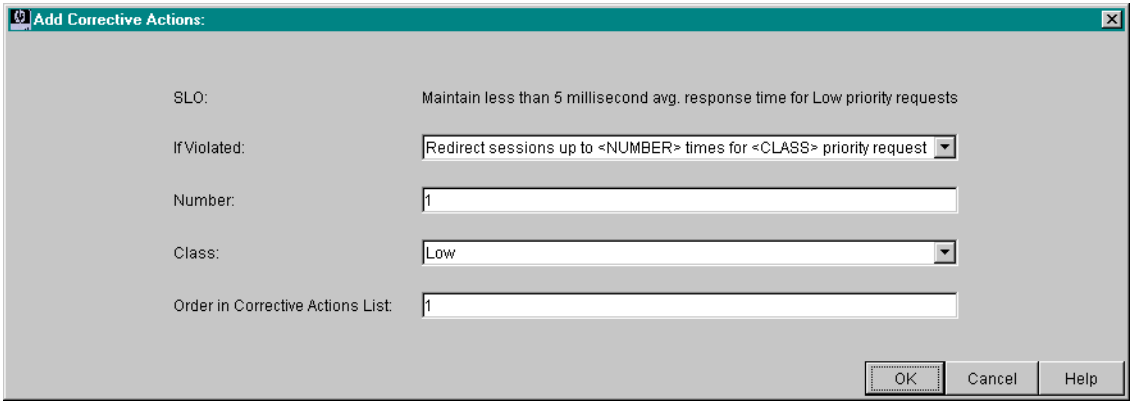

Corrective actions are those actions taken to bring an SLO into compliance. Corrective actions may involve making trade-offs between requests belonging to different request classifications, or making trade-offs between requests going to different sites or services running on the same system. You can add, edit or delete corrective actions by selecting the appropriate button.

You can configure the following corrective actions:

- Redirect sessions up to <NUMBER> times for <CLASS> priority requests
- Reject <CLASS> priority requests
- Defer <CLASS> priority requests

For more information about these corrective actions, refer to the online help and “Corrective Actions for SLO and Threshold Policy Violations” on page 124.

If more than one action is configured for one class, the action with the highest priority occurs first. Priority is based on the order assigned to the action, with order 1 having the highest priority.

For example, a set of corrective actions can be:

Order 1: Redirect low class sessions 3 times

Order 2: Reject low class sessions

A low class session is redirected up to three times. If a session has been redirected three times, it is rejected.

It is recommended that the reject action be configured as the last action (lowest priority) for the class.

You can change the order of the corrective actions by moving their order using [Move Up] and [Move Down].

Threshold Policies

This screen displays the threshold policies configured for the site template.

Figure 4-16 Site Template Threshold Policies

A threshold policy is a type of SLO. It provides capacity protection for your system and site. Both threshold policies and SLOs are displayed in the service level objectives view in the management console's main screen.

When multiple thresholds and SLOs are violated, the corrective action with the highest priority is taken.

To configure a threshold policy, click [Add] or highlight an existing threshold policy and click [Edit].

Add/Edit Threshold Policies

This screen displays the threshold policies you can configure and the corrective actions configured to the threshold policies for the site template.

Figure 4-17 Add Thresholds Dialog Box

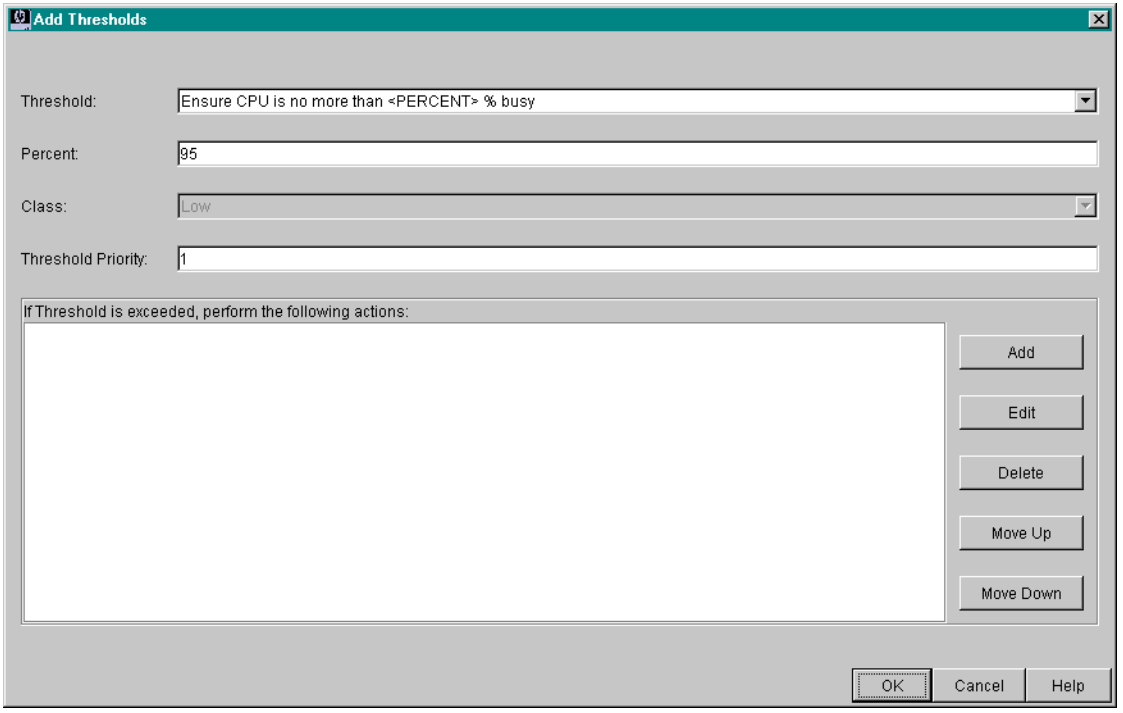

You can configure the following threshold policy:

- Ensure CPU is no more than <PERCENT> % busy

For more information about this threshold policy, refer to the online help and “Threshold Policy” on page 124.

If you wish to add corrective actions for when the threshold policy is violated, go to the next section. Otherwise click [OK].

Add/Edit Corrective Actions This screen displays the corrective actions you can configure for the site template.

Figure 4-18 Add/Edit Corrective Actions

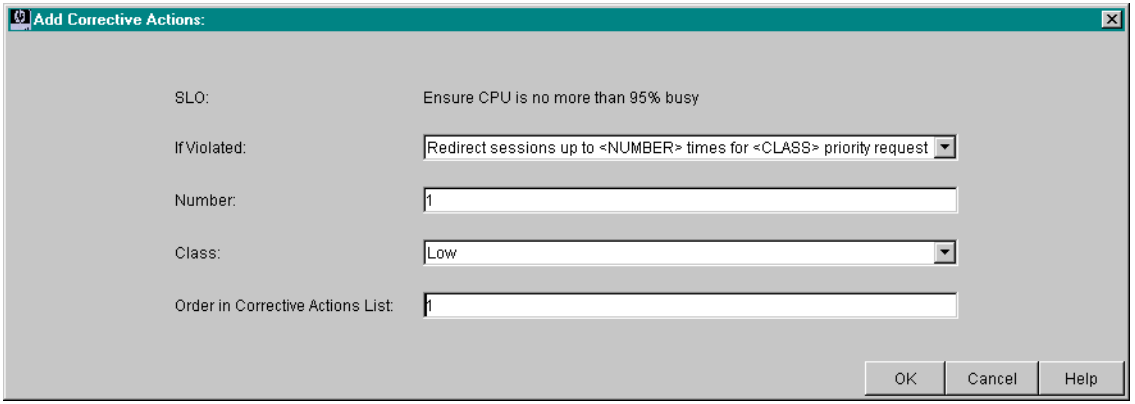

Corrective actions are those actions taken to bring a threshold policy into compliance. Corrective actions may involve making trade-offs between requests belonging to different request classifications, or making trade-offs between requests going to different sites or services running on the same system. You can add, edit or delete corrective actions by selecting the appropriate button.

You can configure the following corrective actions:

- Redirect sessions up to <NUMBER> times for <CLASS> priority requests
- Reject <CLASS> priority requests
- Defer <CLASS> priority requests

For more information about these corrective actions, refer to the online help and “Corrective Actions for SLO and Threshold Policy Violations” on page 124.

If more than one action is configured for one class, the action with the highest priority occurs first. Priority is based on the order assigned to the action, with order 1 having the highest priority.

For example, a set of corrective actions can be:

Order 1: Redirect low class sessions 3 times

Order 2: Reject low class sessions

A low class session is redirected up to three times. If a session has been redirected three times, it is rejected.

It is recommended that the reject action be configured as the last action (lowest priority) for the class.

You can change the order of the corrective actions by moving their order using [Move Up] and [Move Down].

Advanced

From this screen, you can configure session timeouts, web server accept threads, site template corrective actions, and enable/disable site template statistics logging.

Figure 4-19 Edit Site Configuration Advanced Dialog Box

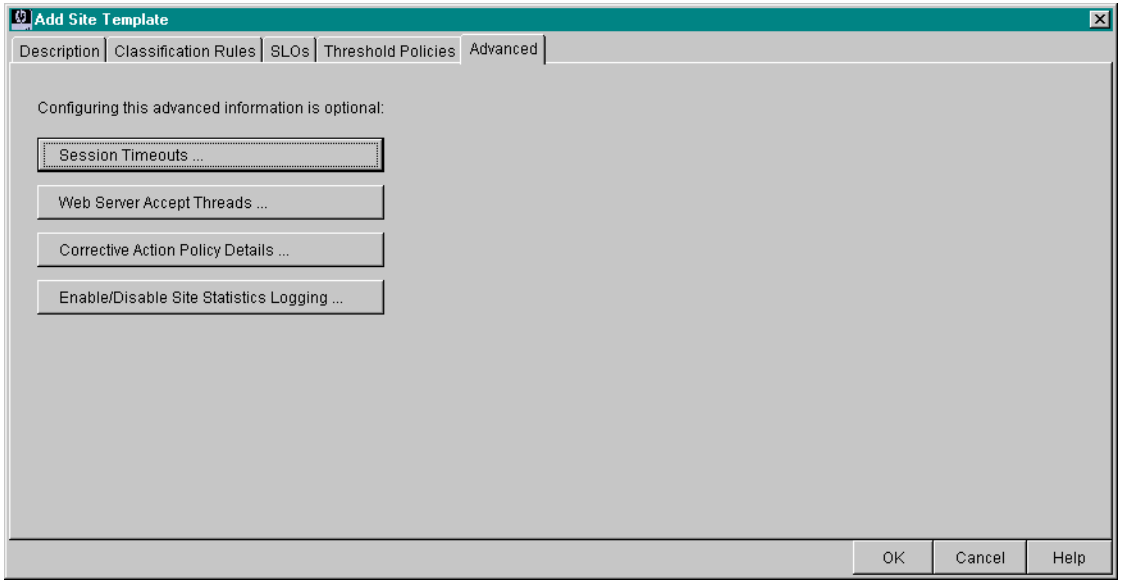

Session Timeouts

From this screen you can configure the maximum allowable time between requests and of a session.

Figure 4-20 Site Advanced Session Timeouts Dialog Box

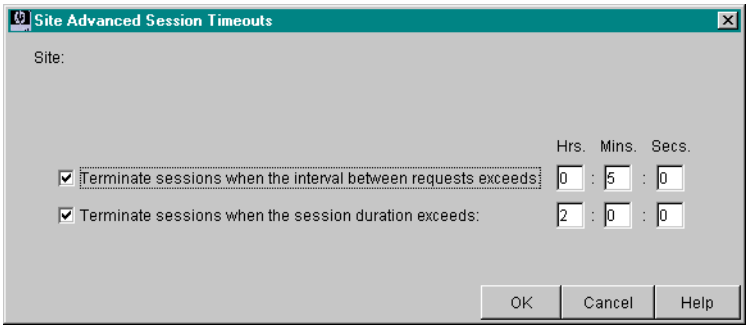

A session consists of all requests from a user within a certain time frame. Once a user has been granted a session, all requests from that user are forwarded to the web server, even when new sessions are no longer being accepted due to an SLO or threshold policy violation. If the session timeouts are exceeded, the current session ends and WebQoS starts a new session unless prevented by a corrective action.

If you have enabled URL encoding and configured the application server session identifier and BroadVision user name in the `qos.conf` file, you cannot configure the interval between requests session timeout.

At least one session timeout must be enabled. Both can be enabled simultaneously.

For more information about session timeouts, refer to the online help.

Web Server Accept Threads

From this screen you can configure the web server accept threads.

Figure 4-21 Site Advanced Accept Threads Dialog Box

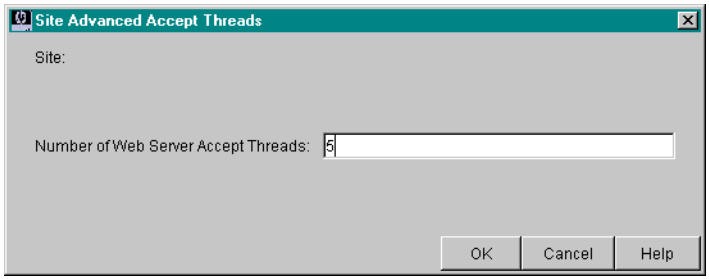

WebQoS web server accept threads are those WebQoS threads that accept incoming connections. The number of WebQoS accept threads affects the speed at which WebQoS can process requests. More accept threads could mean faster processing, but might also increase the load on your system.

Your web server(s) must be restarted in order for this change to take effect.

For more information about web server accept threads, refer to the online help.

Corrective Action Policy Details

From this screen, you can configure how a session is deferred, redirected, or rejected. You must have configured a corrective action to an SLO or threshold policy for these policies to take effect for the site template.

For information on customizing a defer and reject web page, refer to “Customizing the Defer and Reject Web Pages” on page 125.

Defer From this screen you can configure how a defer corrective action is executed for the site template.

Figure 4-22 Defer Corrective Action

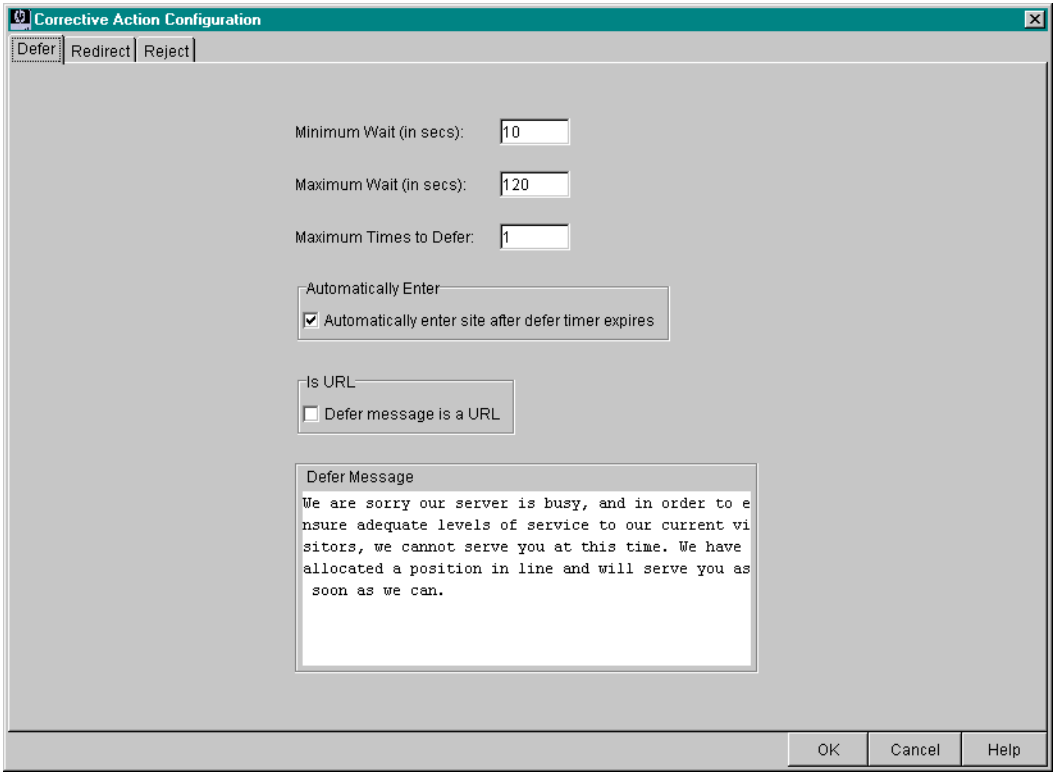

You can configure the minimum and maximum times a session is deferred, the maximum number of times a session is deferred, automatic entry after a defer, and the type of defer message. Refer to the online help for more information.

Redirect From this screen you can configure how a redirect corrective action is executed for the site template.

Figure 4-23 Redirect Corrective Action

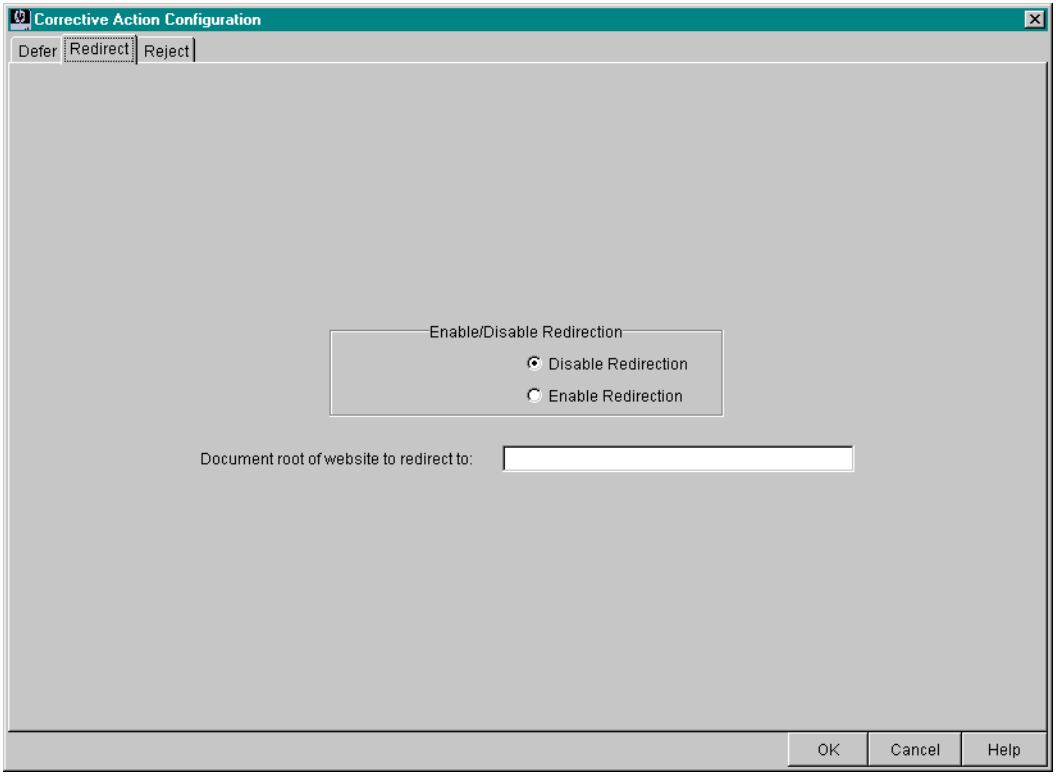

You can enable or disable redirection and configure the URL of a web site to which to redirect. Refer to the online help for more information.

If you have configured a redirect corrective action but disable redirection in this screen, all redirect corrective actions are ignored.

Reject From this screen you can configure how a reject corrective action is executed for the site template.

Figure 4-24 Reject Corrective Action

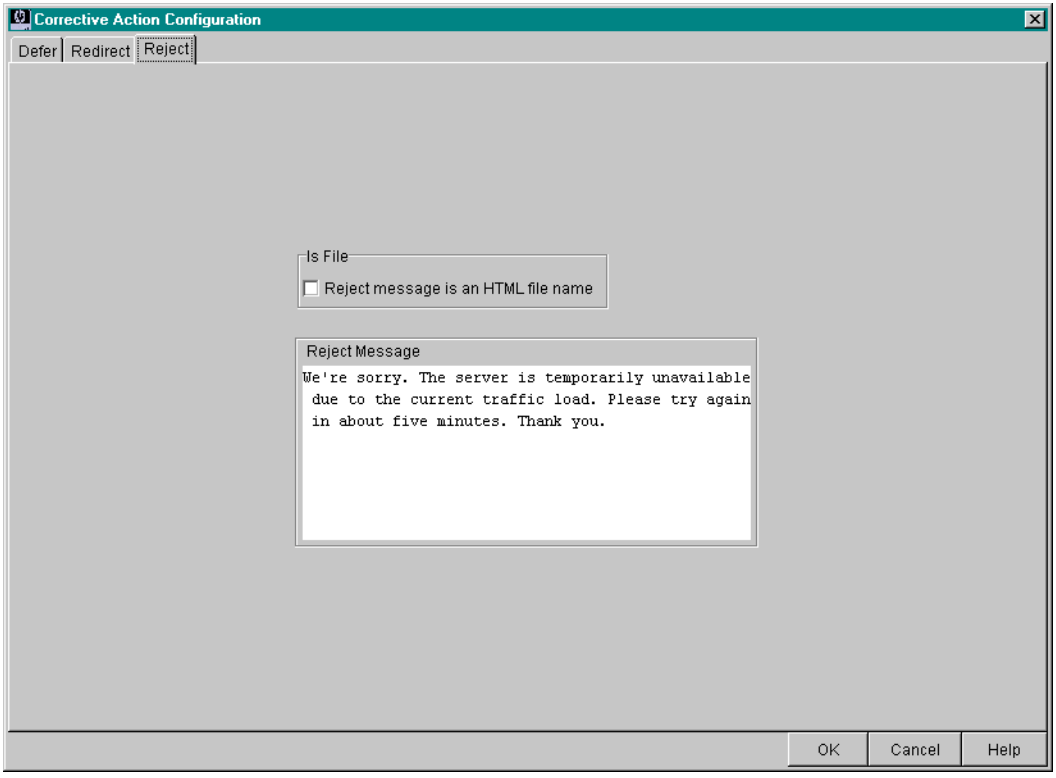

You can configure the message sent to a rejected session. Refer to the online help for more information.

Site Statistics

This screen displays whether or not site statistics logging is enabled for the site template.

Figure 4-25 Site Template Statistics

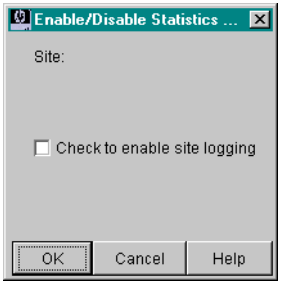

By default, site statistics logging is disabled.

If you enable site statistics logging, WebQoS logs site statistics to a log file named `<site_name>.stat` where *site_name* is the name of the site to which this template is assigned. The log file is located on the web services system (the server on which the SCA is installed)

The log file is created daily until you disable site statistics logging. At midnight, the file is moved to `<site_name>.stat.<previous_date>`.

In UNIX, the log files are located in the `/var/opt/webqos/stats` directory. In Windows, the log files are located in the `<install-directory>\stats` directory.

You can specify how often the statistics are generated by editing the `MeasurementInterval` parameter in the `qos.conf` file. The default is 30 seconds. For more information on the `MeasurementInterval` parameter, see “qos.conf” on page 23.

The following site statistics are logged:

Timestamp
Response High
Response Medium
Response Low
Admissions High
Admissions Medium
Admissions Low
Deferrals High
Deferrals Medium
Deferrals Low

Redirections High
Redirections Medium
Redirections Low
Rejections High
Rejections Medium
Rejections Low
Deferred Outstanding High
Deferred Outstanding Medium
Deferred Outstanding Low
Deferral Time Average High
Deferral Time Average Medium
Deferral Time Average Low
Deferral Time Current High
Deferral Time Current Medium
Deferral Time Current Low

5 Services, Sites, and Systems

This chapter describes services, sites, and systems.

Services

A service is a way to logically group web sites related to a business application. It allows you to view logical groups of web sites as a single entity from the management console.

Add a Service

To add a service, do the following:

1. Select `Actions: Add -> Service`.
2. Select the service template to use.
3. Enter a service name.
4. Select `[OK]`.

Edit a Service

You can edit a service's name and/or the service template to which it is configured. You cannot edit a service that is in the `not connected` state.

To edit a service, do the following:

1. Highlight the service in the services/systems view area.
2. Click on `Edit` in the task bar.
3. Once you have finished modifying the information, select `[OK]`.

Move a Service

You can move a service or multiple services to use another service template.

If you are moving a single service, you can either edit or move the service.

To move a service or multiple services, select `Actions: Move -> Services` and then refer to the online help.

Delete a Service

To delete a service, do the following:

1. Make sure all web sites contained by the service are in the `not connected` state (aqua color code). To achieve this, stop all the web sites of the service.
2. Highlight the service in the services/systems view area.
3. Select `Actions: Delete -> Service`.

NOTE

You cannot delete a service if it is used by the default site template. Either edit the default site template to use another service or define a new default site template.

If a site template other than the default site template uses the service you are about to delete, the template will be changed to use the service configured in the default site template.

4. Verify that you want to delete the service.

NOTE

The sites configured to the service are also deleted. However, any systems containing the service are not deleted.

Sites

A site is a web site or web server instance that is managed by WebQoS.

Add a Site

WebQoS automatically discovers systems and their web sites. For a web site to be automatically discovered, complete the following:

1. Run the `setup` script on the system on which the web site resides. The `setup` script only needs to be run once on the system. If you have already run `setup` on the system, go to the next step.
2. Make sure the web site is configured in the root pathname that you specified while running the `setup` script.
3. Start the web site.

Edit a Site

You can edit a site's name and/or the site template to which it is configured. You cannot edit a site that is in the `not connected` state (aqua color code).

To edit a site, do the following:

1. Highlight the site in the services/systems view area.
2. Click on `Edit` in the task bar.
3. Once you have finished modifying the information, select `[OK]`.

Move a Site

You can move a site or multiple sites to use another site template.

If you are moving a single site, you can either edit or move the site.

To move a site or multiple sites, select `Actions: Move -> Sites` and then refer to the online help.

Delete a Site

To delete a site, do the following:

1. Disable the web site. The site must be in the `not connected` state (aqua color code).
2. Highlight the site in the services/systems view area.
3. Select `Actions: Delete -> Site`.
4. Verify that you want to delete the site.

Microsoft IIS

This section discusses important points about starting and stopping Microsoft IIS.

Starting and Stopping IIS

On a Windows NT system, you can start and stop Microsoft IIS by going to the Control Panel and selecting **`Services`**. Then highlight **`IIS Admin Service`** from the list and press the Start or Stop button.

On a Windows 2000 system, you can start and stop Microsoft IIS by selecting **`Start -> Programs -> Administrative Tools -> Services`**. Then select **`IIS Admin Service`** from the list.

When you stop IIS, all web sites are stopped. Then, when you start IIS, the web sites are not started. Make sure that you start up *all* the web sites that were running when IIS was stopped.

In Windows 2000, if you use the `Restart` command from the `Action` menu, IIS and all web sites are restarted.

iPlanet Web Server

If you add a new instance of a web server on your system, you must run the setup script for the site to be automatically discovered.

If you edit a web server instance by either adding a new object type or enabling cgi-bin processing using the iPlanet Administration Server, you must add the following item in the web server's `obj.conf` file after the object type definition: **`Service fn="SendResponse"`**

Systems

A system is a computer with web sites managed by WebQoS.

Add a System

WebQoS automatically discovers systems and their web sites (see “Auto Discovery” on page 43 for more information). The systems must have the SCA (the Web Services component on UNIX and the Services for IIS component on Windows) installed and the web sites must be running.

In UNIX, to add a system, run the `setup` script on it. Once `setup` has been run, the system's sites (determined by the root pathname you entered in the `setup` script), when started, are also automatically added. For more information about the `setup` script, refer to “WebQoS setup Script (UNIX)” on page 30.

In Windows, to add a system, run the `W2KWebQoS.exe` installation file (if you downloaded the software from the Internet) or the `setup.exe` program on the CD and install the Services for IIS component on the system. For more information, refer to the release note.

Delete a System

You can delete a system from your WebQoS domain using the management console or by running the `setup` script.

If you are deleting a system running the SCO, do the following:

1. Stop the SCO on the system to be deleted. Refer to “Restarting Individual WebQoS Components” on page 30 for more information.
2. On another system in your WebQoS domain, using the `setup` script, install and enable an SCO. You must have one (and only one) SCO running in your WebQoS domain.
3. From the management console, highlight the system in the services/systems view area.
4. Select `Actions: Delete -> System`.
5. Verify that you want to delete the system.

If you are deleting a system running the SCA, do the following:

1. Stop the SCA. Refer to “Restarting Individual WebQoS Components” on page 30 for more information.
2. If this is the only system running an SCA in your WebQoS domain, you must enable an SCA on another system in your WebQoS domain. You must have at least one SCA running in your WebQoS domain.
3. From the management console, highlight the system in the services/systems view area.
4. Select `Actions: Delete -> System`.
5. Verify that you want to delete the system.

If you use the `setup` script, run `setup -r` on that system. Make sure that you have at least one SCA and only one SCO running in your WebQoS domain.

NOTE The sites configured to the system are also deleted.

6 Statistics

Service Statistics

HP WebQoS provides the following service statistics:

- SLO related - average response times per request class and total number of unexpired concurrent sessions
- Session control - percentage of new sessions admitted, deferred, redirected, or rejected for the service during the last 30 minutes
- Defer statistics - total number of outstanding deferred sessions, current defer time, and average defer time for the last 30 minutes

Service SLO Related Statistics

If you have enabled URL encoding and configured the application server session identifier and the BroadVision user name in the `qos.conf` file, the statistics displayed in the SLO Related Statistic screens are based on BroadVision's session counts.

To view the SLO related statistics, do the following:

1. Highlight a service in the services/systems view.
2. Choose `Statistics` from the task bar.

Figure 6-1 Service SLO Related Statistics

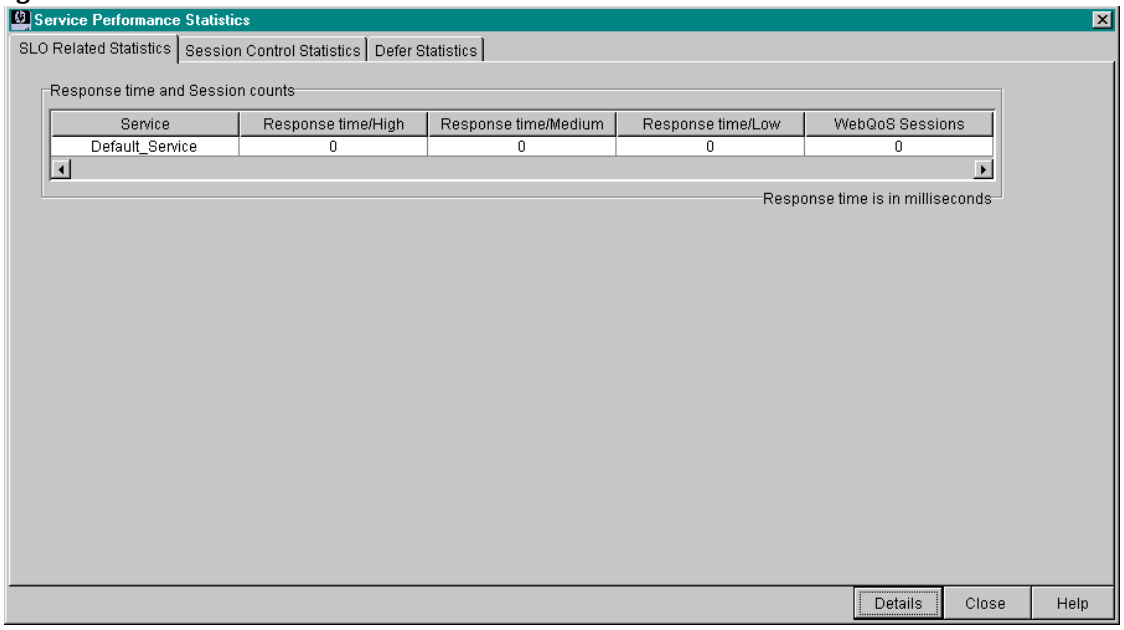

The following statistics are displayed:

- Response Time per request class - the amount of time, in milliseconds, the web server and web server application take to process a request. The response time does not include the time the request spends on the network between the browser and the web server.

  If the web server does not receive any requests during a two minute interval, the response time is displayed as zero.

- WebQoS Sessions - the total number of unexpired concurrent sessions. By default, a session expires after two hours or if the time between requests within a session exceeds 15 minutes. You configure the session timeouts under the Advanced tab of a site template.

There may be a short delay in reporting statistics while WebQoS recalculates the statistics.

To display SLO related statistics for each site in the service, choose `Details` in the `Service SLO Related Statistics` screen.

In the detailed screen, you can sort the information in the columns alphabetically or numerically by clicking on the column heading. For

example, to list the sites in the `Site` column alphabetically, click on the `Site` column heading. To reverse the order, click on the `Site` column heading again.

Service Session Control Statistics

To view the session control statistics, do the following:

1. Highlight a service in the services/systems view.
2. Choose `Statistics` from the task bar.
3. Click on the `Session Control Statistics` tab.

Figure 6-2 Service Session Control Statistics

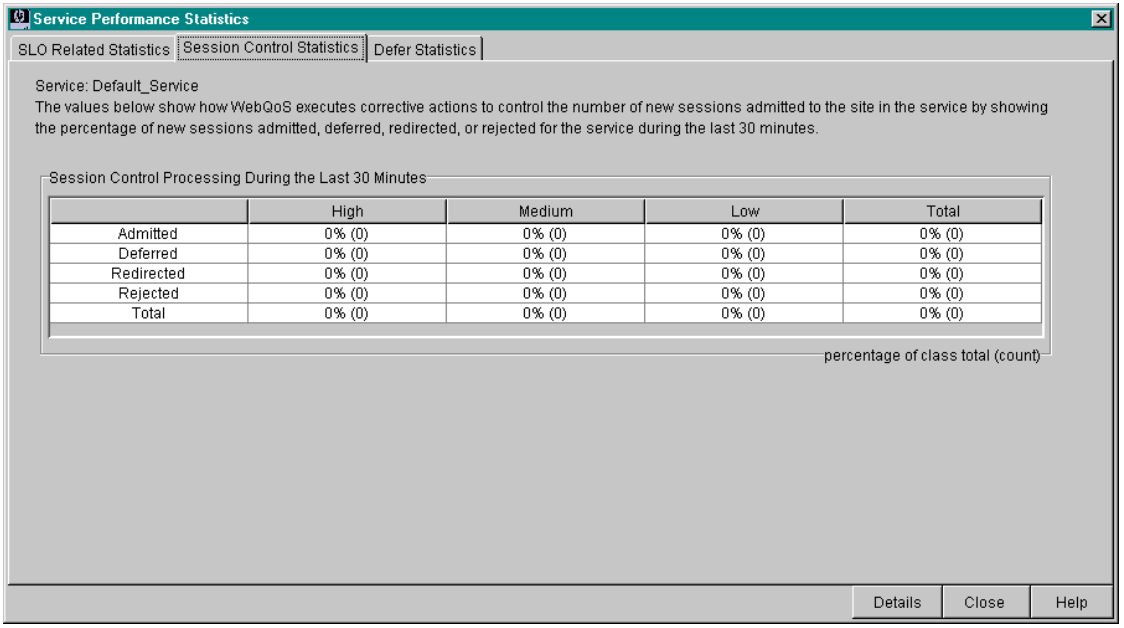

This screen shows how WebQoS executes corrective actions to control the number of new sessions admitted to the service.

Statistics displayed are the percentage of admitted, deferred, redirected, and rejected sessions, broken down by request classification, during the past 30 minutes. Each statistic also includes the number of sessions (shown in parentheses). For example, `90% (117)` means that 90% of all new sessions (or 117 new sessions) were admitted during the last 30 minutes.

There may be a short delay in reporting statistics while WebQoS recalculates the statistics.

To display session control statistics for each site in the service, choose `Details` in the `Service Session Control Statistics` screen.

In the detailed screen, you can sort the information in the columns alphabetically or numerically by clicking on the column heading. For example, to list the percentages in the `Total Admitted` column numerically from highest to lowest, click on the `Total Admitted` column heading. To reverse the order, click on the `Total Admitted` column heading again.

Service Defer Related Statistics

To view the defer related statistics, do the following:

1. Highlight a service in the services/systems view.
2. Choose `Statistics` from the task bar.
3. Click on the `Defer Statistics` tab.

Figure 6-3 Service Defer Related Statistics

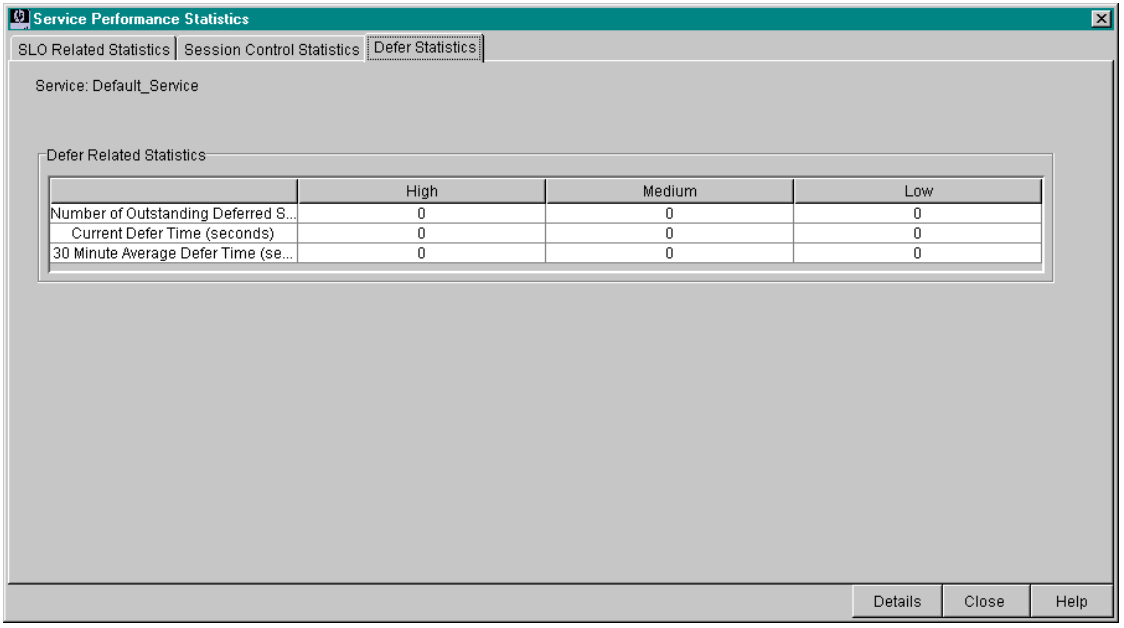

The following statistics are displayed:

- Number of Outstanding Deferred Sessions - the total number of deferred sessions that have not yet been admitted. This number is incremented every time a deferral is initiated.
- Current Defer Time - the amount of time, in seconds, the last deferred session must wait before being admitted. This information is logged every time a deferral occurs and is reset when the SLO is in compliance.
- 30 Minute Average Defer Time - the average amount of time, in seconds, deferred sessions in the last 30 minutes waited before being admitted. The data is collected and the average is updated every 30 seconds.

To display defer related statistics for each site in the service, choose Details in the `Service Defer Statistics` screen.

In the detailed screen, you can sort the information in the columns alphabetically or numerically by clicking on the column heading. For example, to list the sites in the `Site` column alphabetically, click on the `Site` column heading. To reverse the order, click on the `Site` column heading again. To sort the numbers in the `Sessions High` column from highest to lowest, click on the `Sessions High` column heading. To reverse the order, click again on the `Sessions High` column heading.

Site Statistics

HP WebQoS provides the following site statistics:

- SLO related - total number of unexpired concurrent sessions, average response times per request class, total number of outstanding deferred sessions, current defer time, and average defer time for the last 30 minutes
- Session control - percentage of new sessions admitted, deferred, redirected, or rejected for the service during the last 30 minutes

Site SLO Related Statistics

To view the SLO related statistics, do the following:

1. Highlight a site in the services/systems view.
2. Choose Statistics from the task bar.

Figure 6-4 Site SLO Related Statistics

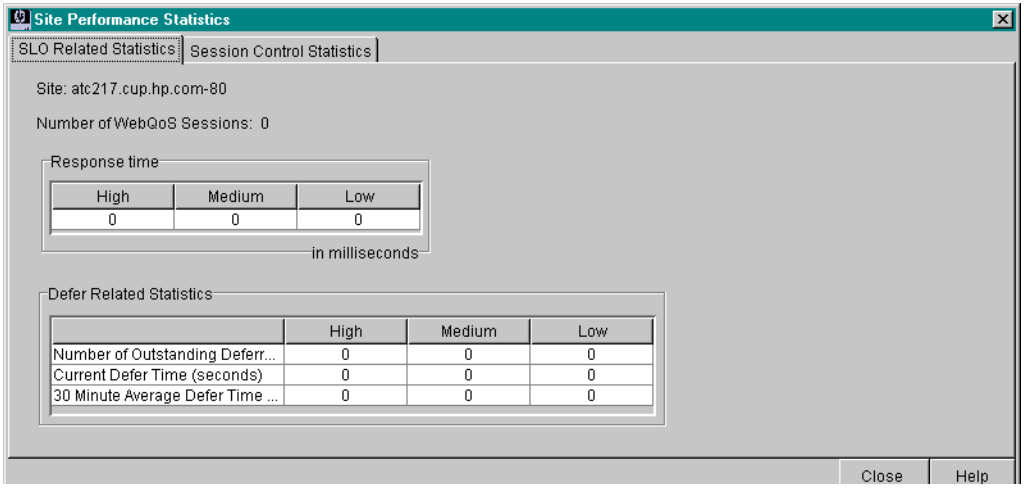

The following statistics are displayed:

- WebQoS Sessions - the total number of unexpired concurrent sessions. By default, a session expires after two hours or if the time between requests within a session exceeds 15 minutes. You configure the session timeouts under the Advanced tab of a site template.
- Response Time per request class - the amount of time, in milliseconds, the web server and web server application take to process a request. The response time does not include the time the request spends on the network between the browser and the web server.

  If the web server does not receive any requests during a two minute interval, the response time is displayed as zero.
- Number of Outstanding Deferred Sessions - the total number of deferred sessions that have not yet been admitted. This number is incremented every time a deferral is initiated.
- Current Defer Time - the amount of time, in seconds, the last deferred

session must wait before being admitted. This information is logged every time a deferral occurs and is reset when the SLO is in compliance.

- 30 Minute Average Defer Time - the average amount of time, in seconds, deferred sessions in the last 30 minutes must wait before being admitted. The data is collected and the average is updated every 30 seconds.

There may be a short delay in reporting statistics while WebQoS recalculates the statistics.

Site Session Control Statistics

If you have enabled URL encoding and configured the application server session identifier and the BroadVision user name in the `qos.conf` file, these statistics cannot be displayed.

To view the session control statistics, do the following:

1. Highlight a site in the services/systems view.
2. Choose `Statistics` from the task bar.
3. Click on the `Session Control Statistics` tab.

Figure 6-5 Site Session Control Statistics

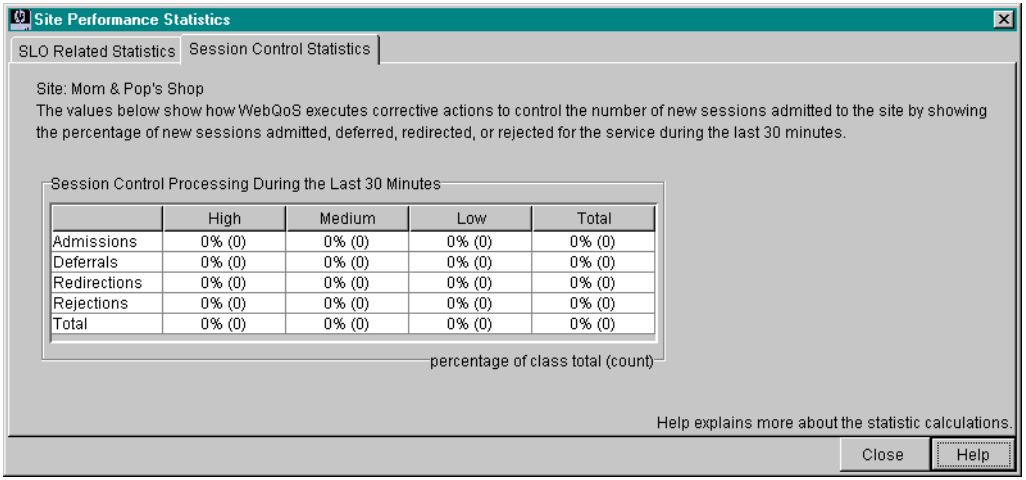

This screen shows how WebQoS executes corrective actions to control the number of new sessions admitted to the site.

Statistics displayed are the percentage of admitted, deferred, redirected, and rejected sessions, broken down by request classification, during the past 30 minutes. Each statistic also includes the number of sessions (shown in parentheses). For example, `90% (117)` means that 90% of all new sessions (or 117 new sessions) were admitted during the last 30 minutes.

There may be a short delay in reporting statistics while WebQoS recalculates the statistics.

7 Logs and Filters

SLOs, Thresholds, and the History Log

You can use the SLO and the history log views to do the following:

- Verify your WebQoS configuration
- Check that your topology and web server configuration meet objectives
- Determine when objectives are not met so you can make changes to your WebQoS configuration, topology, or web server configuration to better meet the objectives
- Determine when objectives are being met so you and your customers know that their requirements are being fulfilled

The Service Level Objectives View

The service level objectives view displays the status of service level SLOs in the top part of the window, and the status of site SLOs and thresholds in the bottom part of the window. You can rearrange the columns if you wish by clicking on the column title and dragging it to its new position.

Figure 7-1 Service Level Objectives View

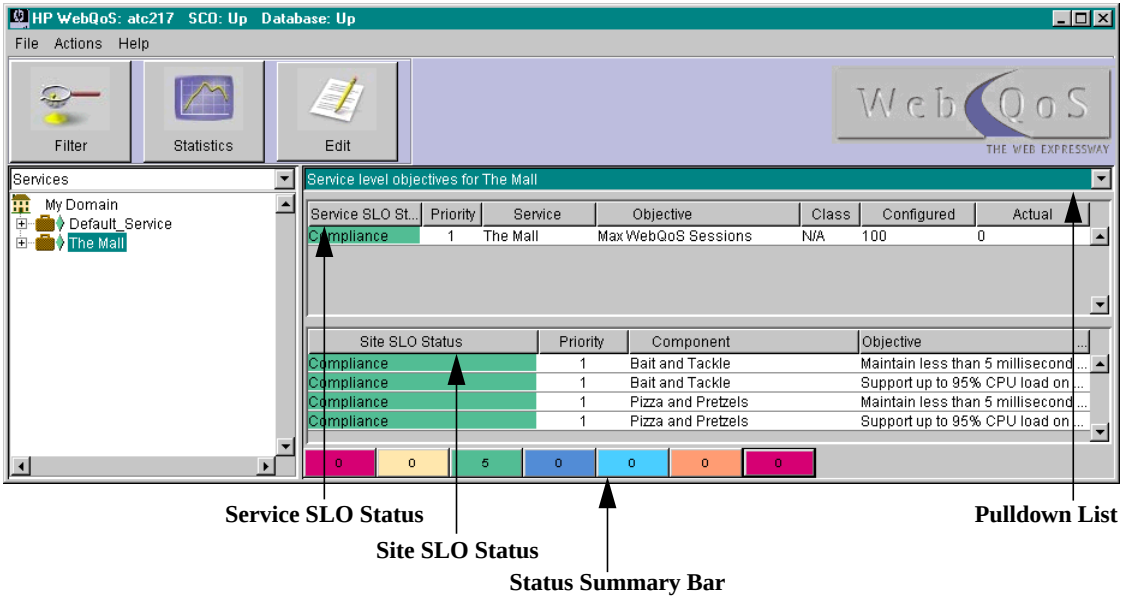

- Status column - Displays the color-coded status associated with SLOs and thresholds.
  - — Red: Violation. The SLO or threshold is in violation and no corrective actions are being executed by WebQoS.
  - — Yellow: At risk. The SLO or threshold is in violation, but WebQoS is executing a corrective action to try and bring it back into compliance.

    Only one SLO or threshold can be yellow on any given system, as only one corrective action is executed at a time. If you select the Services View, the SLO status for SLOs on several systems (if they are configured) would be shown.
  - — Green: In compliance. The SLO or threshold is being met.
  - — Blue: Inactive. The SLO is inactive; WebQoS is disabled.
  - — Aqua: Not connected. The SLO is not currently connected to the system, or that the system is not connected to a web server. For a site, the web server is not connected to the SCA. For a service, the SCO is not running.

    The management console updates the status display every 30 seconds. An SLO can show a status of “not connected” if the display has not yet been updated. Click on another site, service, or system to get an updated status.
  - — Orange: Never connected. The SLO was never connected. When you configure a site, the WebQoS configuration database and the web server running that site are notified. If the web server is off-line, the WebQoS Service Control Agent (SCA) tries to contact the site later. In the mean time, the site status is defined as “never connected.”
  - — Red: In error. The SLO or threshold is in error. This is most likely a state for the site that is inherited by the SLOs. It usually means the web server is not functioning properly. For example, the web server is unable to start up or shut down completely.

NOTE

If there is a problem with the web server, for example if the web server is unable to start or shut down completely, the SLOs inherit this problem and the SLO status is represented with a blue, aqua, orange, or red color code.

- Status summary bar - Displays the number of SLOs and thresholds of each status. The status summary bar is color-coded with the same scheme as the status column described above.
- Priority column - Displays the SLO or threshold priority (for example 1, 2, 3, or 4). WebQoS performs corrective actions defined for the highest priority (1) SLO or threshold policy that is currently violated.
- Component column - Displays the site associated with the SLO or threshold.
- Objective column - Displays a one line description of the SLO or threshold.

The History Log View

The history log view displays a history of SLO, threshold, and corrective action events for the site or service selected in the service/systems view.

NOTE The management console uses the local time zone. If the system where you installed the management console is in a different time zone than the web server and application server, the time stamps are formatted using the local time of the management console.

The history log shows which corrective actions have been taken. The history log entries have time stamps so you can tell how long an SLO or threshold has been violated.

You can use the history log to determine performance trends. You can see how often the objectives of a system, service, or site have not been met over time (for example, during the past month).

If you modify an SLO, all of its events in the history log are deleted.

When a service-level SLO is violated, all corrective actions configured for that SLO are entered into the history log.

Figure 7-2 History Log View

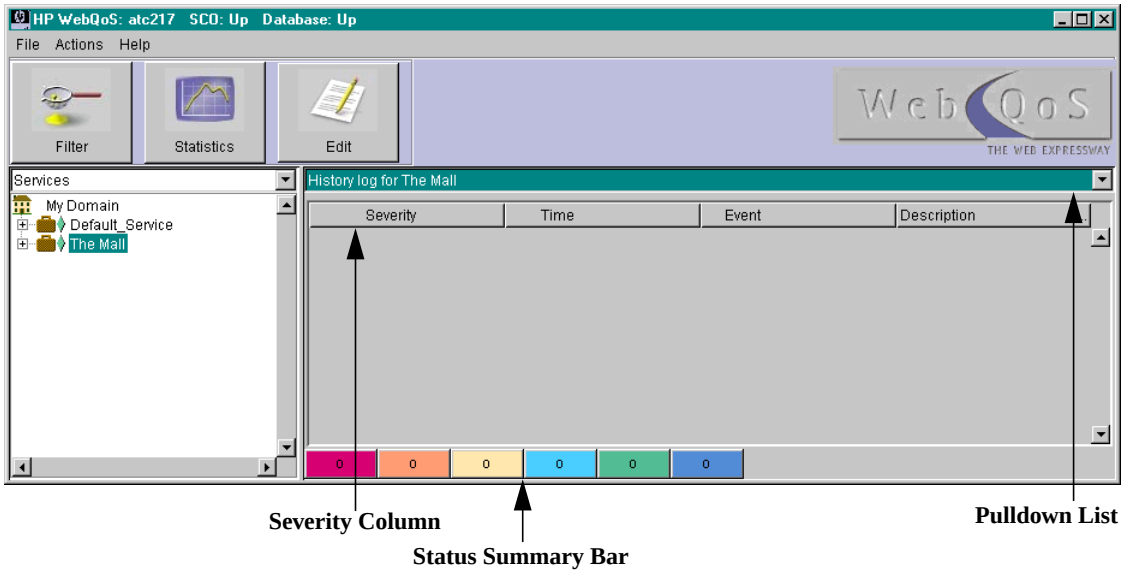

- Severity column - Displays the color-coded severity associated with the events.
  - — Red: Number of critical log entries
  - — Orange: Number of major log entries
  - — Yellow: Number of minor log entries
  - — Cyan (light blue): Number of warning log entries
  - — Green: Number of normal log entries
  - — Blue: Number of unknown log entries
- Status summary bar - Displays the number of events of each status. The status summary bar is color-coded with the same scheme as the severity column described above.
- Time column - Displays the time and date an event occurred.
- Event column - Displays the event type. The following event types can be displayed:
  - — SLO or threshold violation
  - — SLO or threshold risk

  - — Action taken
  - — CPU limit
  - — CPU reset
  - — Disk limit
  - — Disk reset
- Description column - Displays a one line description of the event.

SLO and Threshold Filters

You may be managing a large web complex and you may have a significant number of sites. SLO filtering lets you view SLOs and thresholds selectively to make problem solving easier. WebQoS can only display 50 SLOs and thresholds. You can use the filtering and sorting options to do the following:

- Determine which SLOs or thresholds to display. For example, you might want to display only priority one SLOs and thresholds.
- View SLO and threshold status for an important web site.
- View which SLOs and thresholds are in violation.

To filter your SLOs and thresholds, do the following:

1. Highlight the service, system, or site in the services/systems view and select `Service Level Objectives` from the SLO/history log pull down list.
2. Choose `Filter` from the task bar.

Figure 7-3 SLO Filters

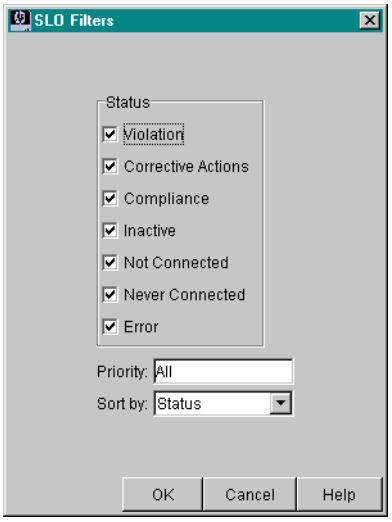

3. Select the appropriate `status` classifications to display SLOs and thresholds with the chosen status. By default, all types are selected. For example, you might choose to display only SLOs and thresholds that are being violated or in error. See the online help for descriptions of the status classifications.
4. Specify the appropriate `Priority` classification. SLOs and thresholds with the requested priority are displayed. These are the priorities that you assigned to the SLO which determine the order the SLOs are checked. For example, you may want to display only priority 1 SLOs and thresholds. The default is to display all priorities.
5. Sort your list of SLOs and thresholds as desired by selecting the appropriate sort item from the `Sort by:` pull down list. For example you may want to sort your SLOs and thresholds by priority.
6. Choose `[OK]`.

History Log Filter

History log filtering allows you to view history log entries selectively to make analysis and problem solving easier. WebQoS can only display 50 history log entries at a time. If there are more than 50 log entries for a service, system, or site, you can use filtering to specify which 50 entries should be displayed. Filtering and sorting the history log helps you to see patterns in the behavior of a service, system, or site.

To filter the history log, do the following:

1. Highlight the system, service or site in the services/systems view and select `History log` from the SLO/History log pull down list.
2. Choose `Filter` from the task bar.

Figure 7-4 History Filters Dialog

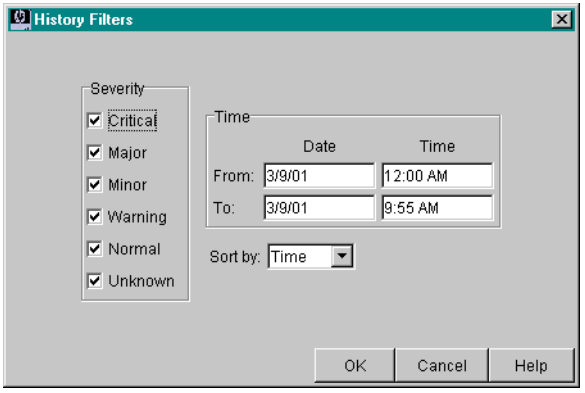

3. Select the appropriate `severity` classifications to display events with the chosen severity. For example, you might choose to display only events that are critical and major. See the online Help for descriptions of the severity classifications.
4. Type the date/time range to display history log entries which occur during this range. (The date format is MM/DD/YY and the time format is HH:MM AM or PM.) For example, you can use the time filter in combination with the severity filter to determine if any of the site's SLOs have been violated in the last month.

NOTE The management console uses the local time zone. If the system where you installed the management console is in a different time zone than the web server and application server, the time stamps are formatted using the local time of the management console.

5. Sort the history log as desired by selecting the appropriate sort item from the `Sort by:` pull down list. For example, you can sort events by severity.
6. Choose `[OK]`.

OpenView Logs

Log files created by WebQoS can also be used by OpenView. These log files keep track of service and site statistics, site measurements, and policy and corrective action events.

Service and Site Statistics

OpenView can use the information logged by a service or site template. Logging must be enabled in the template for this information to be available. In UNIX, the log files are located in the `/var/opt/webqos/stats` directory. In Windows, the log files are located in the *<install-directory>*`\stats` directory.

If you enable logging in a service template, use the following log file(s): *<service_name>*`.stat` and/or *<service_name>*.stat.*<previous_date>*. Refer to “Service Statistics” on page 50 for a list of statistics that are logged.

If you enable logging in a site template, use the following log file(s): *<site_name>*`.stat` and/or *<site_name>*`.stat.`*<previous_date>*. Refer to “Site Statistics” on page 74 for a list of statistics that are logged.

Site Measurements

To log site measurements to a log file, on the system running a web server managed by WebQoS, do the following:

1. Enable the *LogOn* parameter in the `qos.conf` file.
2. Restart the SCA.

Site measurements are logged to *<site_name>*.*<date>*. In UNIX, the log files are located in the `/var/opt/webqos/measures` directory. In Windows, the log files are located in the *<install-directory>*`measures` directory.

The following statistics are listed as a total and broken out by class:

Timestamp
Average number of active sessions
Average response time
Admitted sessions
Redirected sessions

Deferred sessions
Rejected sessions
Average number of deferred sessions
Average defer time

Policy and Corrective Action Events

To log policy and corrective action events, on the system running the SCO, do the following:

1. In UNIX, create the subdirectory `OV` in `/var/opt/webqos`.
2. Add **`SCOLogITOOn 1`** to the `qos.conf` file.
3. Restart the SCO.

In UNIX, policy and corrective action event log files are located in the `/var/opt/webqos/OV` directory. In Windows, the log files are located in *<install-directory>*\OV.

Service policy events are logged to the file `servicepolicyevents`.

The following information is logged:

Date
Time
Service name
Event type
SLO status
SLO type
SLO class
Name

Site policy events are logged to the file `sitepolicyevents`.

The following information is logged:

Date
Time
System name
Site name
IP address
Port number
Event type
SLO status
SLO type
SLO class

Service corrective events are logged to the file `servicecorractevents`.

The following information is logged:

Date
Time
Service name
Event type
Corrective action
Class

Site corrective action events are logged to the file `sitecorractevents`.

The following information is logged:

Date
Time
System name
Site name
IP address
Port number
Event type
Corrective action
Class

Events are logged every 15 seconds. The following fields may have the following values (a value of "-" means that the information is not available):

Field	**Values**
(SLO) Class	• `low` • `medium` • `high` • `-`
Corrective Action	• `redirect` • `reject` • `defer` • `-`
Event Type	• `cor` (corrective action) • `pol` (policy)
Name	• URL group name if the SLO type is url_response_time • external measurement name if the SLO type is external_measure • `-`

Field	Values
SLO Status	• violation • risk • compliance • -
SLO Type	Service SLOs: • external_measure (Ensure <Measurement> is <CONDITION> <NUMBER>) • session_count (Support at most <NUMBER> WebQoS sessions) • url_response_time (Ensure less than <NUMBER> millisecond response time for <CLASS> requests to <URL GROUP>) • - Site SLOs: • cpu_load (Ensure CPU is no more than <PERCENT>% busy) • max_sessions (Support at most <NUMBER> WebQoS sessions) • min_sessions (Support at least <NUMBER> WebQoS sessions) • response_time (Maintain less than <NUMBER> millisecond avg response time for <CLASS> priority requests) • -

8 Troubleshooting

Troubleshooting is broken down by platform: UNIX or Windows. Refer to the appropriate section.

Troubleshooting on UNIX

Make sure you have installed all required patches. Refer to the release note for a list of these patches.

Shared Memory

The following warning in the /var/opt/webqos/logs/scalog file:

`Couldn't attach to shared memory`

means the SCA cannot function correctly without memory segments.

Do the following:

- run **ipcs** to list memory segments, and then run **ipcrm -m** to delete memory segments.

Management Console Error Messages

The following error messages are viewable in the management console:

- `The SCI on <machine> on port <number> is not running.`

  The WebQoS SCI is down. Restart the SCI by running **/opt/webqos/bin/sci_start** and then reconnect to the management station.

NOTE You must be root to run **/opt/webqos/bin/sci_start**

- `Problems writing to the SCI socket`

  The WebQoS SCI is down. Restart the SCI by running **/opt/webqos/bin/sci_start** and then reconnect to the management station.

NOTE You must be root to run **/opt/webqos/bin/sci_start**

- `Problems connecting to the SCI on <machine>. Unknown hostname`

The machine that contains the SCI may be down or the network is unreachable. Also, the intended machine may be mistyped or unknown to the name server. Try the following steps:

1. Verify network connectivity. Run `/usr/sbin/ping <machine>`. Reference the *ping(1M)*man page for more information.
2. Verify the name server configuration for your network. For example, if you are using DNS or `/etc/hosts`, run `/usr/sbin/nslookup <machine>`. Reference the *nslookup(1)* man page for more information.

- `This new threshold put the set of thresholds out of order. At least one threshold won't be violated until a higher priority threshold is violated. You may wish to edit it later to change its priority.`

  Threshold rules are misconfigured.

  Ensure that the lower priority threshold rule is violated after the higher priority rule.

Error Logging and Tracing

The WebQos components use the following log files to record WebQoS related activity. You can use these files to monitor the WebQoS server and help with troubleshooting. The files are:

```
/var/opt/webqos/logs/qoslog
/var/opt/webqos/logs/qoslog.child
/var/opt/webqos/logs/qostrace
/var/opt/webqos/logs/qostrace.child
/var/opt/webqos/measures
/var/opt/webqos/logs/scalog
/var/opt/webqos/logs/scatrace
/var/opt/webqos/logs/scilog
/var/opt/webqos/logs/scitrace
/var/opt/webqos/logs/scolog
/var/opt/webqos/logs/scotrace
```

Logging

Error messages are logged in the `qoslog`, `qoslog.child`, `scalog`, `scilog`, and `scolog` files. You can set the level of logging (number and types of messages) by setting the parameters in the file `/etc/opt/webqos/qos.conf`. For information on how to set the logging level, see "Logging and Tracing Parameters (UNIX Only)" on page 28. More messages will be logged by changing the levels. See "Changing Logging and Tracing Levels" on page 112 for more information.

Tracing

The `qostrace`, `qostrace.child`, `scatrace`, `scitrace`, and `scotrace` files are for use by HP Support.

Measurement Logging

If you have set the *LogOn* parameter in the `/etc/opt/webqos/qos.conf` file to on, WebQoS creates a new measurement log file daily for each WebQoS site and places them in `/var/opt/webqos/measures`. The measurement log file name is the site name that was entered in the management console, followed by the date. For information on how to set the parameter, see "SCA Parameters" on page 26.

Statistics Logging

If you have enabled service and site statistics logging in the management console, WebQoS generates statistics and creates files in `/var/opt/webqos/stats`.

The statistics log file names are the site name or service name that you entered in the management console followed by the date. For information on how to enable Statistics Logging, and for a description of which statistics are generated, see "Statistics" on page 50 for service statistics logging, and "Site Statistics" on page 74 for site statistics logging.

Changing Logging and Tracing Levels

The product ships with default levels. For logging, the default level is informative messages. Tracing is turned off. It is recommend that you run the product with these levels. Turning on more detailed levels will quickly fill up the log files. The levels are set in the `/etc/opt/webqos/qos.conf` file. There is more information on the levels in this file. To change the levels and see more detail do the following:

1. Edit the `/etc/opt/webqos/qos.conf` and set the level.
2. Restart the WebQoS components.

For more information on modifying the parameters in the `/etc/opt/webqos/qos.conf` file, see “qos.conf” on page 23.

Troubleshooting on Windows

The WebQoS management console displays the status of the service or site SLO. If there is a problem with the web server, the SLO status will likely reflect this problem. For example, if the Status column of the Service Level Objectives view shows “Never Connected” (orange color), the web server may be offline. In this case, check the web server status. See “The Service Level Objectives View” on page 96 for descriptions of the Status column entries and color codes.

Problems with Shared Memory

If you see the following warning in the *<install-directory>*\logs\qoslog.txt file:

- `Couldn’t attach to shared memory`

  The Service Control Agent (SCA) cannot function correctly without memory segments. Restart the SCA (see “Restarting Individual WebQoS Components” on page 30).

Resetting the Management Console Password

When resetting the management console password (see “Password” on page 44), you may receive the following error message after typing the command:

- `Exception in thread “main” java.util.zip.ZipException: The system cannot find the file specified`

  ```
  at java.util.zip.ZipFile.open(Native Method)
  at java.util.zip.ZipFile.<init>(Unknown Source)
  at javautil.jar.JarFile.<init>(Unknown Source)
  at java.util.jar.JarFile.<init>(Unknown Source)
  ```

Make sure you are in the *<install-directory>* before typing the command.

Error Messages Viewable From the Management Console

The following error messages appear in popup windows in the management console:

- `The SCI on <machine> on port <number> is not running.`

  The WebQoS SCI is down. You need to exit the management console, restart the SCI on the WebQoS Management Server (see “Restarting Individual WebQoS Components” on page 30), then restart the management console.

- `Problems writing to the SCI socket.`

  The WebQoS SCI is down. You need to exit the management console, restart the SCI on the WebQoS Management Server (see “Restarting Individual WebQoS Components” on page 30), then restart the management console.

- `Problems connecting to the SCI on <machine>. Unknown hostname.`

  The machine that contains the SCI may be down or the network is unreachable. Also, the intended machine may be mistyped or unknown to the name server. Do the following:

  1. Verify network connectivity. Run `ping <machine>`. If there is no response, then check the Network connections from the Control Panel.
  2. Verify the name server configuration for your network in the Control Panel

- `This new threshold put the set of thresholds out of order. At least one threshold won’t be violated until a higher priority threshold is violated. You may wish to edit it later to change its priority.`

  Misconfiguration of threshold rules.

  Fix: Ensure that the lower priority threshold rule is violated after the higher priority rule.

Logging and Tracing

The WebQos components use the following log files to record WebQos related activity. You can use these files to monitor the WebQoS server and help with troubleshooting. The files are:

<install-directory>\logs\qoslog.txt
<install-directory>\logs\qostrace.txt
<install-directory>\logs\scalog

Logging

Error messages are logged in the *<install-directory>*\logs\qoslog.txt and scalog files. The default logging level is informative messages. More messages will be logged by changing the level. See "Changing Logging and Tracing Levels" for more information.

Tracing

The *<install-directory>*\logs\qostrace.txt file is for use by HP Support.

Changing Logging and Tracing Levels

The product ships with default levels. For logging, the default level is informative messages. Tracing is turned off. It is recommended that the product be run with these levels. Turning on more detailed levels will quickly fill up the logfiles.

WebQoS tracing and logging levels are stored in the Windows Registry. You can change WebQoS tracing and logging levels with the Windows Registry Editor.

To change the WebQoS tracing or logging level:

1. From the Windows Start menu, choose the Run option.
2. Run the Registry Editor program by entering **regedt32.exe** in the Run window.
3. Click on the Registry Editor window HKEY_LOCAL_MACHINE.
4. Click on the icons for Software, Hewlett-Packard, then WebQoS.

Figure 8-1 Registry Editor Window for WebQoS

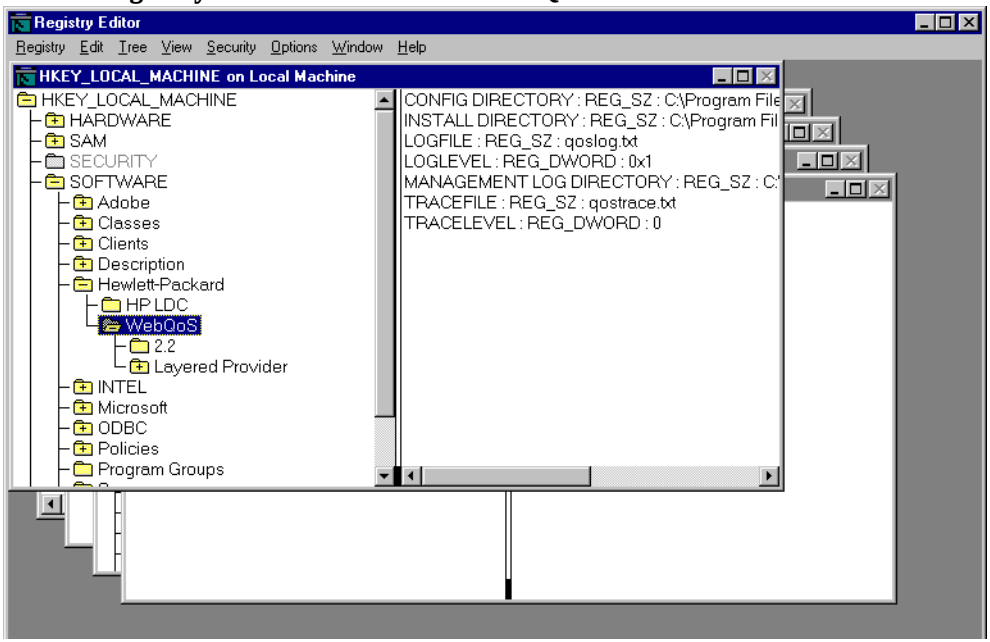

5. In the left side of the window, double-click on the LogLevel or TraceLevel values.
6. In the popup window, enter the new value for the selected item.

   The default tracing level is 0, which means no tracing. A value from 0-3 can be specified; the level of tracing increases with each value. The default value for logging is 3, which means informative logging. A value from 0-4 can be specified; the level of logging increases with each value.

Deleting WebQoS Log and Trace Files

The size of the WebQoS log and trace files is fairly small, unless you turn on detailed log or trace levels. If you need to delete log or trace files, you must first stop Microsoft IIS, then stop the WebQoS SCA and SCO (see “Restarting Individual WebQoS Components” on page 30). You can then delete the WebQoS log and trace files.

A Policy and Rule Descriptions

This appendix describes request classification rules, service level objectives, and threshold policies.

Request Classification Rules

In Windows and Solaris, if there is no request classification applied to a request, the default class is low for SSL *and* non-SSL requests.

If more than one class rule match is found in the request, the highest classification is chosen.

IP masking is allowed when entering an IP address in a request classification rule. That is, you can enter an * (asterisk) instead of a number in the last three octets of the IP address. The * represents the octets from 0 to 255 and must be entered in consecutive positions to the last octet. For example, if you enter the IP address of 1.2.3.*, you are referring to IP addresses 1.2.3.0 to 1.2.3.255. If you enter the IP address of 1.2.*.*, you are referring to IP addresses 1.2.0.0 to 1.2.255.255. If you enter the IP address of 1.*.*.*, you are referring to IP addresses 1.0.0.0 to 1.255.255.255. The following are NOT VALID IP addresses using IP masking: 1.*.3.4, 1.*.3.*, 1.*.*.4, and 1.2.*.4.

The request classification rules supported by WebQoS are listed below.

- `Requests to URL document path PATH are CLASS` (default)

  You need to enter a URL document path in the `Path` input box. A request classification (`High`, `Medium`, or `Low`) must be selected. This rule then assigns the chosen class to all requests going to the URL document path defined in the `Path` input box. Refer to the online help for information about using wildcards in the URL document path.

- `Requests to IP address ADDRESS are CLASS`

  You need to enter an IP address in the `Address` input box. A request classification (`High`, `Medium` or `Low`) must be selected. This rule then assigns the chosen class to all requests going to the IP address defined in the `Address` input box. This option is most useful in hardware virtual server configurations.

- `Requests from IP address ADDRESS are CLASS`

  You need to enter an IP address in the `Address` input box. A request classification (`High`, `Medium` or `Low`) must be selected. This rule then assigns the chosen class to all requests coming from the IP address defined in the `Address` input box.

- `Requests to port number NUMBER are CLASS`

  You need to enter a port number in the `Port Number` input box. A request classification (`High`, `Medium`, or `Low`) must be selected. This rule then assigns the chosen class to all requests going to the `Port Number` specified in the input box.

- `Requests to virtual server NAME are CLASS`

  You need to input a virtual server name in the `Name` input box. A request classification (`High`, `Medium` or `Low`) must be selected. This rule then assigns the chosen class to all requests going to the virtual server defined in the `Name` input box.

  The virtual server concept allows you to configure several “sites” on a single http port. All the virtual sites are handled by the same web server instance, but they appear to the user to be distinct sites.

  This rule is only supported if the `HTTPVersion` web server parameter in the `qos.conf` file is set to HTTP/1.1 (this is the default setting).

Policies Set in WebQoS

There are two types of policies for WebQoS: business-oriented SLOs and operations-oriented Threshold Policies. You define these two types of policies and prioritize them relative to one another. WebQoS trades off meeting these policies based on their priorities.

WebQoS uses priorities to determine which corrective actions to take when one or more of the SLOs or Threshold policies are violated. When multiple policies are violated, WebQoS executes the corrective actions configured for the highest priority policy.

The scope of priorities apply across both SLOs and thresholds.

SLOs and thresholds involving the CPU are system-wide. All other SLOs and thresholds are site-specific.

Service Level Objectives for Site

Service level objectives (SLOs) are measures that are typically negotiated with a customer or business unit and specify the level of service they expect you to provide. WebQoS helps you to meet these expectations by monitoring compliance and taking corrective actions. SLOs put limits on response time delays and minimum concurrent session capacity.

The SLOs supported by WebQoS are:

- `Maintain less than NUMBER millisecond avg response time for CLASS priority requests`

  You can choose an appropriate `NUMBER` to input for the average response time (in milliseconds) for each request classification (High, Medium or Low) supported by the site.

- `Support at least NUMBER WebQoS sessions`

  This SLO lets you support a specific `NUMBER` (for example, 20) of concurrent sessions. It is tied to the average response time SLO described in the previous bullet. It can only be violated if one or more SLOs or thresholds are violated. The request classification `CLASS` information is not supported.

  Concurrent sessions means the total number of sessions that have not

expired (based on the configuration of the advanced features discussed in “Session Timeouts” on page 69) and are possibly active. This number does not indicate the total number of sessions currently active on the system.

- `Support at most NUMBER WebQoS sessions`

  This SLO lets you support a specific `NUMBER` (for example, 1000) of concurrent sessions. The request classification `CLASS` information is not supported.

  Concurrent sessions means the total number of sessions that have not expired (based on the configuration of the advanced features discussed in “Session Timeouts” on page 69) and are possibly active. This number does not indicate the total number of sessions currently active on the system.

Service Level Objectives for the Service

WebQoS offers the following SLOs at the service level:

- `Support at most NUMBER WebQoS sessions`

  This SLO lets you support a specific `NUMBER` (for example, 1000) of concurrent sessions. The request classification `CLASS` information is not supported.

  Concurrent sessions means the total number of sessions that have not expired and are possibly active.

- `Ensure less than NUMBER millisecond response time for CLASS requests to URL GROUP`

  For requests to the specified URL(s), make sure the response time is met for the specified request classification. Refer to the online help for information about using wildcards in the URL(s).

- `Ensure MEASUREMENT is CONDITION NUMBER`

  Make sure the specified measurement meets the specified condition.

Threshold Policy

The threshold policy is an “internal” measure used to ensure that your system is operating in a reasonable performance range and does not get overloaded. It puts limits on CPU load.

If CPU load gets too high, your server may get into a thrashing state causing few or possibly no requests to be satisfied in a timely fashion. If WebQoS queues begin to grow abnormally, this may signal a problem with your web server or an application/database process that it relies upon.

The threshold policy supported by WebQoS is:

- `Ensure CPU is no more than PERCENT % busy`

  This threshold policy lets you specify the upper threshold of how busy the CPU is. The `CLASS` request classification information is not required.

  You can define five threshold policies for up to five CPU thresholds.

Corrective Actions for SLO and Threshold Policy Violations

Corrective actions are those actions taken to bring an SLO or threshold policy into compliance. These actions are only performed on new sessions. Existing sessions that are already admitted into the system are not affected. The corrective actions supported by WebQoS are:

- `Redirect sessions up to NUMBER times for CLASS priority request`

  This corrective action limits the number of times a session can be redirected. You enter the `NUMBER` of times for the respective priority request. It is highly recommended that the specific URL all have mirrored sites. You can redirect the session to the URL of a website (for example, `http://www.bigcompany.com`) which would redirect sessions to this specified location.

  If the `NUMBER` of redirection is met, the next corrective action for this class is executed. If no other corrective action for this class is found, the session is admitted into the system. Therefore, it is highly recommended that every redirection policy has a rejection action of the same class following it in the corrective action list.

- `Reject CLASS priority requests`

  This corrective action rejects sessions for the respective priority request.

  When a session is rejected, WebQoS sends the rejected session a web page indicating that the server is not available.

  If both redirection and rejection actions are configured for any class, the rejection action should be the last in the corrective action list. For example, for any class, the order of its corrective action class could be: “redirect, reject” but not “reject, redirect.”

  Refer to “Customizing the Defer and Reject Web Pages” on page 125 for information about how to customize the Reject web page.

- `Defer CLASS priority requests`

  This corrective action defers sessions for the priority request. The deferred session is admitted instantly after the deferred time has passed if the `Automatically Enter` checkbox is selected.

  When a session is deferred, WebQoS sends the deferred session a customized web page indicating that the web site is very busy, and that the request will be fulfilled soon. The requestor sees a timer that indicates when the request will be re-submitted automatically.

  Refer to “Customizing the Defer and Reject Web Pages” below for information about how to customize the Defer web page.

Customizing the Defer and Reject Web Pages

The defer and reject web pages can be customized by editing the following files:

- `Defer.html`
- `Reject.html`

In UNIX, these files are located in the `/etc/opt/webqos/html` directory. In Windows, these files are located in the *<install-directory>*`\html` directory.

The customizable defer and reject web pages can be edited to contain foreign language contents. In Windows, these customizable web pages do not support Unicode.

You must restart the web server if you modify these files or path names.

To use the customizable reject web page, you must enable `Reject message is an HTML file name` and enter the exact path name(s) and filename in the `Reject Message` box (for example, `/etc/opt/webqos/html/Reject.html` or `C:\Program Files\Hewlett-Packard\HPWebQoS\html\Reject.html`).

You can also use your own reject web page. Be sure to enter its exact path name(s) and filename in the `Reject Message` box and restart the web server. If you enter the path name(s) or filename incorrectly, WebQoS uses the `Reject.html` file instead.

To use the customizable defer page, you must enable `Defer message is a URL`. The following are examples of customizable defer pages. One example shows how to incorporate the configured URL (from the Defer Corrective Action screen) into the defer page.

Using a Defer Timer If you are creating a Defer HTML page and want to use the defer timer mechanism, you need to place the following form in your page:

```
<FORM NAME="clock_form">
<P>
<INPUT TYPE=TEXT NAME="clock" SIZE=6
value="MINUTEKEY:SECONDKEY"> (minutes:seconds).
</b>
</P>
</FORM>
```

Figure A-1 shows the default defer page with a defer timer. The defer timer displays the time remaining before the request is retried.

Figure A-1 Default Defer Page

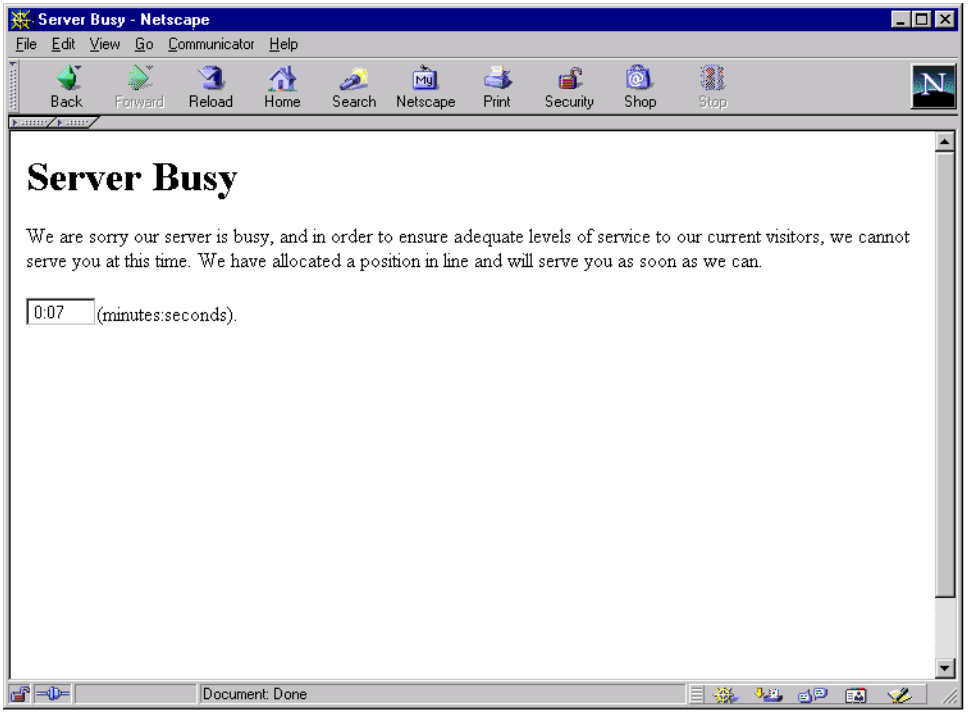

Using the Is URL Link Your defer page can also include a link to another page that the user can go to while waiting for the original request to be retried (see Figure A-2).

Figure A-2 Example of Custom Defer Page with Link

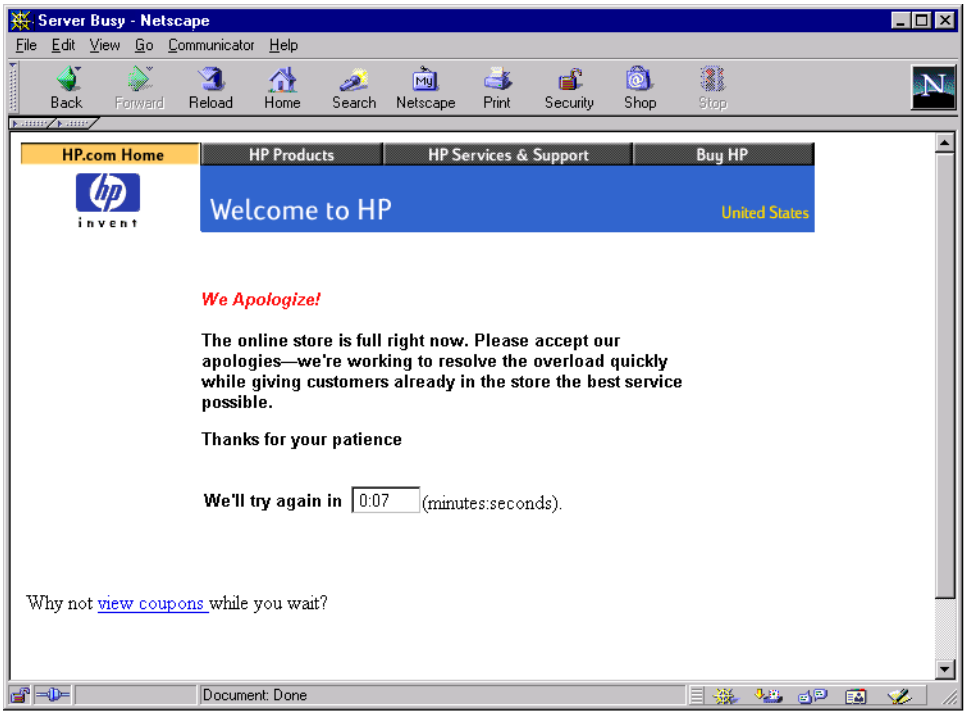

If you want to use the link you configured in the `Defer Corrective Actions` screen (by selecting the `Is URL` checkbox, entering a URL in the `Defer Message` box, and restarting the web server) within the defer page you are creating, you must code the link using Javascript. This retains the request's position in the defer session queue. The Javascript dialog is shown below:

```
<SCRIPT LANGUAGE="java-script">
<!-- Hide from non java-script browsers

function clockTick()
{
        now = (new Date()).getTime();
        remaining = Math.max(Math.round((finalTime - now) /
```

```
1000),0);
        minutes = Math.floor(remaining / 60);
        seconds = remaining % 60;

        if(seconds<10) secondString = "0" + seconds;
        else secondString = "" + seconds;

        document.clock_form.clock.value = " "+ minutes + ":"
+ secondString;
        document.clock_form.clock.blur();

        if(remaining > 0) setTimeout("clockTick()", 1000);
        else dialog();
}

finalTime = (new Date()).getTime() + (TOTALSECKEY * 1000);
clockTick();
// End of clock -->
</SCRIPT>

<P> Why not
<a href="DEFERURLKEY" target=_blank><i>
view coupons</i></a> while you wait?
</P>
```

Using a Popup Window Your defer page can also include a popup window that prompts the user to enter the site after the defer timer has expired (see Figure A-3).

Figure A-3 Example of Custom Defer Page with Popup Dialog

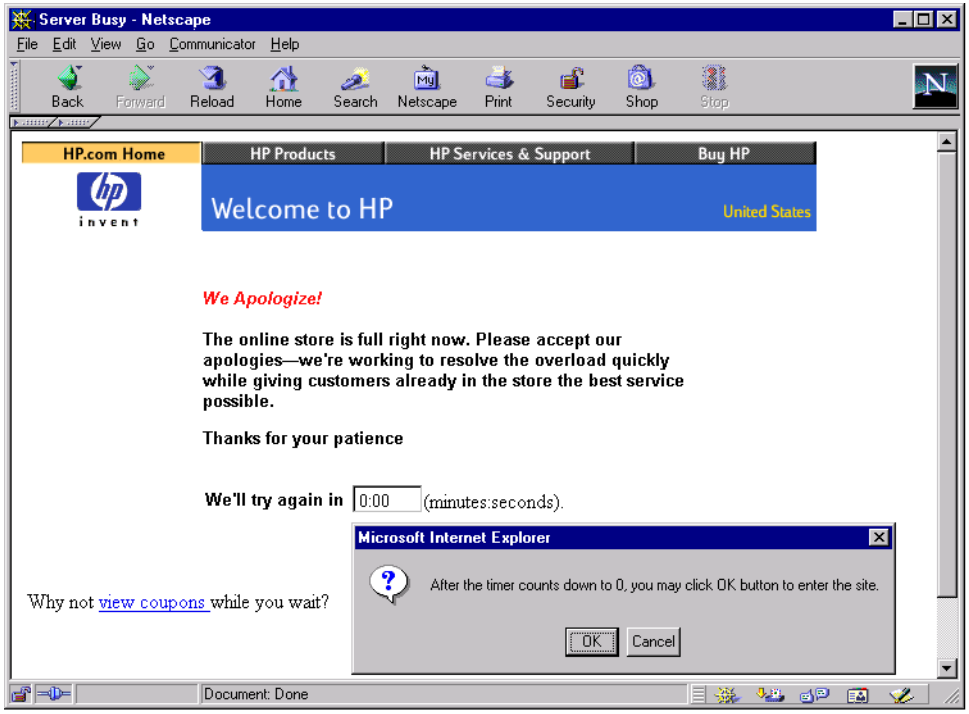

You must code the popup using Javascript. This is shown below:

```
<SCRIPT LANGUAGE="JavaScript">
<!-- Hide from non JavaScript browsers
function dialog()
{
    var msgs = "After the timer counts down to 0, you may
click OK button to enter the site.";
    if (window.confirm(msgs))
         parent.location="TARGETURLKEY";
}
```

```
function clockTick()
{
    now = (new Date()).getTime();
    remaining = Math.max(Math.round((finalTime - now) /
1000),0);
    minutes = Math.floor(remaining / 60);
    seconds = remaining % 60;

    if(seconds<10) secondString = "0" + seconds;
    else secondString = "" + seconds;

    document.clock_form.clock.value = " "+ minutes + ":" +
secondString;
    document.clock_form.clock.blur();

    if(remaining > 0) setTimeout("clockTick()", 1000);
    else dialog();
}

    finalTime = (new Date()).getTime() + (TOTALSECKEY *
1000);
clockTick();
// End of clock -->
</SCRIPT>

<P> Why not
<a href="DEFERURLKEY" target=_blank><i>
view coupons</i></a> while you wait?
</P>
```

B External Measurement API

This appendix describes how to use the external measurement API.

Using the External Measurement API

In Windows and Solaris, you must install the WebQoS API component.

The API allows you to send measurements to a customizable service level SLO. The SLO is: `Ensure <Measurement> is <CONDITION> <NUMBER>`. It is configured in the service template using the management console (refer to the online help for more information about this SLO).

To use the API, do the following:

1. Choose a unique measurement name. The name you use in your program must match the name you use in the service level SLO. If you are configuring more than one SLO based on different measurements, make sure you choose unique measurement names.
2. Write your program. Note the following:
   - Include the following header file in your program:

     In UNIX, the header file is
     `/opt/webqos/include/qosExternMeasApi.h`

     In Windows, the header file is
     `<install-directory>\include\qosExternMeasApi.h`
   - Call the following three methods in this order (`wqSendMeasurement()` can be called multiple times within a single `wqCreateConnection()` and `wqCloseConnection()` pairing):

     a. `wqCreateConnection()`

     b. `wqSendMeasurement()`

     c. `wqCloseConnection()`
   - Link to the following library when compiling your program:

     In UNIX, link to the file
     `/opt/webqos/lib/libqosExternMeasApi.sl`

     In Windows, link to the file
     `<install-directory>\lib\hpqosExternMeasApi.lib`
3. Define the service level SLO in the service template using the management console.

4. If you want to create a log file, on the system running the SCO component, enable the *SCOLogExternMeas* parameter in the qos.conf file by setting it to 1. Restart the SCO.
5. Run your program.

wqCreateConnection()

Establishes a connection to the SCI component of WebQoS.

Synopsis

```
#include </opt/webqos/include/qosExternMeaseAPI.h> or
#include <install-directory>\include\qosExternMeasApi.h

wqStatus wqCreateConnection(char *hostname, wqSocket *socket)
```

Arguments

hostname	The name of the system running the WebQoS SCI component.
socket	The socket descriptor returned by WebQoS and used by the two other methods to identify the connection to WebQoS.

Return Value

`wqCreateConnection()` returns a non-negative value representing the method status.

Return Value	Status
WQ_STATUS_OK	A connection with WebQoS has been established.
WQ_STATUS_FAILURE	Could not establish a connection with WebQoS.

wqSendMeasurement()

Sends the timestamp, measurement name, and measurement value to WebQoS.

Synopsis

```
#include </opt/webqos/include/qosExternMeaseAPI.h> or
#include <install-directory>\include\qosExternMeasApi.h

wqStatus wqSendMeasurement(wqSocket socket, char *xmlString)
```

Arguments

socket	The socket descriptor returned by `wqCreateConnection()` used to identify the connection to WebQoS.
xmlString	Timestamp, measurement name(s), and measurement value(s) in the following format:

```
<timestamp>TIMESTAMP</timestamp><measPair><measName>MEASUREMENT_NAME</measName>
<measIntegerValue>MEASUREMENT_VALUE</measIntegerValue></measPair>
```

where *TIMESTAMP* is the unformatted timestamp (in milliseconds), *MEASUREMENT_NAME* must match the measurement name you configured or will configure in the service level SLO, and *MEASUREMENT_VALUE* is the value of the measurement. More than one `measPair` can be included in the *xmlString*.

xmlString cannot be greater than 1024 characters (as defined by WQ_MAX_XML_STRING_SIZE in the header file).

Return Value

`wqSendMeasurement()` returns a non-negative value representing the method status.

Return Values	Status
WQ_STATUS_OK	The measurement has been sent to WebQoS.
WQ_STATUS_PACK_ERR	Packing error.
WQ_STATUS_UNPACK_ERR	Unpacking error.
WQ_STATUS_READ_ERR	Error reading from the socket.
WQ_STATUS_DIR_ERR	Unrecognized data read.
WQ_STATUS_MEM_ERR	Memory allocation error.
WQ_STATUS_PART_READ_ERR	Partial read error.
WQ_STATUS_WRITE_ERR	Error writing to the socket.

wqCloseConnection()

Closes the connection to the SCI component of WebQoS.

Synopsis

```
#include </opt/webqos/include/qosExternMeaseAPI.h> or
#include <install-directory>\include\qosExternMeasApi.h

wqStatus wqCloseConnection(wqSocket socket)
```

Arguments

socket	The socket descriptor returned by `wqCreateConnection()` used to close the connection to WebQoS.

Return Value

`wqCloseConnection()` returns a non-negative value representing the method status.

Return Value	Status
WQ_STATUS_OK	The connection with WebQoS has been closed.

Sample Program

This sample UNIX program is an example of how the three methods can be used. It prompts for the name of the system to connect to and the measurement name and value to supply to the service level SLO.

To successfully run this program, you must first configure the service level SLO. The measurement name you type in response to the measurement name prompt must match the name you configured for the SLO.

```
#include <stdio.h>
#include <time.h>
#include <math.h>
#include <stdlib.h>
#include </opt/webqos/include/qosExternMeasApi.h>
#define MAX_TOKEN 64

/*****************************************************************************
******************************************************************************
* SOURCE FILE : Test.c
*
* PURPOSE : A sample test program that invokes the methods in the external
*           measurement API and changes the statistic values. It makes the
*           following three calls to the API.
*
*          1. wqStatus wqCreateConnection(char *hostname, wqSocket *socket):This
*           call sets up the connection with the SCI. This method returns an
*           integer indicating success or failure.
*
*          2. wqStatus wqSendMeasurement(wqSocket socket, char *xmlString): This
*           call sends the xml string to the SCI, again returning an integer to
*           success or failure.
*
*          3. wqStatus wqCloseConnection(wqSocket socket): This call closes the
*           connection between the SCI and the API.
*
******************************************************************************
*****************************************************************************/

void main()
{
   char xmlString[WQ_MAX_XML_STRING_SIZE];
   char measName[MAX_TOKEN];                // Measurement Name
```

```
   char measValue[MAX_TOKEN];                // Measurement Value as a char array
    char hostname[MAX_TOKEN];                 // SCI System Name
    char str_timestamp[MAX_TOKEN];            // timestamp as a char array
    wqStatus status = 0;                      // status code
    wqSocket socket;                          //socket Descriptor
    char flag = 'y';
    double timestamp;
    long l_timestamp;

    // get the SCI system name
    printf("Enter the host name:\n");
    gets(hostname);

    // connect to the SCI
    status = wqCreateConnection(hostname, &socket);

    if (status == WQ_STATUS_OK)
    {
    // calculate the timestamp, once the connection is set up
    timestamp = (double)time(NULL); /* convert seconds to milliseconds */
    l_timestamp = ceil(timestamp);
    // convert the timestamp into char*
    ltoa(l_timestamp, str_timestamp,10);

    // see the Schema file or the DTD to create an xml string.
    // create the xml string with the appropriate values and tags.
    strcpy(xmlString,"<timeStamp>");
    strcat(xmlString, str_timestamp);
    strcat(xmlString, "</timeStamp>");
    while(flag == 'y')
    {
     fflush(stdin);
     strcat(xmlString, "<measPair><measName>");
printf("Enter the Measurement Name :\n");
gets(measName);
printf("Enter the Value:\n");
gets(measValue);
strcat(xmlString,measName);
strcat(xmlString,"</measName><measIntegerValue>");
     strcat(xmlString,measValue);
strcat(xmlString,"</measIntegerValue></measPair>");
printf("Do you want to Continue :(y-yes)\n");
scanf("%c", &flag);
    }
    // send the xml string to the SCI
    status = wqSendMeasurement(socket, xmlString);
```

```
   }
   // close the connection once everything is done
   if (status == WQ_STATUS_OK)
      status = wqCloseConnection(socket);

}
```

Glossary

A

admission control - A type of access control that is focused on protecting the system from overload. Based on one or more system load thresholds, it makes decisions about what type of requests to admit and which ones to defer, redirect or reject.

C

class - *See request classification rule.*

co-hosting - The sharing of a server/cluster by several sites whose contents are typically owned and supplied by different companies or business units. Service isolation features within the system and the web server prevent different sites from interfering with each other.

co-location - The sharing of a physical premise by servers/clusters each hosting a different site, typically owned by different companies or business units. The businesses or business units save money by having one organization (Information Technology or a service provider) manage and maintain the systems. The machines, however, are considered to be owned by the business or business unit.

concurrent sessions - The total number of sessions that have not expired (based on the configuration of the advanced features under “Session Timeouts”) and are possibly active. This number does not indicate the total number of sessions currently active on the system.

corrective action - A set of one or more actions that are invoked when an SLO is violated. The set of potential actions are intended to alleviate the condition by causing a lower priority SLO to be violated, or by reducing load on the system via session deferral or rejection.

D

database - Stores configuration and event information, statistics, and component status. Resides on only one system in your WebQoS domain.

deferral - The postponement of the servicing of a user request. A deferral is performed by an admission controller, typically in response to heavy load on the server. The request is given a time slot at which it is expected that the server will have enough capacity to handle the entire session.

differentiated access - Access based on classification rules. Requests associated with one classification rule might be admitted while those associated with another are not.

differentiated performance - Differentiated response time and throughput achieved via queuing, process control, and flow control. Differentiation may be based on request classifications.

I

IP precedence - Refers to the precedence field in the ToS byte (see ToS byte) in IP. See RFC 1812. In WebQoS, the low-order 2 bits are used.

L

load balancer - A network element that directs traffic to servers within a server farm based on the current load being processed by the servers.

M

management console - Allows you to define performance objectives for the components of the service and to monitor performance levels and SLOs.

N

network QoS protocols - Protocols that are used for network quality of service. For example, IP ToS, RSVP, 802.1p, etc. Also known as QoS behavior.

P

persistent connection - An HTTP 1.1 feature in which a connection established with a web browser remains open for the servicing of additional subsequent requests coming from that browser.

priority - Request classifications determine how quickly a request is processed. Priorities are designated as high, medium, or low.

R

redirection - Redirection of a session to a specified URL that has WebQoS configured and running.

rejection - The decision not to satisfy a user request. A rejection is performed by an admission controller, typically in response to heavy load on the server. The user may come back at a later time when the server is less busy. The rejection may contain special incentives for the user to return later, such as a personalized promotion or an electronic coupon.

request classification rule - Determines the processing priority and access priority for requests submitted to a web site. During periods of heavy system load, low priority requests may be deferred or rejected. (This depends on your specific policies.) If a request is accepted, it is placed on a processing queue, based on its request class priority. Based on your specified policies, the request might be immediately processed, or may "sit" in the queue while other higher priority requests are processed first.

S

service - Something of value that a business is offering. A service is made up of a collection of service components (see service components), such as web sites, ftp sites, mail sites, customer databases, pricing rules, application logic modules, and so on

service component - A component that provides business specific data or processing rules. Examples include Internet service sites (for example, web sites, ftp sites), databases, business logic modules, business specific rules and scripts, and so on.

Service Control Agent (SCA) - Monitors site level SLOs, classifies user requests, collects performance information, and performs corrective actions. Resides on at least one system in your WebQoS domain.

Service Control Interface (SCI) - Receives configuration information from the management console and API and monitors the SCO and SCA(s). Resides on at least one system in your WebQoS domain.

Service Control Operator (SCO) - Monitors service level SLOs. Resides on only one system in your WebQoS domain.

Service Level Agreement (SLA) - A negotiated agreement between service providers (or information technology) and businesses or business divisions. SLAs contain one or more service level objectives (SLOs) that describe performance and capacity requirements for a given service hosted by the service provider. The SLA will often specify any penalties associated with violation of the objectives.

Service Level Objective (SLO) - A performance (response time), capacity, or availability objective for a service. Typically, there are several SLOs, each stating a very specific objective, possibly tied to a particular request classification rule. Action policies state corrective actions that may be taken when a SLO is in jeopardy.

service-oriented request classification - A request classification that focuses on the destination (specific module or content) of a request.

session - In WebQoS, a session is composed of one or more requests to a web site. These requests may arrive over one, or possibly several connections. Once a session has been granted, a user's remaining requests are guaranteed to be forwarded to the web server, unless the session times out. Under periods of heavy system load, it is possible that new sessions will not be granted. (This is driven by your specific policies.) However, existing sessions continue uninterrupted.

shared server/cluster - A system or cluster of systems that host services and content belonging to several businesses or business units.

T

threshold policy - An SLO defined by an IT/ISP administrator for the purposes of monitoring and maintaining the health and well being of the complex as well as keeping critical services up and running. This SLO is private to the service provider and are distinct from business SLOs which are explicitly negotiated with a hosting customer.

ToS byte - Type of Service byte in the IP header. In RFC1812, it is divided into three sections: the precedence field (high-order 3 bits), a field that is customarily called type of service or ToS (next 4 bits), and a reserved bit (the low order bit).

transparent classification - The classification performed by examining the attributes of a request such as whom it came from, where it is destined, or the type of service for which the request is constructed. The application is unaware that the classification has been performed by WebQoS.

U

URL document path - The directory path and filename of the URL used to access the site. For example, given the URL http://some.system.domain.com/my_path/index.html, the URL document path is /my_path/index.html