IBM WebSphere Performance Pack: Caching and Filtering with IBM Web Traffic Express

Marco Pistoia, Poh Yee Tiong

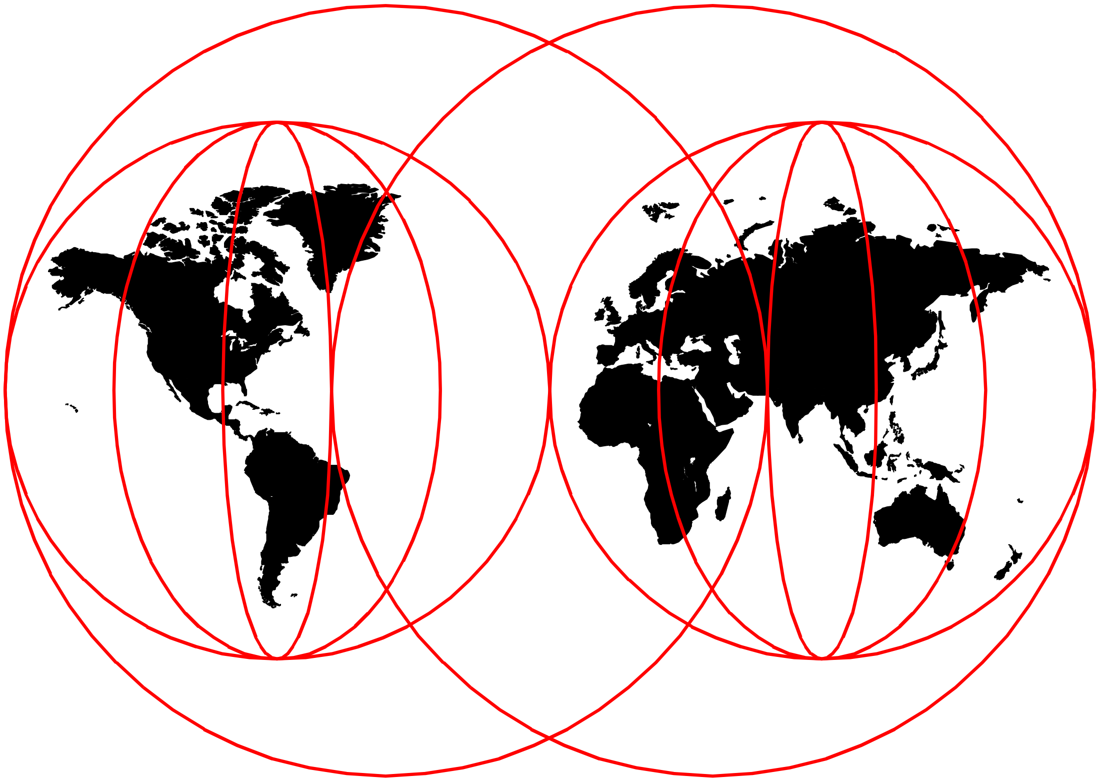

International Technical Support Organization

http://www.redbooks.ibm.com

SG24-5859-00

International Technical Support Organization

SG24-5859-00

IBM WebSphere Performance Pack: Caching and Filtering with IBM Web Traffic Express

October 1999

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, “Special Notices” on page 311.

First Edition (October 1999)

This edition applies to IBM Web Traffic Express, the Caching and Filtering component of WebSphere Performance Pack Version 2, for use with the AIX, Solaris, and Windows NT operating systems.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8A Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Preface

IBM WebSphere Performance Pack is Web infrastructure software that addresses the scalability, reliability and performance needs of e-business applications in both local and geographically distributed environments. Its functions incorporate leading-edge and robust caching, file management and load balancing, that together compensate for the inherent weakness of the Internet to support critical business applications and expectations.

This redbook will give you a clear understanding of the features of IBM Web Traffic Express, the Caching and Filtering component of IBM WebSphere Performance Pack. It shows how to plan for, install, configure, use, tune and troubleshoot each component and offers specific implementation examples. Moreover, it helps explain how to build complex scenarios that involve all the components of IBM WebSphere Performance Pack, to give you a better understanding of the technologies involved.

Note that this publication was written in conjunction with two other volumes about WebSphere: *IBM WebSphere Performance Pack: Web Content Management with IBM AFS Enterprise File System*, SG24-5857, and *IBM WebSphere Performance Pack: Load Balancing with IBM SecureWay Network Dispatcher*, SG24-5858. To realize the most benefit, all three volumes should be obtained.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

The leader of this project was Marco Pistoia.

Marco Pistoia is an Advisory Software Engineer, working as a project leader at the International Technical Support Organization Raleigh Center. He writes extensively and teaches IBM classes worldwide on all areas of the e-business Application Framework, WebSphere, Java and Internet security. Marco holds a Master of Science degree with honors in Pure Mathematics from the University of Rome. Before joining the ITSO, he was a System Engineer in IBM Italy. He received an Outstanding Technical Achievement Award in 1996.

Poh Yee Tiong is an Advisory IT Specialist in IBM Singapore. He joined IBM in 1990. He is responsible for RS/6000 AIX Services and Support for the ASEAN countries. He is an IBM-certified HACMP and RS/6000 SP specialist. He is also an Advanced Technical Expert for AIX.

Thanks to the following people for their invaluable contributions to this project:

Rick Schenck, Jerry Gschwind, Susan Hanis, Martin Presler-Marshall, Barbara Kemper
IBM Research Triangle Park, North Carolina

Catherine Milligan
IBM Transarc Laboratory, Pittsburgh

Corinne Letilley
IBM Canada

Shawn Walsh, Jorge Ferrari, Tim Kearby, Margaret Ticknor, Pat Donleycott, Tate Renner, Linda Robinson, Gail Christensen
International Technical Support Organization, Raleigh Center

Vincenzo Iovine, Stefano Pischedda
IBM SEMEA Sud, Italy

Karen Gelveles
IBM Boca Raton

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 327 to the fax number shown on the form.
- Use the online evaluation form found at `http://www.redbooks.ibm.com/`

  Send your comments in an Internet note to `redbook@us.ibm.com`

Part 1. Web Traffic Express Usage and Administration

Chapter 1. IBM WebSphere Performance Pack Concepts

IBM WebSphere Performance Pack is Web infrastructure software that addresses the scalability, reliability and performance needs of e-business applications in both local and geographically distributed environments. Its functions incorporate leading-edge and robust caching, file management and load balancing, that together compensate for the inherent weakness of the Internet to support critical business applications and expectations.

IBM WebSphere Performance Pack has been developed using IBM's extensive experience with very demanding Web sites.

IBM WebSphere Performance Pack is composed of three main components, which permit you to reduce Web server congestion, increase content availability and improve Web server performance:

1. **File sharing**

   The file-sharing component, known as IBM AFS Enterprise File System (AFS), is an enterprise file system that enables cooperating hosts (clients and servers) to efficiently share file system resources across both local area networks (LANs) and wide area networks (WANs). It provides non-disruptive real-time replication of information across multiple servers, which guarantees data consistency, availability, global stability and administrative efficiency, necessary by large distributed Web sites or by Web sites with volatile content requiring considerable administrative effort to maintain content links and URLs to file I/O mapping.

2. **Caching and filtering**

   The caching and filtering component, known as IBM Web Traffic Express (WTE), is a caching proxy server that provides highly scalable caching and filtering functions associated with receiving requests and serving URLs. With tunable caching capable of supporting high cache hit rates, this component can reduce bandwidth costs and provide more consistent rapid customer response times.

3. **Load balancing**

   The load balancing component, known as IBM SecureWay Network Dispatcher (ND), is a server that is able to dynamically monitor and balance TCP servers and applications in real time. The main advantage of the load balancing component is that it allows heavily accessed Web sites to increase capacity, since multiple TCP servers can be dynamically linked in a single entity that appears in the network as a single logical server.

WebSphere Performance Pack V2 offers a fourth component named Common Configuration. This new feature allows for centralized configuration of the main components and can be run remotely on a separate machine connected to the network.

The WebSphere Performance Pack components, which were previously unavailable in a single Internet software offering, can increase the scalability, availability and reliability of your Web site while reducing infrastructure costs.

Installation procedures permit selection of which components to install, and specification of on which machine(s) the selected component(s) should be located. Subject to installation needs and operating platforms, components can coexist on a single machine or can be distributed over multiple machines.

IBM WebSphere Performance Pack Version 2 is supported on the following platforms:

- Any IBM RISC/6000-based machine running IBM AIX 4.2.1 or later and Java Runtime Environment (JRE)[1] 1.1.6 or later
- Any SPARC workstation running Sun Solaris 2.6 or later and JRE 1.1.6 or later – JRE 1.1.7B or later is required to use ND
- Any Intel x86 PC running Microsoft Windows NT[2] 4.0 and JRE 1.1.6 or later – JRE 1.1.7B is required to use ND

If you are planning to use the Common Configuration, then you need the Java Development Kit (JDK)[3] 1.1.6 or higher and a Java 1.1-enabled Web browser, such as:

- Netscape Communicator V4.08 or higher
- Microsoft Internet Explorer V4.01 with the fix pack, or higher
- Sun HotJava V1.1 or higher

1.1 AFS Enterprise File System in Distributed Computing Environments

AFS has provided scalable file administration and file sharing for large enterprises for many years, based upon its use of a virtual name space to make naming and logical directory structures of files independent of their physical location. AFS clients and AFS servers are used to establish this virtual name space capability. In typical local area network (LAN) file systems, this is achieved by installing AFS clients on user workstations, communicating with an AFS server that manages the I/O operations associated with the actual files. In a Web site, the AFS clients can be installed on HTTP servers to reduce the administrative effort associated with maintaining URL to file I/O mapping relationships. In addition, HTTP servers that are simultaneously AFS clients can significantly increase the connectivity capacity to Web server content and can provide local and geographically distributed access efficiency.

AFS is a central and scalable file system:

- It is *central* because AFS brings together all of the files within the file system into a single name space. Every AFS user shares this same name space, making all AFS files easily available from any AFS machines. With AFS, the name of a file is independent of both the file’s and the user’s physical location, contributing to ease of file sharing and resource management.
- It is *scalable* because AFS is able to manage a very large number of files, spread across many geographical locations. When remote files, residing on AFS servers, are accessed by remote AFS clients, they are cached on the client machines to improve performance. This makes remote working across global distances feasible, since it is possible to access your own files from sites many thousands of miles away as if they were local.

[1] You can download the JRE from `http://www.javasoft.com`.

[2] IBM WebSphere Performance Pack can be installed on either Windows NT Server or Workstation V4.0. However, Windows NT Server is required for PCs running the ND functions Interactive Session Support (ISS) Nameserver and Observers.

[3] You can download the JDK from `http://www.javasoft.com`.

Both small and large-scale distributed environments benefit from AFS mechanisms to reduce server and network load:

- AFS caches data on client machines to reduce subsequent data requests directed at file servers, substantially reducing network and server loads. Servers keep track of data given to clients through *callbacks*, guaranteeing cache consistency without constant queries to the server to see if the file has changed. It is important to underline that AFS also allows disk cache, not just memory cache. This is a key advantage of AFS over other shared file systems.
- The AFS remote procedure call (RPC) reads and writes data to an RPC stream, further improving the efficiency of data transfer across a local or wide area network.

Extended security is guaranteed through Kerberos authentication and access control lists (ACLs). AFS Kerberos-based authentication requires that users prove their identities before accessing network services. Once authenticated, AFS access control lists give individual users or groups of users varying levels of permission to perform operations on the files in a directory.

The AFS component also offers replication techniques for file system reliability. Multiple copies of frequently accessed (but infrequently changed) data are replicated on multiple file servers within a cell. When accessing this information, a client will choose among the available servers that house replicas. If one server is unavailable or unreachable, the client will go to another server. Replication also reduces the load on any particular server by placing frequently accessed information on multiple servers.

Moreover, management utilities are provided to ease the load of system administrators in growing environments. Backup, reconfiguration and routine maintenance are all done without any system down time. Files remain available to users during these operations. This is done by creating online clones of volumes.

AFS commands are RPC-based. Administrative commands can be issued by any authenticated administrator from any client workstation. System databases track data location information, authentication information and protection groups. These databases are replicated on multiple servers, and are dynamically updated as information changes. Server processes accomplish many tasks automatically, such as restarting servers, tracking file locations and updating file servers with new binaries and configuration files.

1.2 Caching and Filtering to Manage Internet Traffic and Bandwidth Demand

WTE, the caching and filtering component of IBM WebSphere Performance Pack, is both a caching proxy server and a content filter. The advanced caching of this component minimizes network bandwidth and ensures that end users spend less time when retrieving the same content multiple times.

This component acts as a gateway for multiple clients and performs basic Web server duties, such as receiving requests and serving URLs.

A traditional proxy server receives a request for a URL from a client and it forwards the request to the destination content server. WTE does something more; it can save or cache the Web documents it retrieves, and serve subsequent requests for those documents from its local cache. The client gets the requested information faster and network bandwidth is reduced.

This component of IBM WebSphere Performance Pack also offers other key features of advanced caching, such as:

- The ability to handle very large caches
- An option to automatically refresh the cache with the most frequently accessed pages
- The possibility to cache even those pages where the header information says to fetch them every time
- Configurable daily garbage collection, to improve server performance and ensure cache maintenance
- Remote Cache Access (RCA), a new function – available for the first time with IBM WebSphere Performance Pack V2 – that allows multiple WTE machines to share the same cache by using a distributed file system, such as AFS, therefore reducing redundancy of cached content

Moreover, WTE allows you to set content filtering at the proxy server level, rather than or in addition to the browser level, where content filtering could be easily compromised or over-ridden. This way, offensive contents will not be displayed on the client’s browser, depending on the parameters used in the configuration. Content filtering in WTE can use:

- Platform for Internet Content Selection (PICS) rules guiding use of rating labels - such as Recreational Software Advisory Council on the Internet (RSACi) criteria for inappropriate language, nudity or violence - placed in HTML or HTTP headers or third-party content rating label distributions
- Lists of URLs/sites for which access is to be blocked
- APIs for filtering applications

1.3 Load Balancing and Server Monitoring Capabilities

The Internet has grown so rapidly over the last few years, that you are probably looking for a way to handle your company's share of that traffic. If this growth is not properly handled, users get slow response or refused connections, creating an unsatisfactory user experience which may cause the user never to visit your site again. Internet sites can become unstable or even fail under critical load conditions. What is needed is a solution that balances the load effectively and protects the user from these bad experiences.

ND, the load balancing component of IBM WebSphere Performance Pack, has been developed to address these limitations and provide customers with advanced functions to meet their site’s scalability and availability needs. It consists of three functions: the Dispatcher, Interactive Session Support (ISS) and the new Content Based Routing (CBR) function, available for the first time with IBM WebSphere Performance Pack V2. These three functions can be deployed separately or together in various configurations to suit a wide variety of customer application requirements:

- You can use the Dispatcher function to balance the load on the server within a local area network or wide area network using a number of weights and measurements that are set dynamically.
- ISS is a DNS-based load monitoring component (daemon) that can be installed on each of your servers. This group of daemons is called an *ISS cell*. One of the members of the cell becomes a spokesman for the load monitoring service.

  You can use the ISS function to balance the load on servers within a local area network or wide area network using a domain name server round-robin approach or a more advanced user-specified approach. ISS periodically monitors the level of activity on a group of servers and detects which server is the least heavily loaded.

  ISS provides an *observer* interface to enable other applications to use the load monitoring service. Observers watch the cell and initiate actions based on the load. Application servers with the ISS load monitor daemon installed can pass periodic load reports to the Dispatcher using the Dispatcher observer. The results of these reports can be factored into the load-balancing performed by Dispatcher.
- You can use the CBR function (which must be installed and configured to work together with WTE) to load balance traffic based on the content of a client's URL request. CBR also offers a Cookie Affinity feature that allows requests from a particular HTTP client session to be load balanced to the same server for a specified time period. The client session will maintain affinity for a particular server without relying on the IP address of the client.

These three components offer a high availability feature:

- The Dispatcher high availability feature involves the use of a secondary machine that monitors the main, or primary, machine and stands by to take over the task of load balancing, should the primary machine fail at any time. This feature is available on all the platforms where IBM WebSphere Performance Pack is supported, without using High Availability Cluster Multi-Processing (HACMP).
- In ISS high availability, all the nodes in a site work together to eliminate any single point of failure.
- Multiple CBR machines can be load balanced in turn using ND, therefore granting CBR high availability.

The high availability feature provided by the Dispatcher function can be successfully used even in other configurations, for example to guarantee firewall high availability.

1.4 Building Record-Breaking Web Sites

IBM WebSphere Performance Pack allows you to design several architectures to enhance the performance of your Web site. Figure 1 on page 8 offers an idea of the multiple configurations that can be obtained combining the WebSphere Performance Pack components:

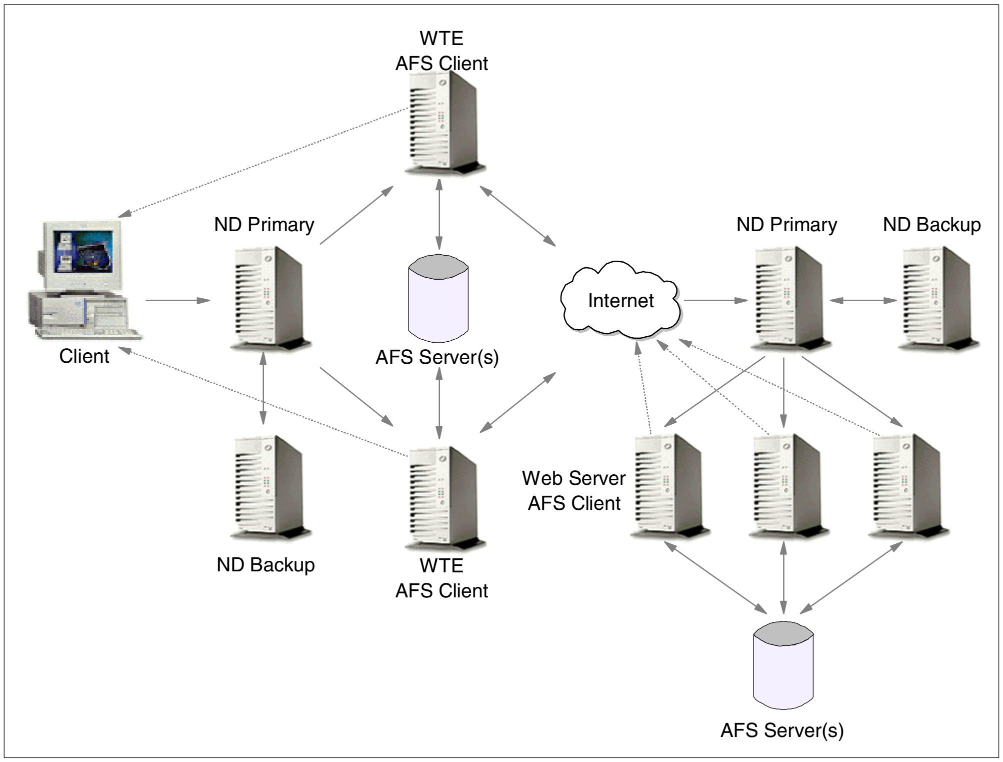

Figure 1. How to Implement a WebSphere Performance Pack Environment

- The WTE component minimizes response time and network bandwidth utilization by providing Web content caching. It also ensures reliable content filtering at the proxy server level. Multiple WTE servers can be load balanced.
- The ND component distributes the load between multiple clustered Web servers and WTE servers. As soon as a client request arrives, the load balancing machine uses sophisticated monitoring tools and then forwards the request to the least loaded server. High availability is provided by a backup Dispatcher machine, which monitors the state of the primary machine and takes over the primary machine if it fails.
- The clustered Web servers can share the same content by using AFS. This ensures scalability, high availability, reliable access to replicated data and an efficient security model for access and group management.

  AFS can also be integrated with RCA, so that clustered WTE servers can share the same Web cache therefore eliminating the need to fetch the same Web pages multiple times.
- The server selected by the ND machine can then respond directly to the client without any further involvement of the ND machine. Since there is no need for the server response to go back through the same physical path, a separate high-bandwidth connection can be used.

IBM used the WebSphere Performance Pack technology to create a scalable and reliable system that efficiently handled unprecedented traffic volumes. On February 17th, 1998, at 12:41 (Japan Standard Time), the official Web site of the Olympic Winter Games in Nagano made Internet history by logging a staggering 98,226 hits per minute. Less than a week later, a new all-time record was established with a peak load of more than 103,400 hits per minute, while still providing normal response time. The Internet site of the 1998 Nagano Olympic Winter Games is recognized by the Guinness Book of World Records.

In addition to the Winter Olympics site, IBM has built some of the largest Web sites in the world. For example, IBM hosted the Deep Blue chess match, the 1996 Olympic Games, the U.S. Open tennis tournament, Wimbledon, the Masters golf tournament, and the official Web site of the 1998 French Open tennis championship.

By using WebSphere Performance Pack on your Web site, you will have an efficient Web site, capable of providing fast responses. Your Web site will be able to handle very large amounts of simultaneous requests, without any major delay. Furthermore, the high availability features built into WebSphere Performance Pack will make the Web services available even when one or more of your server machines should unexpectedly fail.

The following figure shows an example of what your Web customers would not see if your Web site uses WebSphere Performance Pack:

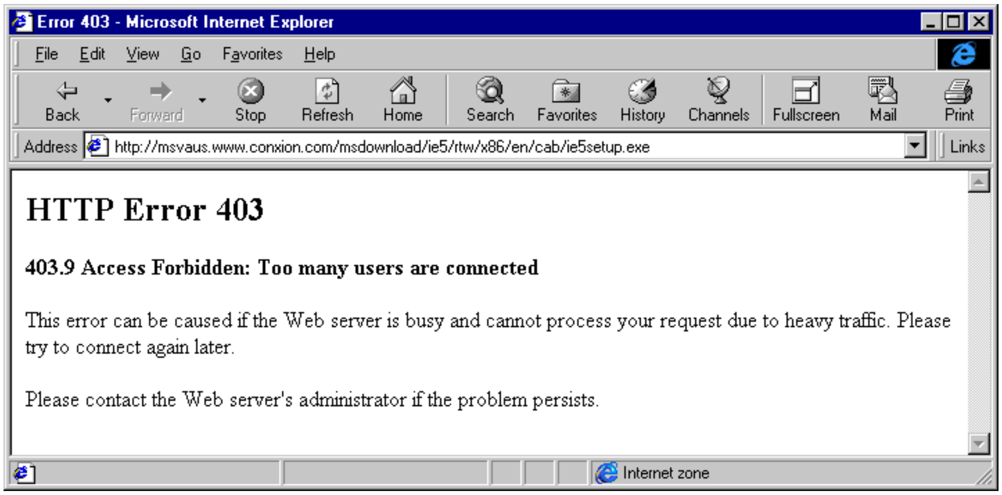

http://msvaus.www.conxion.com/msdownload/ie5/rtw/x86/en/cab/ie5setup.exe

HTTP Error 403

403.9 Access Forbidden: Too many users are connected

This error can be caused if the Web server is busy and cannot process your request due to heavy traffic. Please try to connect again later.

Please contact the Web server's administrator if the problem persists.

Figure 2. Don't Let This Happen to You!

1.5 What Is New in Version 2

In addition to the functions already available in Version 1.0, described in the IBM redbook *IBM WebSphere Performance Pack Usage and Administration*, SG24-5233, IBM WebSphere Performance Pack enhancements in Version 2 include:

- Quality of service enhancements
- New ND functions
- New WTE functions
- New AFS functions
- New functions available by combining components
- Integrated configuration assistance

1.5.1 Quality of Service Enhancements

WebSphere Performance Pack now provides support for differentiated quality of service, as discussed in the following sections.

1.5.1.1 User Classes

User classes give you the ability to change the level of service depending on the identity of the user. Preferred customers would expect to get better service. In Version 1.0, ND provided this capability by allowing rules-based load-balancing based on the client IP address, and WTE allowed PICS filtering based on the client IP address or the user and group. In Version 2, ND and WTE add the ability for rules-based load-balancing based on HTTP headers including `Cookie`, `Referer`, and `User-Agent`.

Differentiated service based on user identity lets you determine what level of service should be provided. For example, a frequent buyer can be routed to a higher capacity server or given access to additional content (for example, sale information) not made available to an unknown customer. The information used to identify the user comes from cookies or header information in the HTTP request.

1.5.1.2 Service Classes

Server classes give you the ability to change the level of service depending on the information or application requested by the client. Services involving a purchase would get better service than requests for information. In Version 1.0, WTE provided this capability through its reverse proxy function by allowing requests to be redirected to different servers based on the contents of the URL. In Version 2, ND and WTE add the ability for rules-based load-balancing based on the protocol, host, and path portion of the URL.

Differentiated service based on the service requested lets you give preferential service to some requests. For example, you may want to give preferential service to a *proceed to checkout* request over a *search* request.

1.5.2 New ND Functions

ND has been enhanced to include several new features. These are discussed in the following sections.

1.5.2.1 Server Directed Affinity API

In previous releases, administrators could define a particular port to be configured as *sticky*, meaning all requests to that port remain with a particular physical server for a short period of time. When a client connects to a sticky port, an entry for that client's IP address is made in the affinity table, and a time stamp is set. If a new connection from the same client arrives, and the time stamp has not expired, then the new connection is sent to the same server as before. This function is still available in Version 2, but it is no longer the only option.

Customers can now write their own software using the *Server Directed Affinity* (SDA) Application Programming Interface (API) to implement an SDA agent, which communicates with a listener in the Dispatcher. This software can then manipulate the Dispatcher affinity tables to:

- Query the contents
- Insert new records

- Remove records

1.5.2.2 Binary Logging and Statistics

Dispatcher now provides a log in binary format. The command line and the GUI have been enhanced to provide access to the binary log information. A sample Java program is provided to allow customers to access and manipulate the log.

1.5.2.3 Remote Administration

In the previous release, the Dispatcher administration GUI could only be run on the same machine where Dispatcher was installed. Now remote administration is possible, since the Dispatcher administration GUI can be run on a separate machine.

1.5.2.4 Authenticated Administration

Authentication is provided to make remote administration more secure. The communication between the ND server and an ND administration client is authenticated using a key pair.[4] These keys are generated when the ND server is started for the first time.

1.5.2.5 Wildcard Cluster

Wildcard cluster is the capability to define a cluster which will receive traffic which is not destined for this particular machine, or which is destined for this machine, but for an address which is not defined as a cluster. The traffic will be intercepted and routed to a default cluster. This feature also makes it easier to configure multiple aliases on the same Dispatcher machine to use the same port and server configuration.

1.5.2.6 Wildcard Port

Wildcard port is the capability to define default actions when no port matches a particular request. You can use this to create a load balancing configuration for traffic to any port that has not been explicitly defined in your Dispatcher configuration, for example in a firewall load balancing environment, or to discard requests for ports that have not been configured.

1.5.2.7 ISS GUI

The GUI for ISS is essentially an `isscontrol` command generator much like the Dispatcher GUI. In other words, the user makes a change in the GUI, and an isscontrol command is generated and issued. Asynchronously, the issd daemon tells the GUI that the configuration has changed and the GUI refreshes itself.

1.5.3 New WTE Functions

The enhancements to the WTE component of WebSphere Performance Pack are described in the following sections.

1.5.3.1 Transparent Proxy

Transparent proxying means that the client software is totally unaware of the existence of the intermediate proxy server. Normally, if a client browser uses a proxy server, then the browser must be configured to specify the address and port of the proxy server. This is no longer necessary with transparent proxy, in fact the client is unaware that an intermediate proxy is in the network.

[4] A *key pair* is a matching pair of public and private keys, used for digital signatures and asymmetric encryption.

To use transparent proxy, the router, which may be a Dispatcher machine, is programmed to redirect requests to the WTE transparent proxy. WTE then intercepts all HTTP requests on port 80 that are targeted at some server out in the Internet. The request is parsed and processed, and may be satisfied from the transparent proxy's cache.

Note that in WebSphere Performance Pack V2, transparent proxy is only supported on the AIX platform and works for HTTP requests only.

1.5.3.2 Proxy Autoconfiguration Support

WTE now supports *automatic proxy configuration*, a feature of Netscape Navigator 2.0 (and later) and Microsoft Internet Explorer V 4.0 (and later). This feature provides a form of transparency, in that clients do not have to configure their browser to point to a specific proxy or SOCKS server, but to an automatic configuration file, as shown in the following figure:

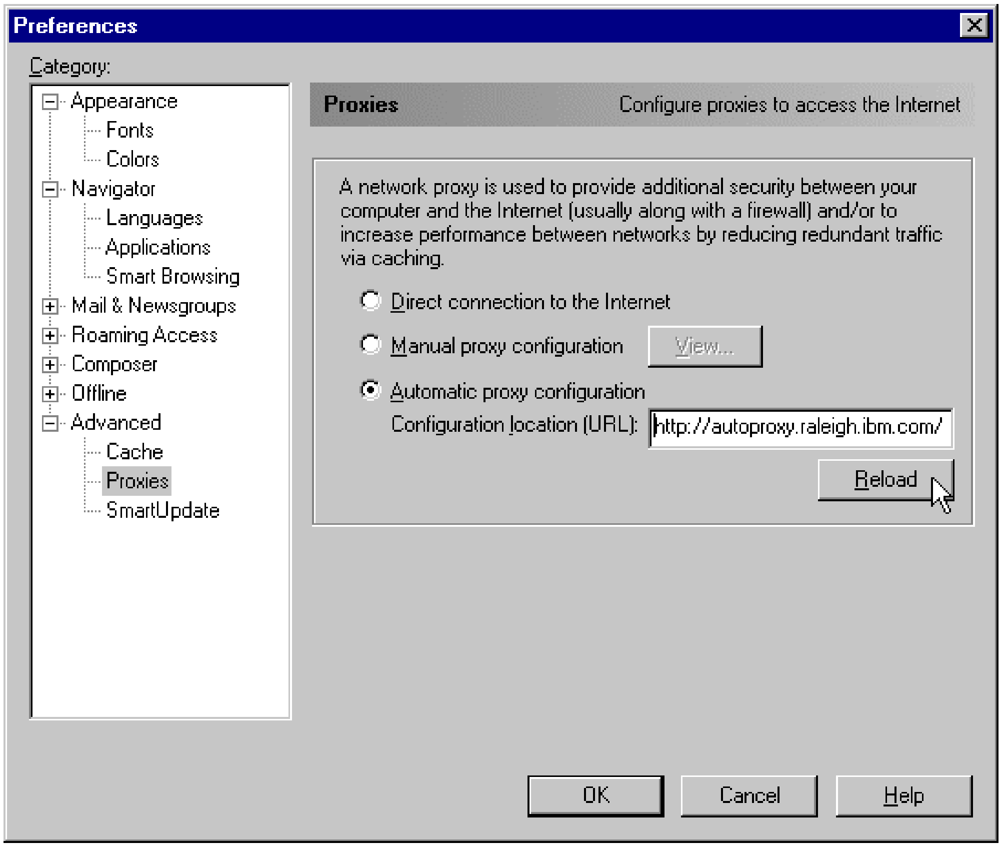

A network proxy is used to provide additional security between your computer and the Internet (usually along with a firewall) and/or to increase performance between networks by reducing redundant traffic via caching.

○ Direct connection to the Internet

○ Manual proxy configuration

● Automatic proxy configuration

Configuration location (URL): http://autoproxy.raleigh.ibm.com/

Figure 3. Automatic Proxy Configuration in Netscape Navigator

This lets the system administrator modify the configuration with little impact to the clients, who update their automatic configuration files and are automatically directed to the new configuration. Server administrators can use this to reroute requests when servers are down, to balance workload, to send requests for specific URLs to specific proxies, or other reasons specific to their installation.

1.5.3.3 FTP Proxy Enhancements

The WTE FTP proxy code now includes FTP PUT capability, improved authentication to prompt for the user ID and password (instead of requiring it in the URL), and a configuration directive to allow the user to specify whether FTP URLs will be treated as relative URLs or absolute URLs.

1.5.3.4 SNMP Subagent and MIB Support

WTE provides a Simple Network Management Protocol (SNMP) management information base (MIB) and SNMP subagent so you can use any SNMP-capable network management system, such as Tivoli NetView or Tivoli Distributed Monitoring to monitor your proxy server's health, throughput, and activity. The MIB data describes the proxy server being managed, reflects current and recent server status, and provides server statistics.

1.5.3.5 Performance Improvements

The cache architecture has been restructured to map URLs onto the cache file system more efficiently. This speeds retrieval of cached objects, uses disk space more efficiently within the cache, and speeds cache garbage collection. The cache also uses write-behind techniques for greater throughput.

Additionally, WTE now includes caching of Domain Name System (DNS) server lookup results, which can improve response time and reduce network load.

1.5.3.6 HTTP 1.1-Compliant Proxy Server

WTE is now an HTTP 1.1-compliant proxy server. WTE identifies itself as an HTTP 1.1 server and sends HTTP 1.1 in the outbound flows to origin servers. Persistent connections are supported from the client to the proxy, and from the proxy to the origin server. HTTP 1.1 cache control headers are processed and used to determine if a Web document is able to be cached. WTE receives and processes chunked data sent by HTTP 1.1 origin servers; WTE will unchunk data before giving control to a data filter or transmogrifier plug-in.

WTE provides new directives that allow the administrator to override certain HTTP 1.1 cache control headers. For example, query strings – URLs with a question mark (`?`) in them – are not considered cacheable by the HTTP 1.1 protocol, but WTE provides a directive that allows the administrator to specify which query strings should be cached.

WTE also provides an aggressive caching directive that allows the administrator to override the `Cache-Control: no-cache` header in Web documents.

1.5.3.7 Customization Exits

Several enhancements have been made to the WTE API to simplify writing an application plug-in. New request steps (*exit points*) have been added: *transmogrifier* and *garbage collection* (GC) advisor.

The transmogrifier gives the application write access to the outgoing data stream while the GC advisor allows the plug-in to influence garbage collection decisions:

- The transmogrifier step is intended to be used by applications that wish to perform transformations on the HTTP response data stream. Examples include converting Adobe PDF files to HTML, converting high resolution images to lower resolution quality, or translating pages from one language to another. The transmogrifier step is an extension of the data filter step.

  The WTE enhancements allow the application to specify multiple transmogrifiers, thus allowing the application to have multiple plug-ins, each of them performing different transformations on the data.

WTE introduces a correlator mechanism that eases state maintenance in the plug-in; WTE now automatically determines the content length of the response data, and the application no longer has to buffer the data to determine the content length. WTE also makes HTTP header processing easier; the response headers can now be extracted and set using API variables.

- The GC Advisor step is called for each file in the cache during the garbage collection process and allows the application to influence which files are kept and which are discarded.

1.5.3.8 Tivoli Ready

WTE is Tivoli Ready, which means it can be managed through either the Tivoli Enterprise Console (TEC), or through Tivoli Global Enterprise Manager (GEM). Supported Tivoli configurations include managing IBM applications that are installed on Tivoli-managed nodes, PC-managed nodes and endpoints in a distributed environment.

With the purchase of an IBM software product that carries the Tivoli Ready logo, you have the ability to manage your IBM software products through the Tivoli Enterprise management products, allowing you to automatically discover, monitor, and inventory one or more Tivoli Ready applications.

This Tivoli Ready instrumentation, when configured, provides you with the ability to:

- Graphically view the health of WTE through Tivoli GEM 2.2, TEC 3.1 or higher consoles
- Inventory WTE using Tivoli Inventory Version 3.2

1.5.3.9 Variant Caching

Variant caching extends the capabilities of the transmogrifier and allows WTE applications to request that WTE cache another version (*variant*) of the original document retrieved from the Web. This need arises when a plug-in performs a transformation on the original page retrieved from the Web. For example, if a plug-in translates a page from English to Italian, it would be advantageous to be able to cache not only the original document but the variant, as well. In fact, transformations of Web data are very CPU-intensive and degrade the performance of the proxy server.

By caching the variant, the number of transformations required is decreased, which improves overall proxy caching performance. The WTE API provides new predefined functions, HTTPD_variant_lookup() and HTTPD_variant_insert(), to allow the application to find a variant that has already been cached, and to insert a new variant into the cache.

1.5.3.10 Enhanced Log Maintenance

To effectively manage the space requirements for the four logs (error, access, proxy access, and cache access) generated by WTE, two improved functions have been implemented. They are compression and purging:

- Compression can be specified for the logs by date. The logs are compressed and stored with no regard for storage constraints.
- Purging is done by date and size of each file. The maximum size for each log is set during configuration of WTE. When the logs are purged by date, if the

maximum size for a log is still exceeded, daily logs are removed until the maximum size is no longer exceeded.

1.5.3.11 Error Message Personalization

The error messages sent to the client's browser are kept in HTML pages, which lets you change or alter the message. You can use this to give additional information, give additional instructions, or include things such as logos.

1.5.3.12 Secure Request Filtering

Filtering for Secure Sockets Layer (SSL) requests is done by using PICS filtering on the non-secure home page of the URL and applying the receive or block decision to the secure request.

1.5.3.13 Configuration Enhancements

Configuring the server using the browser, local or remote, is improved and more user friendly. The configuration screen is divided into three areas: navigation, workspace, and header. Navigation between different areas of the configuration process is always available in the navigation portion of the screen with a click of the mouse. The forms to be filled in are displayed in the workspace. The header portion contains a Help button and a Restart Server button. The server can be restarted anytime with the Restart Server button. When the Help button is clicked, another instance of the browser that does not interfere with the form being used displays the help information. The help information is divided into three types: field definition of the form being viewed, task oriented, and server general help.

1.5.4 New AFS Enterprise File System Function

The AFS component of WebSphere Performance Pack now provides some new features.

1.5.4.1 AFS Server for Windows NT

WebSphere Performance Pack Version 1.0 included AFS client and server for AIX and Solaris, but only the client was included for Windows NT. WebSphere Performance Pack V2 includes AFS Version 3.5 for Windows NT, which offers AFS server for Windows NT. With this component, you can store AFS files and directories, and run processes that provide servers on a Windows NT machine.

1.5.4.2 AFS Control Center

The AFS Control Center is a set of Windows NT-based tools for managing AFS cells. The AFS Control Center includes a User Administration graphical user interface (GUI) for account management and a Server Manager GUI for volume management. The Control Center helps simplify AFS server and account administration by letting administrators manage entire environments from a single Windows NT workstation, for example, monitoring file server utilization, transferring collections of files across servers, performing load balancing, or managing AFS accounts.

1.5.5 New Functions Available by Combining Components

As we said, WebSphere Performance Pack is composed of three main components: ND, WTE and AFS. Although these components can be installed separately, they are not different products; especially in this new version of WebSphere Performance Pack, new functions have been provided to integrate these components together and supply services that were not previously available.

1.5.5.1 Content Based Routing

Content Based Routing (CBR) is a new ND function that provides load balancing enhancements to the reverse proxy capabilities of WTE. CBR works in conjunction with WTE, which must be installed on the same machine.

The same rules used for Dispatcher can be used with CBR to load-balance requests over different sets of servers based on the client IP address, the entire URL, the protocol portion of the URL, the host portion of the URL, the path portion of the URL, the Referer HTTP header, or the User-Agent HTTP header.

1.5.5.2 Peak Load Management

Peak load management is the ability to detect and react to sudden increases in activity. In Version 1.0, ND allowed rules-based routing to alter the load-balancing algorithm based on the current connection rate, the number of active connections, and the time of day. In Version 2, an advisor for Dispatcher enhances load management on WTE by preventing Dispatcher from sending new requests to a WTE node that is engaged in garbage collection or cache refresh.

1.5.5.3 Remote Cache Access

RCA provides a way to share cache content among proxies. In many scenarios, multiple proxies are deployed *near* (in network terms) each other. Typically, they will have a load balancer in front of them for load balancing and high availability reasons. Each proxy has its own cache, and if the cached data cannot be shared, this results in cache space being wasted as multiple copies of the same document are stored. Also, since each cache is smaller than the sum of all the disks, the cache hit rate is lower (due to the smaller cache size).

RCA is a new, powerful feature implemented in WTE, and not available in the previous release of IBM WebSphere Performance Pack. RCA allows multiple proxy servers to cooperate to form cache arrays. Using RCA, multiple proxy servers can distribute the cache contents across their combined, logical cache to improve hit rates and reduce redundancy of cached content.

WTE uses new caching algorithms and information-sharing technologies to enable an Internet service provider (ISP) to manage its servers more efficiently by storing information where it is more likely to be needed and delivering it more efficiently to customers. These enhancements reduce transmission costs and eliminate the need for ISPs to replicate information in redundant proxy servers.

Although RCA is a new feature of WTE, it is best used in conjunction with the other components of WebSphere Performance Pack. In fact, RCA enables multiple peer WTE proxy servers, load balanced by one ND machine, to share the contents of their caches utilizing a shared file system, such as AFS, DFS, NFS, or Windows NT file sharing. We recommend the use of AFS, since it offers nondisruptive real-time replication of information across multiple servers, data consistency, availability, global stability and data consistency.

1.5.5.4 Integrated Configuration Assistance

To facilitate proxy caching, the install program automatically sets up caching with WTE. WebSphere Performance Pack also provides wizards to make configuration easier. Invoked from a browser, these wizards can configure key

WebSphere Performance Pack scenarios that involve all of the WebSphere Performance Pack components.

1.6 Who Can Benefit

IBM WebSphere Performance Pack allows you to design and use multiple architectures. The scenarios described in this section provide specific examples of how various ISP implementations can benefit from the use of IBM WebSphere Performance Pack.

1.6.1 Content Hosting Internet Service Providers

Content hosting ISPs can use WebSphere Performance Pack to more efficiently support and distribute the content from their own Web sites, and to provide more response- and cost-effective access to other sites.

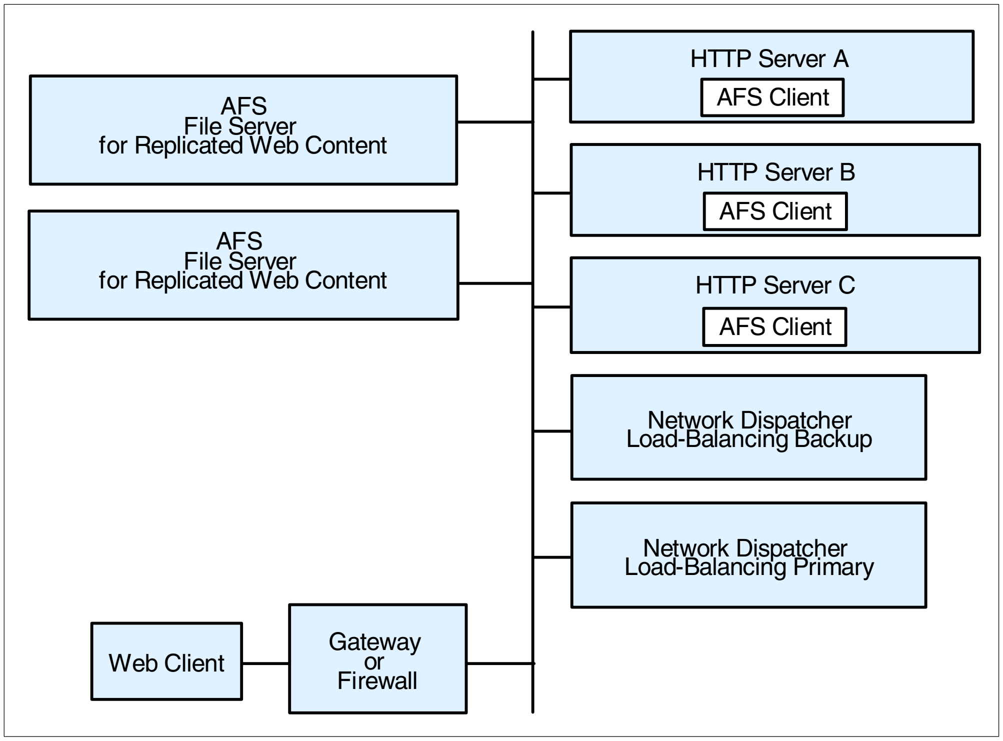

Figure 4. Content Hosting Internet Service Providers

Within a content hosting *server farm*, the WebSphere Performance Pack components can be configured to provide high availability and accessibility as follows:

- Content for hosted Web sites can be distributed over multiple volumes using the AFS virtual name space to simplify administration of page content. Because all of the Web servers need to have equal access to the Web content, a shared file system is an obvious choice for manageability of the Web content. AFS is a superior file system in this environment because of its replication capabilities, which provide improved availability and scalability. The disk caching capability of the AFS client ensures that network accesses to the file server are minimized.

- Scalability of access to the content can be achieved by adding multiple HTTP servers that are simultaneously AFS clients to access AFS server content, and by using efficient load balancing to dispatch requests to the HTTP server with the best capacity to handle the requests. ND, the load balancing component of IBM WebSphere Performance Pack guarantees improved availability and scalability by allowing a farm of Web servers to provide a single Web site image to clients.
- High availability can be maintained by using AFS file replication capabilities, and by configuring a hot standby ND component.
- Proxy caching can be used to provide more responsive access to content from other sites, as well as to optimize backbone network traffic capacity.
- Both availability and performance may be further enhanced by geographically distributing HTTP servers with AFS clients, together with proxy caching for other Internet content closer to user access points, such as points of presence (POPs) and network access points (NAPs).

Thus, content hosting service providers or corporate webmasters can benefit from the non-disruptive replication and distribution capabilities of the file sharing functions, local and wide area load balancing, and proxy caching. Flexible configuration of these components can ensure that requests are directed to the most appropriate local or remote location, and can enable location outages or routine maintenance schedules to be handled without disrupting customers.

The IBM WebSphere Performance Pack components can be used in conjunction with firewalls and authentication gateways to provide secure access where desired, and the load balancing function of WebSphere Performance Pack can also be used to scale these capabilities.

1.6.2 Corporate Web Sites and Content Aggregators

Use of IBM WebSphere Performance Pack by corporate Web sites and content aggregators is similar to that of content hosting ISPs.

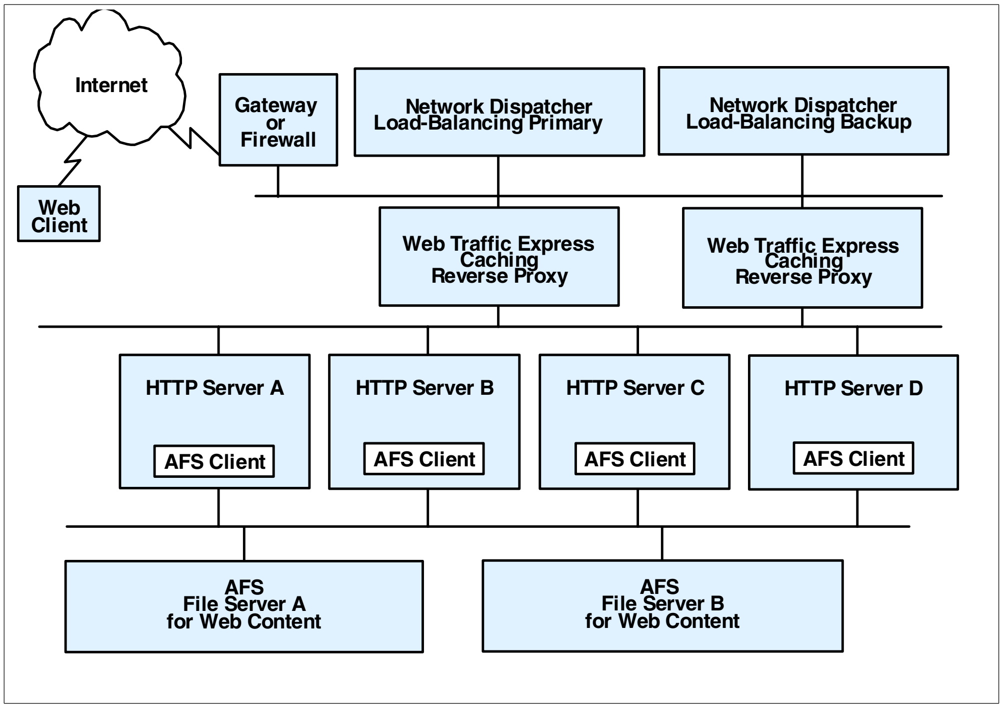
Internet
Web Client
Gateway or Firewall
Network Dispatcher Load-Balancing Primary
Network Dispatcher Load-Balancing Backup
Web Traffic Express Caching Reverse Proxy
Web Traffic Express Caching Reverse Proxy
HTTP Server A
AFS Client
HTTP Server B
AFS Client
HTTP Server C
AFS Client
HTTP Server D
AFS Client
AFS File Server A for Web Content
AFS File Server B for Web Content

Figure 5. Corporate Web Sites and Content Aggregators

In many cases, corporate Web sites and content aggregators need to maintain a demilitarized zone (DMZ) to ensure that access to Web content by employees, business partners, or customers does not expose internal computer resources to unauthorized users or hackers. However, firewalls cannot be deployed between AFS clients and their servers. In such instances, AFS replication can be used to establish read-only AFS servers within the DMZ. Here multiple AFS clients and the load balancing ND function can be used to provide the degree of scalability necessary to satisfy users. In addition, WTE caching and filtering proxy servers may be deployed on the same machines or on separate systems to filter or optimize access from within the corporation to external Web sites.

Many corporate Web sites are located at head office or regional locations, while branch offices or business partners may need frequent access to the content. Most offices of this type have relatively low line speed (somewhat less than 1.5 Mbps) network access to the regional or corporate sites. Relatively few users can, with concurrent usage, use all the available bandwidth with resulting erratic response times. At these locations HTTP servers with AFS clients combined with general purpose caching can provide more consistent user response time.

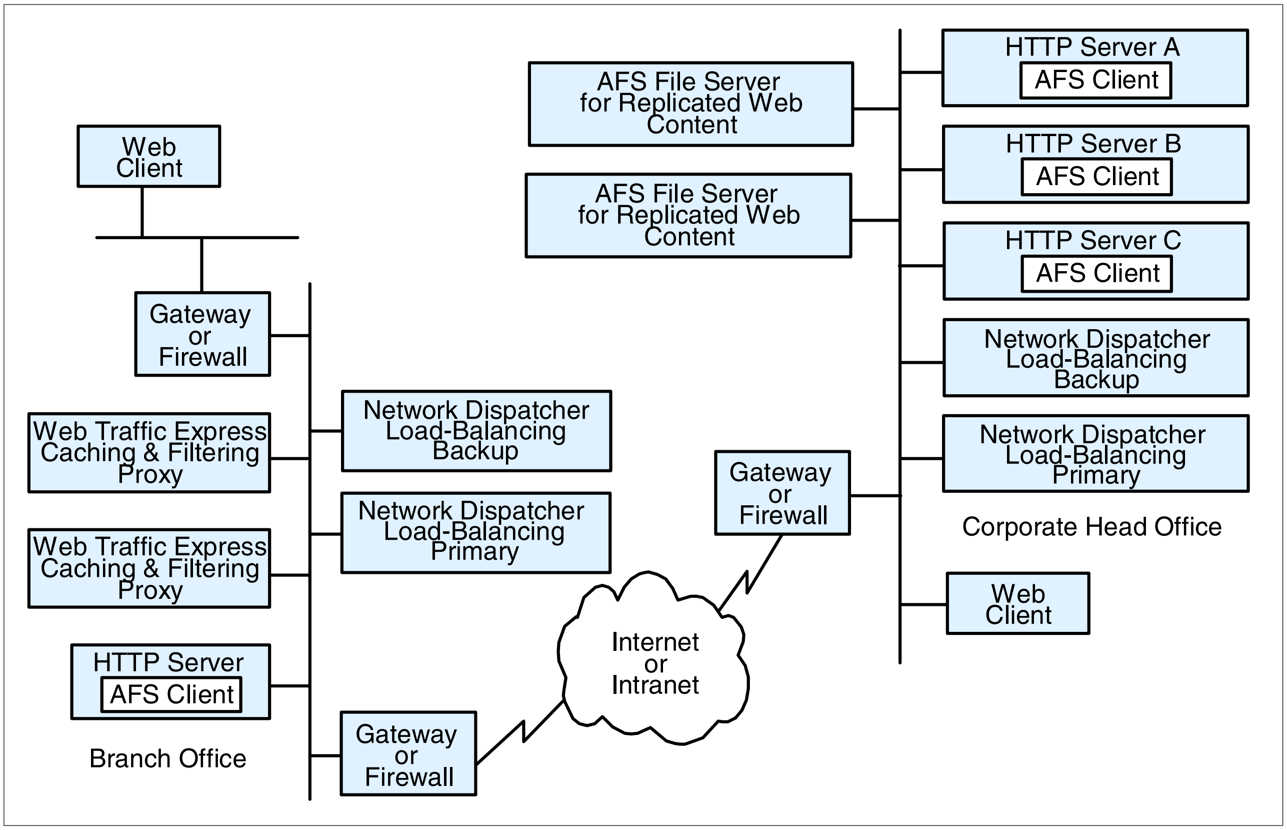
HTTP Server A
AFS Client
AFS File Server for Replicated Web Content
HTTP Server B
AFS Client
Web Client
AFS File Server for Replicated Web Content
HTTP Server C
AFS Client
Gateway or Firewall
Network Dispatcher Load-Balancing Backup
Web Traffic Express Caching & Filtering Proxy
Network Dispatcher Load-Balancing Backup
Network Dispatcher Load-Balancing Primary
Gateway or Firewall
Network Dispatcher Load-Balancing Primary
Corporate Head Office
Web Traffic Express Caching & Filtering Proxy
Web Client
Internet or Intranet
HTTP Server
AFS Client
Branch Office
Gateway or Firewall

Figure 6. Head Office - Branch Office

Some industries also have small branch offices in addition to larger branch or regional offices. A simple deployment of WTE caching and filtering on a single platform, eventually combined with a firewall, is probably sufficient to provide more consistent response time for employees in the small branch offices.[5] In the larger branches or regional offices, it may be desirable to have more than one WTE proxy server and to use the load balancing ND functionality to provide better resource management.

Two approaches can be used to reduce redundant page storage and to address the effectiveness of caching in these locations:

1. Hierarchical caching

   Optimize a primary (initial) cache for higher page hit rate by favoring more files of a smaller size, and a hierarchical cache for higher hit byte rate by favoring fewer files of a larger size.

2. RCA

   Configure RCA together with the load balancing function of ND and with AFS shared file storage to reduce the index size of each proxy and to increase the scalability of the caching storage.

[5] IBM has recently released a new product in the WebSphere family, called WebSphere Performance Pack WebSphere Cache Manager. This product is a caching and filtering proxy server obtained by the same code base as WTE. Cache Manager is available on Windows NT and Linux. Only a subset of the functions available in WTE are implemented in Cache Manager. For this reason, Cache Manager is particularly indicated for small companies or for small branch offices of large industries.

1.6.3 Corporate Headquarters Buildings or Large Campuses

On large campuses or in corporate headquarter buildings, the size of the campus or number of personnel frequently lead to the creation of smaller local area networks (LANs) interconnected by routers and a backbone LAN. Busy LAN servers combined with increasing use of Web server applications can result in congestion on the backbone LAN segments. This can be reduced by installing AFS client-enabled HTTP servers and general proxy caching on user LAN segments. The AFS clients can provide caching for corporate Web content, while the WTE proxy server can provide caching and filtering for external Internet access.

Design considerations for the smaller LAN segments are similar to those for small branch offices discussed earlier.

1.6.4 Backbone Internet Service Providers

Backbone ISPs typically provide co-location and/or peering services for other ISPs in addition to content hosting for large national or international corporations. In many cases they may provide *virtual ISP services* for other service providers such as content hosting ISPs. Backbone ISP customers are increasingly demanding both high availability and differentiated service levels, and backbone ISPs are responding by enhancing their Internet infrastructures. Features demanded by backbone ISPs include:

- Load balancing for a variety of traffic (for example, mail and FTP in addition to Web traffic)

  In addition, requests are made for traffic balancing management and high availability for authentication servers and management systems.

- Proxy caching

  To minimize the effects of *hot potato* routing and Web traffic *surges*, backbone ISPs are typically installing highly scalable caches at major peering points and network interconnections points such as network access points (NAPs).

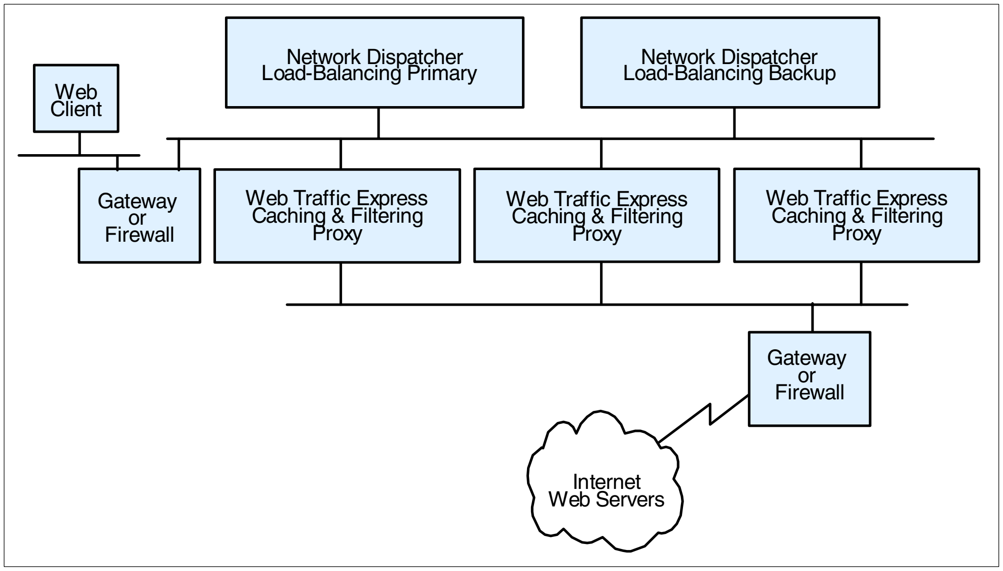
Network Dispatcher Load-Balancing Primary
Network Dispatcher Load-Balancing Backup
Web Client
Gateway or Firewall
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Gateway or Firewall
Internet Web Servers

Figure 7. Hierarchical Caching

International ISPs in particular need caching to reduce the costs associated with trans-oceanic links. Such installations can make very effective use of the RCA feature of WTE, a variation of the traditional caching function that when used together with load balancing and shared file storage can dramatically reduce redundancy of page storage.

Typical configurations for backbone ISPs would include load balancing for authentication gateways (such as authentication servers and subscriber management application servers) as well as for WTE caching proxy servers. Because of the amount of traffic and the desire for high availability, such caching servers would likely use the RCA feature and would thus also be configured as AFS clients. Therefore, there would also be an AFS server with the file content.

1.6.5 Access Internet Service Providers

Access ISPs need to provide both more consistent response time to their customers and to conserve their backbone network link and access charges. Thus caching at POPs would address these needs. These configurations would be similar to those for backbone ISP solutions.

Because access service providers target many of the small- to medium-size business customers, there is also an opportunity for them to create revenue producing services, by deploying smaller caching devices on customer premises but configuring and managing them as an ISP service. For this scenario, the caching would look much as it does for corporate branch offices.

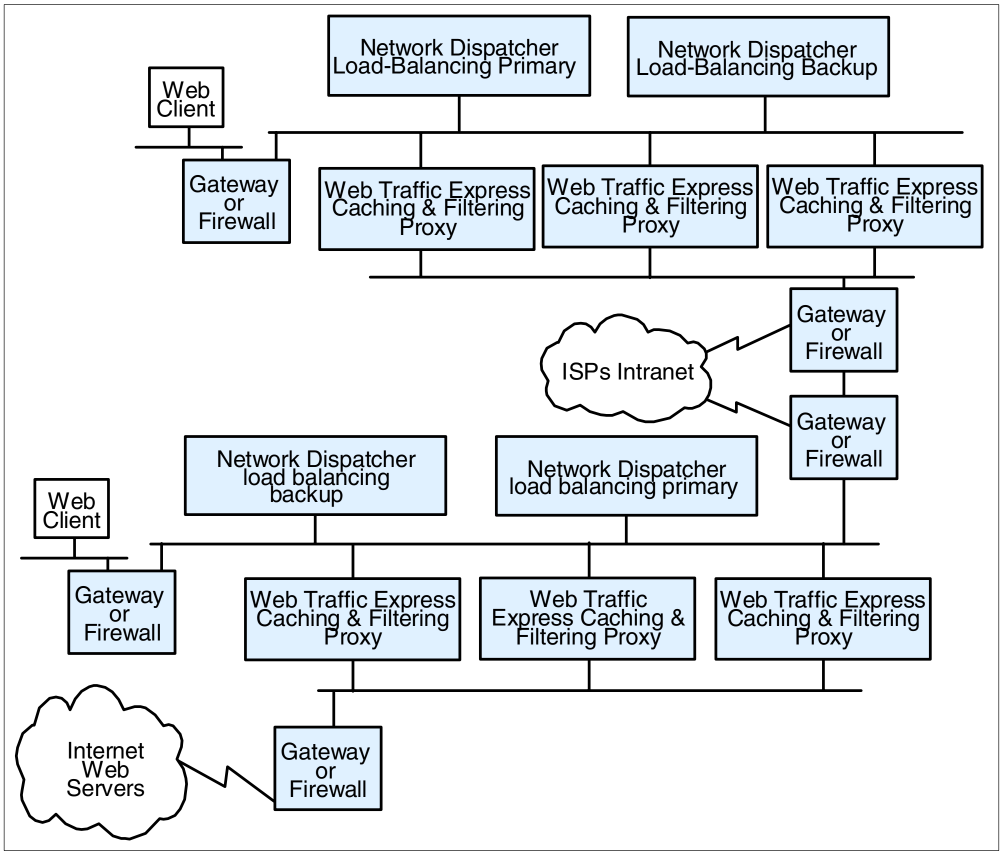
Network Dispatcher Load-Balancing Primary
Network Dispatcher Load-Balancing Backup
Web Client
Gateway or Firewall
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Gateway or Firewall
ISPs Intranet
Gateway or Firewall
Network Dispatcher load balancing backup
Network Dispatcher load balancing primary
Web Client
Gateway or Firewall
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Internet Web Servers
Gateway or Firewall

Figure 8. Access Internet Service Providers

1.6.6 Access ISPs with Subscriber Home Page Hosting

Access ISPs providing subscriber home page hosting typically do so on multi-homed servers, allowing subscribers to have their own domains. Such servers may have frequent changes, requiring the ISP to dedicate considerable time to administering the servers from a directory and backup perspective. For example, if the ISP is successful in recruiting new users, therefore having to expand the server capacity, it may involve reorganizing the various servers to spread the load and to allow room for individual sites to expand. The file management services provided by IBM WebSphere Performance Pack enable the ISP to reduce this maintenance workload, and to provide high availability replication for subscribers willing to pay for continuous availability. They also allow the ISP to scale up the site capacity with additional AFS clients, without the need to reorganize the underlying Web servers. ND load balancing can ensure that despite the growth users get rapid response time. In addition, when combined with subscriber management routings, to intelligently assign DHCP IP addresses, it becomes possible to offer differentiated classes of services for customers paying a premium.

1.7 Other IBM WebSphere Offerings

This section describes the other IBM WebSphere products that can interact with IBM WebSphere Performance Pack in an e-business Application Framework environment: IBM WebSphere Application Server and IBM WebSphere Studio.

Notice that IBM WebSphere Performance Pack can be integrated with a number of other IBM e-business products, such as IBM eNetwork Firewall, IBM HTTP Server and Lotus Domino Go Webserver. In this redbook, there are several examples and scenarios where all these e-business products are used together to create a powerful and secure Web site.

1.7.1 IBM WebSphere Application Server

IBM WebSphere Application Server lets you achieve your *Write Once, Run Anywhere* goal for Java servlet development. The product consists of a Java-based servlet engine that is independent of both your Web server and its underlying operating system.

WebSphere Application Server offers a choice of server plug-ins that are compatible with the most popular server APIs. The supported Web servers are:

- IBM HTTP Server
- Apache Server
- Domino
- Lotus Domino Go Webserver
- Netscape Enterprise Server
- Netscape FastTrack Server
- Microsoft Internet Information Server

IBM WebSphere Application Server V2.0 is available in Standard Edition and Advanced Edition.[6] In addition to the servlet engine and plug-ins, WebSphere Application Server Standard Edition provides:

- Implementation of the JavaSoft Java Servlet API, plus extensions of and additions to the API

[6] See http://www.software.ibm.com/webservers/appserv/.

- Sample applications demonstrating the basic classes and the extensions
- The IBM WebSphere Application Server Manager, a graphical interface making it easy to:
  - Set options for loading local and remote servlets
  - Set initialization parameters
  - Manage servlets
  - Specify servlet aliases
  - Create servlet chains and filters
  - Administer and monitor Enterprise Java Services (EJS) components
  - Enable Lightweight Directory Access Protocol (LDAP) directory support
  - Log servlet messages
  - Enable JVM debugging
  - Monitor resources used by Application Server
  - Monitor loaded servlets, active servlet sessions, and JDBC connections
  - Monitor errors, events, exceptions, and log output
  - Create dumps and data snapshots
  - Dynamically enable and disable tracing
- A connection management feature that caches and reuses connections to your JDBC-compliant databases

  When a servlet needs a database connection, it can get one from the pool of available connections, eliminating the overhead required to open a new connection for each request.
- Additional Java classes, coded to the JavaBeans specification, that allow programmers to access JDBC-compliant databases

  These data access beans provide enhanced function while hiding the complexity of using relational databases. They can be used in a visual manner in an integrated development environment.
- Support for dynamic page content called JavaServer Pages (JSP)

  JSP technology lets you produce dynamic Web pages with server-side scripting. The result is to separate your presentation logic (for example, the HTML code that defines your Web site structure and appearance) from your business logic (for example, the Java code that accesses a database for information to display on the Web site). For flexibility, JSP files can include any combination of inline Java, `<SERVLET>` tags, National Center for Supercomputing Applications (NCSA) tags, and JavaBeans.
- Enablement for LDAP supported directory services
- Modules and a command line interface for integrating Application Server and Apache Server into the Tivoli environment for distributed monitoring and operations
- eXtensible Markup Language (XML) Document Structure Services

WebSphere Application Server Advanced Edition provides all the features of Standard Edition, plus:

- Enterprise Java Services

  This function is provided to run and manage applications coded to Sun's Enterprise JavaBeans (EJB) specification.

- Common Object Request Broker Architecture (CORBA) support, enhanced to provide both bean-managed and container-managed persistence.

IBM WebSphere Application Server will also be available in an Enterprise Edition, which will include the same features as the Advanced Edition, plus:

- TXSeries support, IBM's world-class transactional application environment
- Component Broker (CB) support, with its fully distributed object and business process integration capabilities

1.7.2 IBM WebSphere Studio

IBM WebSphere Studio is a suite of tools that can be used by people involved in creating and maintaining Web sites. It allows your team to:

- Easily create Java beans, database queries, and Java servlets using the Studio wizards
- Group your Web site files into projects and folders
- Edit and update the files with your preferred tools
- Publish all or part of the Web site on any of your WebSphere Application Server systems
- Maintain the files locally on individual workstations, or in a central location using a source control system

IBM WebSphere Studio is the tool complement of IBM WebSphere Application Server. This suite of tools makes it easier to design, develop and maintain dynamic, interactive Web sites.

The tools provided by IBM WebSphere Studio are:

- Studio workbench

  The workbench helps you manage and maintain your Web site applications and files. It uses projects and folders to group and organize the files and lets you perform all other functions from this central location. The Studio supports all file types. You can create files, open them, edit them, copy them, move them around, delete them, publish them. And, you can do all this right from the workbench using the workbench menu functions, the wizards, and your preferred Web development software.

- Studio wizards

  The wizards are the fastest way to add dynamic content to your Web pages. They help you retrieve information from common databases, use server-side JavaBeans, capture information about your customers, and register Web visitors. You don't have to be expert at SQL syntax or Java programming. The wizards walk you through step-by-step and then generate sophisticated servlet code for you.

- Companion products

For added convenience the Studio comes with an integrated set of companion products. Everything is included for you to build, manage, and publish complete Web sites:

- NetObjects ScriptBuilder

  Use this text-based editor for files that contain common markup, scripting, and programming languages, such as HTML, DHTML, and JavaServer Page (JSP) extensions, JavaScript, JScript, and Java. Its features make the scripting process easy. You can quickly preview scripted Web pages, reference and add language elements, and navigate to embedded functions and objects.

- NetObjects Fusion

  Build your Web sites in a visual manner, dragging and dropping pages to create an interconnected hierarchy. You can view individual pages or the entire site structure. This powerful tool lets you include images, multi-media, and dynamic HTML, and control the visual appearance of an entire site from one central location.

- NetObjects BeanBuilder

  BeanBuilder allows you to put an applet in your Web site in a very short time. This visual authoring tool lets you quickly combine Java beans into new applets.

- VisualAge for Java, Professional Edition

  If you are familiar with Java programming, you can use this robust, full-function environment to create and customize Java components. This award-winning tool includes advanced functions, such as incremental compilation and the ability to invoke methods while debugging. You can use it to build sophisticated Java beans that you can use in the Studio basic servlet wizard. And, in turn, you can also use it to modify the Java servlets and beans generated by the Studio wizards.

- WebSphere Application Server

  What good is server-side Java logic without a server that can handle it? For your convenience, the best place to publish your Web sites is right in the package. The WebSphere Studio comes with a complimentary copy and a developer's license for the WebSphere Application Server.

- Base HTTP servers

  The WebSphere Application Server runs on several HTTP servers and more than one operating system. But, just in case you don't have one handy, the Apache Web server and a base HTTP server are included for good measure.

Chapter 2. IBM Web Traffic Express Concepts

IBM Web Traffic Express Version 2, the caching and filtering component of IBM WebSphere Performance Pack V2, is a caching proxy server that provides highly scalable caching and filtering functions associated with receiving requests and serving URLs. With tunable caching capable of supporting high cache hit rates, this component can reduce bandwidth costs and provide more consistent rapid customer response times. Using WTE, `Unable to Connect to Server` errors, which today are so frequent, will be dramatically reduced.

- WTE is a *proxy server*, which means that it assumes the responsibility for retrieving Internet data for multiple browser clients. Client requests are sent to the Web servers through the proxy. In other words, the client may be configured to send its request to the proxy first, and then it is the proxy that forwards the client's request to the Web server, acting on behalf of the originating client. The Web server does not even see the IP address of the client in the packet headers, but only the IP address of the proxy server. Once the proxy receives the information from the Web server, it forwards the information to the requesting client. This way the proxy function can be used to provide address security and optionally, through specific proxy features, to support additional functions, such as request filtering or modification.
- A traditional proxy server receives a request for a URL from a client and then forwards the request to the destination Web server. WTE has a highly configurable *caching* functionality. This means that it can save, or cache, the Web documents it retrieves that are considered cacheable according to the HTTP protocol and administrator defined overrides. It can then serve subsequent requests for cached documents from its local cache. The client gets the information faster and network bandwidth utilization is reduced.
- WTE has a *content filtering* functionality, implemented through Platform for Internet Content Selection (PICS) labels, which rate Web material by criteria such as language, nudity or violence. This is a consistent way to implement filtering on a broad scale and control the content you are providing. You can apply these filters in addition to, or instead of, the filtering set by the browsers. It is especially useful where end users do not (or should not) have access to those controls. Content filtering in WTE can use:
  - PICS rules to guide the use of rating labels - such as Recreational Software Advisory Council on the Internet (RSACi) criteria for inappropriate language, nudity or violence - placed in HTML or HTTP headers or third-party content rating label distributions
  - Lists of URLs/sites for which access is to be blocked
  - APIs for filtering applications

WTE uses an innovative caching scheme from IBM research laboratories to improve end user response time and/or optimize bandwidth utilization and costs. It allows you to customize its caching features to your own benefit and it offers key features of advanced caching. You can specify which pages are cached, when the information on a page will expire, how large to make the cache and when to update it.

When used in conjunction with ND and AFS, WTE becomes a massively scalable solution. Add IBM eNetwork Firewall to your architecture and you will have a scalable and secure solution for caching and filtering Internet information.

2.1 Why Do I Need IBM Web Traffic Express?

In this section we describe the major advantages of using WTE. To do this, we focus on the three major functions implemented in Web Traffic Express: proxy, caching and filtering.

2.1.1 Proxy Function

The WTE functions as a proxy server. It accepts HTTP, FTP and Gopher requests from clients, sends them to the respective Web servers, retrieves the data from the target servers on behalf of the originating clients and finally forwards the data to the requesting client. Security is ensured in this case as the client's IP address is transparent to the Web servers. Also, the machine load associated with making a URL request is transferred to a dedicated machine, thus allowing more control of the network.

WTE also supports automatic downloading of the proxy configuration file by the browser.

2.1.2 Caching Function

WTE functions as a caching server. Traditional proxy servers only receive a request from a client and forward the request to the designated Web servers. When the WTE receives the data from the target servers, it saves a copy in its local file system (caching). If another request is made for the same URL, WTE does not need to go back to the same target server. This greatly speeds up response time and frees up network bandwidth.

2.1.3 Filtering Function

WTE functions as a filtering server. It provides content filtering functionality, implemented through the Platform for Internet Content Selection (PICS) labels, to prevent clients using the proxy server from accessing certain types of data (typically, offensive Web contents). PICS rules can also be configured on some browsers; however, implementation would be tedious and the rules could easily be changed by a knowledgeable user.

WTE looks for PICS labels:

- In the data returned from the server (if the content is self-rated, it will have embedded labels, for example, using the RSACi rating system).
- By requesting information from a service bureau. It matches the PICS labels against the WTE PICS rules to decide whether to return the data to the client.

2.2 What Is New in Web Traffic Express Version 2

IBM WebSphere Performance Pack V2 offers several new features, many of which involve WTE V2 and make this product a unique offering on the market today.

2.2.1 Transparent Proxy

If WTE is configured to operate as a *transparent proxy*, the client software is totally unaware of the existence of the intermediate proxy server. Normally, if a client browser uses a proxy server, then the Web browser must be configured to specify the address and port of the proxy server. This is no longer necessary with transparent proxy, because the client is unaware that an intermediate proxy is in the network.

To use transparent proxy, the router - which may be an IBM 2216 or even, for low traffic volumes, SecureWay Network Dispatcher - is programmed to redirect requests to the WTE transparent proxy. WTE then intercepts all HTTP requests on port 80 that are targeted at some server out in the Internet. The request is parsed and processed, and may be satisfied from the transparent proxy's cache.

The transparent proxy feature is *currently* supported only on the AIX platform, and applies only to HTTP requests that do not require authentication.

2.2.2 Web Proxy Auto Discovery Support

As an alternative to transparent proxy, Web Proxy Auto Discovery (WPAD) is a mechanism to permit Web clients to locate nearby Web proxy caches. WPAD is currently an IETF Internet draft, but WTE V2 is already WPAD-compliant. You can find more details on WPAD at `http://eggplant.rte.microsoft.com/wpad/wpad.txt`.

On the client side, WPAD is supported on Microsoft Internet Explorer Version 5 or greater. Therefore, if your client platform is running this level of Microsoft Internet Explorer, you can launch the Automatic Discovery of Proxy Server feature of the browser, as shown in the following figure:

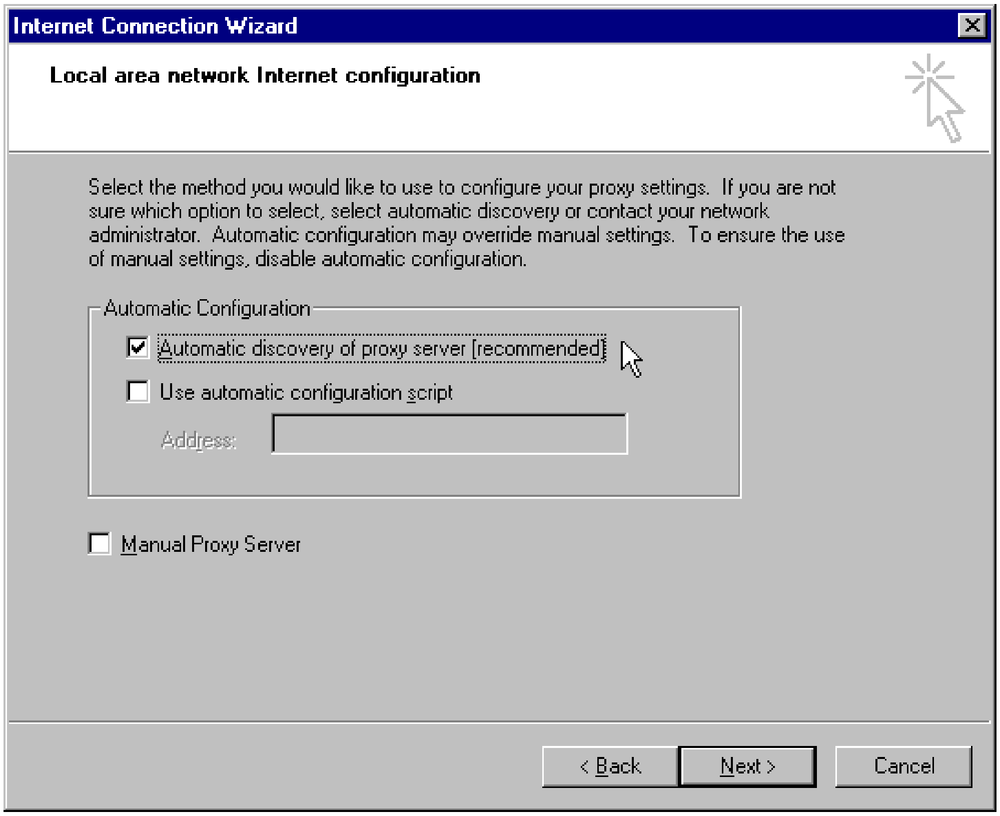
Internet Connection Wizard
Local area network Internet configuration

Select the method you would like to use to configure your proxy settings. If you are not sure which option to select, select automatic discovery or contact your network administrator. Automatic configuration may override manual settings. To ensure the use of manual settings, disable automatic configuration.

Automatic Configuration

☑ Automatic discovery of proxy server [recommended]

□ Use automatic configuration script

Address:

□ Manual Proxy Server

Figure 9. Automatic Discovery of Proxy Server

This will automatically detect a WPAD-compatible proxy server in the network, such as WTE, to which the browser will forward all Web requests, without the need for the browser administrator to issue any other configuration.

2.2.3 FTP Proxy Enhancements

This line item consists of several enhancements to the WTE FTP proxy code, including the addition of FTP PUT capability, improved authentication to prompt for the user ID and password instead of requiring it in the URL, and a configuration directive to allow the user to specify whether FTP URLs will be treated as relative URLs or absolute URLs.

2.2.4 Proxy Autoconfiguration Support

Netscape Communicator V2 and Microsoft Internet Explorer V4, and later versions, have a feature called *automatic proxy configuration*, that is supported by WTE. This function provides a form of transparency, in that the users on the client-side do not have to configure their Web browsers to point to a specific proxy or SOCKS server, but to the automatic configuration file instead. This feature also allows the system administrator to modify the proxy configuration with little impact to the clients. Proxy autoconfiguration support offers a form of proxy high availability, because a secondary proxy can be specified in the configuration file, which will take over in case the primary proxy should fail.

2.2.5 Remote Cache Access

In many deployment scenarios, multiple proxies are located *near* each other in network terms. Typically they will be front-ended by an ND machine for load balancing and proxy high availability reasons. In WebSphere Performance Pack V1, each proxy had to have its own cache, and the cached data could not be shared. This resulted in cache space being wasted to store multiple copies of the same document. A more serious problem is that if you have N nodes in your cluster, then each file will be transferred across the network link up to N times. When the link is very expensive, such as a transatlantic link, or if you are paying for each packet sent across the link, or the link itself is very slow, you may not want to send the same packet on the network more than one time. The Remote Cache Access (RCA) function of WTE addresses those problems by providing a means for sharing cache content between proxies.

2.2.6 Tivoli Ready

WTE V2 is *Tivoli Ready*. IBM software products that are Tivoli Ready can be managed through either Tivoli Enterprise Console (TEC) or Tivoli Global Enterprise Manager (GEM).

Tivoli Ready instrumentation, when configured, provides you with the ability to:

- Graphically view the health of the WTE server through Tivoli GEM V2.2, TEC V3.1 or higher consoles.
- Inventory WTE using Tivoli Inventory V3.2.

2.2.7 SNMP Subagent and MIB Support

WTE V2 provides an SNMP management information base (MIB) and SNMP subagent so you can use any SNMP-capable network management system, such as Tivoli NetView or Tivoli Distributed Monitoring to monitor your proxy server's health, throughput, and activity. The MIB data describes the proxy server being managed, reflects current and recent server status, and provides server statistics.

2.2.8 Performance Improvements

The cache architecture has been restructured in WTE V2, giving a more efficient mapping of URLs onto the cache filesystem. This provides quicker retrieval of cached objects, more efficient use of disk space within the cache, and quicker cache garbage collection. The cache also uses write-behind techniques for greater throughput. Additionally, WTE now includes caching of Domain Name Server (DNS) lookup results, giving quicker response time and further reducing network load.

2.2.9 Customization Exits (ICAPI Extensions)

Several enhancements have been made to the WTE API to simplify the job of writing an application plug-in. New request steps (exit points) have been added: the transmogrifier and the garbage collection (GC) advisor.

- The *transmogrifier* gives the application write access to the outgoing data stream. It is intended to be used by applications that wish to perform transformations on the HTTP response data stream. Examples include converting PDF files to HTML, converting high resolution images to lower resolution quality, or translating pages from one language to another. The transmogrifier step is an extension of the data filter step. The WTE enhancements in WebSphere Performance Pack V2 allow the application to specify multiple transmogrifiers, thus allowing the application to have multiple plug-ins that each perform different transformations on the data. In WebSphere Performance Pack V2, WTE introduces a correlator mechanism that eases state maintenance in the plug-in, and now that WTE automatically determines the content length of the response data, the application no longer has to buffer the data to determine the content length. The new version of WTE also makes HTTP header processing easier. The response headers can now be extracted and set using API variables.

- The *GC Advisor* allows the plug-in to influence garbage collection decisions. It is called for each file in the cache during the garbage collection process and allows the application to influence which files are kept and which are discarded.

2.2.10 HTTP 1.1 Compliance

This line item adds *HTTP 1.1 compliance* to WTE as a proxy server. Prior to the 2.0 release, WTE partially implemented the HTTP 1.1 RFC by supporting client-to-proxy persistent connections and handling receipt of HTTP 1.1 requests from the client.

With the full implementation of HTTP 1.1 for WTE 2.0, WTE now identifies itself as an HTTP 1.1 server. Persistent connections are supported not only from the client to the proxy, but also from the proxy to the content server. The HTTP 1.1 cache control headers are processed and used to determine if a Web document is cacheable. WTE receives and processes chunked data sent by HTTP 1.1 content servers. WTE will unchunk data before giving control to a data filter or transmogrifier plug-in.

In WebSphere Performance Pack V2, WTE provides new directives to allow the administrator to override certain HTTP 1.1 cache control headers. For example, query strings (URLs with a question mark `?` in them) are not considered cacheable by the HTTP 1.1 protocol, but WTE provides a directive to allow the administrator to specify that query strings should be cached. WTE also provides an aggressive caching directive that allows the administrator to override the following header in Web documents:

```
Cache-Control: no-cache
```

2.2.11 Variant Caching

Variant caching extends the capabilities of the transmogrifier and allows WTE applications to request that WTE cache a variant of the original document retrieved from the Web. This need arises when a plug-in is performing a transformation on the original page retrieved from the Web. For example, if a plug-in is translating a page from English to Italian, it would be advantageous to be able to cache not only the original document but the translated variant of the original as well. Transformations of Web data are very CPU intensive and degrade the performance of the proxy server. By caching the variant, the number of transformations required is decreased, thus improving overall proxy caching performance.

The WTE API provides new predefined functions, HTTPD_variant_lookup() and HTTPD_variant_insert(), to allow the application to find a variant that has already been cached, and to insert a new variant into the cache.

2.2.12 Cancel Control

The *cancel control* function lets the administrator set a parameter that allows the proxy server to cache certain documents even if the user disconnects or selects **Stop** before the entire document has been retrieved. The parameter value represents a percentage of the file being received. If the amount of the file already received is equal to or greater than the value, the whole file will be received and stored.

2.2.13 Peak Load Management

In Version 1, IBM eNetwork Dispatcher[1] allowed rule-based routing to alter the load-balancing algorithm based on the current connection rate, the number of active connections, and the time of day.

In Version 2, a WTE advisor for SecureWay Network Dispatcher enhances load management on WTE by preventing ND from sending new requests to a WTE node that is engaged in garbage collection or cache refresh. This new feature is called *peak load management* and is part of the quality-of-service enhancements offered by IBM WebSphere Performance Pack V2.

2.2.14 Content Based Routing

Content Based Routing (CBR) is a new function, created by combining the two components, WTE and ND, of IBM WebSphere Performance Pack. It provides load-balancing enhancements to the reverse proxy capabilities of WTE. Both the ND CBR component and WTE must be installed on the same machine. ND rules can be written to load-balance requests over different sets of servers based on the client IP address, the entire URL, the protocol portion of the URL, the host portion of the URL, the path portion of the URL, the Referer HTTP header, or the User-Agent HTTP header.

2.2.15 Integrated Configuration Assistance

WTE configuration is stored in a text configuration file. Rather than manually editing this file, it is advisable to use the Configuration and Administration Forms, Java-enhanced Web pages that can be accessed from a Web browser only by authorized and authenticated users. The Configuration and Administration Forms offer a user-friendly interface and take care of storing the exact directives in the configuration file, eliminating the possibility of inadvertent syntax mistakes.

Moreover, a set of *task guides* or *wizards* has been added to IBM WebSphere Performance Pack V2 in order to make configuration easier. These task guides are invoked from a Web browser as well. In particular, a task guide has been added to set up PICS filtering rules for WTE and achieve the proxy filtering value:

[1] In IBM WebSphere Performance Pack Version 1, this was the first name of the Load Balancing component.

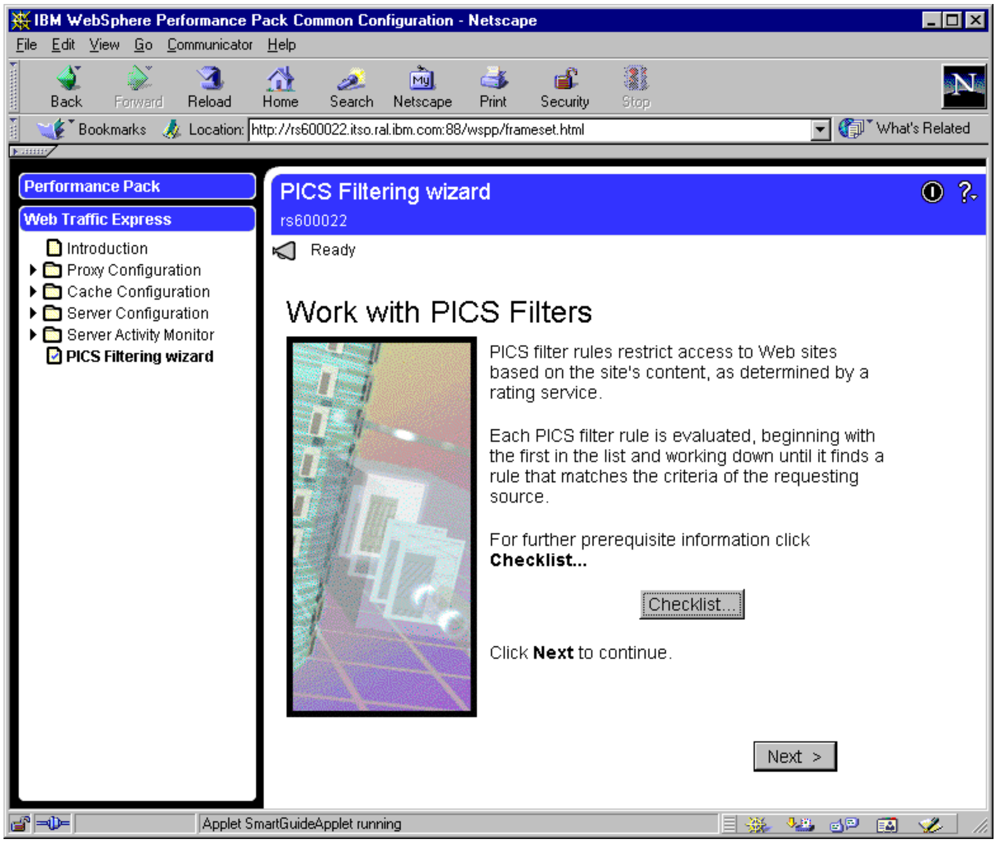
IBM WebSphere Performance Pack Common Configuration - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com:88/wspp/frameset.html
What's Related
Performance Pack
Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filtering wizard
PICS Filtering wizard
rs600022
Ready
Checklist...
Next >
Applet SmartGuideApplet running

Work with PICS Filters

PICS filter rules restrict access to Web sites based on the site's content, as determined by a rating service.

Each PICS filter rule is evaluated, beginning with the first in the list and working down until it finds a rule that matches the criteria of the requesting source.

For further prerequisite information click **Checklist...**

Click **Next** to continue.

Figure 10. PICS Filtering Wizard

2.3 Who Can Benefit

WTE provides a valuable and scalable solution to some of the major traffic management problems. These are the main advantages it offers:

- Reduction of costs and constraints on network bandwidth, particularly during periods of peak concurrent activity
- Scalable infrastructure that provides cost-effective growth paths and essentially unlimited capacity potential with minimum redesign or disruption
- Bandwidth management capabilities based upon content filtering and proxy functions
- Content management capabilities based upon industry-standard filtering technologies to restrict information or to enhance information provided to users
- Functional openness to permit evolution and exploitation of emerging Internet capabilities with minimal impact on overall network architecture
- Multiple platform support to simplify implementation planning and skill requirements

2.3.1 Caching Proxy Function

The caching proxy function provided by WTE is valuable to customers and/or ISPs needing to optimize line costs and performance associated with accessing remote Web sites. Customers that can benefit from using the caching and proxy function of a caching and filtering proxy server include:

- Service providers needing to provide good response time to clients from Web sites accessible only via expensive or distant links that experience significant propagation delay time.

  They need to be able to provide nondisruptive access to information on servers located within their networks as well as those external to their networks. Many ISPs must have infrastructures capable of cost-effective expansion to handle growth rates greater than 10% per month.

- Enterprises with significant external Web access (Internet and/or intranet) wanting to optimize wide-area line usage.

  Large corporations and university campuses are the most likely to benefit from caching, to improve response time while optimizing external links.

2.3.2 Filtering Function

The filtering function of WTE is particularly valuable to customers and/or ISPs who want to filter out content from Web sites on the basis of defined codes, rules or APIs. Such customers include:

- Service providers who want to enable PICS filtering for consumers or to enforce regulations prohibiting certain types of content.
- Service providers who want to filter content to suit the needs of closed user groups, such as subscriber-based communities of interest. Such filtering would likely reflect non-standard rules, programmed to suit the preferences of a particular interest group.
- Service providers who want to filter or intercept content to enable additional processing such as premium content (authentication or billing) or language translation.
- Business network managers who want to minimize traffic on backbone network links during peak periods by filtering traffic according to application or content priorities.
- Education institutions, particularly grade schools, that want to control student access to nonapproved content.

2.4 Increasing Hit Bytes and Hit Rates

Caching effectiveness is typically expressed in terms of *hit bytes* or *hit rates*, referring to the number of bytes retrieved from a cache (and therefore not transmitted over the network) and the number of cached objects retrieved as a percentage of all objects served. In fact, hit bytes are more commonly utilized as a measurement in operational environments such as backbone networks, where bandwidth and uplink network access costs are the reason for caching. Hit rates, on the other hand, are the measurement most likely to be referenced in those environments such as user access points, where consistency and speed of user response time is the key reason for caching.

WTE has several features that result in an enhancement of caching effectiveness:

- Increase cache capacity

WTE supports multiple platforms, including highly scalable clustered nodes such as the SP2.

- Preload cache with known high demand content

  WTE permits preloading of information on the basis of named content and prefetching of content with automatic location and fetching of linked objects to a prespecified *depth*.

- Monitor usage to support tuning

  WTE supports logging that can be used as input to preloading and other administrative activities.

- Purge content least costly to retrieve, and keep content most expensive to retrieve (that is, consider size of object in addition to usage indicators)

  WTE enables network administrators to specify weighting criteria to reflect their specific cost trade-offs between cache capacity and bandwidth costs. WTE also permits administrators to purge or keep content based upon `time-to-live` headers and the relative size of the objects.

- In larger networks, use caching servers in structured hierarchy, that is, second level caching

  WTE can be configured to recognize a *parent cache* that will be searched next in the event a specific request is not found in the local cache. Multiple tiers can be supported, subject to trade-offs in latency for the end users between cache processing time versus network access.

2.5 Cache Content Management

Most network managers that use a caching function to minimize network bandwidth costs want to get the best *hit rate* possible, that is, to increase the probability that user requests can be satisfied from content within the cache versus a network request.

There are several decisions an administrator needs to make:

- Which documents are kept in the cache?
- How many documents can be cached?
- How long are they considered current?
- How are the documents indexed?
- When is the cache refreshed?

This section gives an overview of how caching works.

2.5.1 Controlling Which Documents Are Kept in the Cache

By configuring WTE, you can specify which documents should be cached, how long they should be cached, and which documents should never be cached.

Notice, however, that some files are never cached:

- Documents that were requested through HTTP methods other than `GET`, such as `POST` or `PUT`

- All the documents that were obtained after authentication or payment

- All the documents that were dynamically generated by CGI-BIN scripts or Java servlets

- Any information passed on an SSL connection, because the proxy cannot decrypt the data passing through it
- Any URL containing a question mark (`?`) in it

2.5.2 Cache Freshness

It is important that the WTE administrator ensure that cached documents are consistent with the original documents located at the originating Web server. In other words, WTE must ensure *cache freshness*.

For each document that has been cached, WTE computes a time at which the document will expire:

- For HTTP documents, it is the header of the document generated by the Web server that contains the expiration information.
- For FTP documents, it is WTE that generates its own `Last-Modified:` header information to compute expiration times. The reason for this is that the FTP protocol does not include expiration information, unlike the HTTP protocol.

The Web server can indicate the expiration time in several ways, putting header information in the HTTP response. The permitted header information is in the following order of preference:

1. The Web server can specify for how long the document is good after it has been received.
2. The Web server can specify the exact time at which the document should be considered expired.
3. The Web server can indicate the time when a document was last modified. Then the caching and filtering proxy server performs a sequence of operations based on a class of parameters specified in the configuration file and calculates how long the document will be good.

When a document is found in the cache, but it has expired, WTE issues a special request to the Web server. This request is known as `if-modified-since`. The Web server sends back the document to the caching and filtering proxy server only if such a document has been modified since it was last received by the proxy. Otherwise, the Web server only sends a message indicating that the document has not been modified, and does not send the entire file. At this point, the caching and filtering proxy server can serve the page to the client.

The caching and filtering proxy server administrator is also allowed to specify how long to keep unused cached files.

2.5.3 Cache Size and Garbage Collection

Disk space and file maintenance are common concerns when using a cache. WTE allows you to control the amount of disk space used for the entire cache.

Another important feature implemented in WTE is the nightly cleanup process known as *garbage collection*. This process examines the files in the cache directory and attempts to remove old, expired or unused files to make room for more current files.

There are two algorithms that the garbage collection process can use when deciding which files to remove and which files to keep in the cache. One algorithm maximizes the cache to improve user response time and the other maximizes the cache to minimize network bandwidth:

1. When you are tuning your cache to minimize response time, larger files are given a higher priority for deletion and, therefore, are more likely to be removed during garbage collection.
2. When you are tuning your cache to minimize network bandwidth, larger files are given a lower priority for deletion and they are less likely to be removed during garbage collection.

2.5.4 Cache Indexing

WTE implements a cache directory structure and lookup methods that are different from many other proxy servers. It creates an index of the files in the cache to keep in memory as each page is added. RAM memory is used instead of other media, so that the lookup operation and retrieval times are faster.

The index separates the cached files into a set of sub-caches, and for each file in the cache the index stores in memory the file name, URL and expiration information. For this reason, the RAM memory required is directly proportional to the number of files in the cache.

When the caching and filtering proxy server receives a request from a client, the proxy checks the index in memory for that particular URL:

- If the file is not in the index, the request is made to the destination server. The retrieved URL is then checked to ensure that the document is cacheable, and the document is cached if this operation is permitted. The index is then updated with the new URL, sub-cache and expiration information.
- If the file is in the index, the expiration information is checked to see if the URL is stale. If the URL has expired, the caching and filtering proxy server contacts the content destination server, and the URL is replaced by the newly retrieved document with expiration and sub-cache information updated in the index. If the URL is still consistent, the document is served.

The cache contains shadow files that mirror the index information for the proxy to use only when the proxy server is started. The garbage collection process updates the cached document index files.

2.5.5 Automatic Cache Refreshing

Typically, the most common proxy servers cache a particular page only after a user requests it. WTE, in addition to this default caching, has a *cache agent* that provides automatic caching and gives more control to the administrator. The cache agent can retrieve specified URLs even before they are effectively requested and refresh the cache automatically. The cache refresh takes place when the proxy server activity is low (by default, every night at midnight, local time) and all the retrieved pages are ready in the cache to provide faster service even the first time a user requests them.

The automatic cache refreshing has two sources it can use to refresh the cache:

1. It can load specific URLs defined by the administrator. In this way, the administrator can specify a certain set of pages that must be loaded by the cache agent when it starts.

2. It can load the most popular URLs from the previous day’s activity. To obtain this, the cache agent checks the cache access log, sorts it by frequency of requests and then picks the most popular pages. It can refresh the top number of pages as specified by the administrator.

Notice that the cache agent can use both sources of input.

Optionally, the cache agent can follow a specified level of HTML links on the pages it is loading and cache all of those linked pages. This operation is also known as *delving*. It is not necessary that the linked pages reside on the same host as the source page, since the cache agent can retrieve them even if they reside on other hosts.

The cache agent offers a very useful service. Using the cache agent, caching is performed even before cached pages are effectively requested, so the average response time is minimized. Moreover the cache is built before user activity gets busy, typically at night. However, turning the cache agent on forces the caching and filtering proxy server machine to be busy even during hours of low activity. Moreover, configuring the cache agent to perform delving requires giving more control to the caching and filtering proxy server administrator. For example disabling delving from high-level pages, such as Web indexes or search sites, is recommended, or multiple requests for large numbers of pages will be generated.

2.5.6 Remote Cache Access

Remote Cache Access (RCA) is a new, powerful feature implemented in WTE V2, and not available in the previous release of IBM WebSphere Performance Pack. RCA allows multiple proxy servers to cooperate to form cache arrays. Using RCA, multiple proxy servers can distribute the cache contents across their combined, logical cache to improve hit rates and reduce redundancy of cached content.

WTE uses new caching algorithms and information-sharing technologies to enable an ISP to manage its servers more efficiently by storing information where it is more likely to be needed and delivering it more efficiently to customers. These enhancements reduce transmission costs and eliminate the need for ISPs to replicate information in redundant proxy servers.

2.5.6.1 How RCA Works

RCA enables peer WTE proxy servers, in close physical proximity to each other, to share the contents of their caches utilizing a shared file system, such as AFS, DFS, NFS, or Windows NT file sharing. We recommend the use of AFS, since it offers nondisruptive real-time replication of information across multiple servers, data consistency, availability, global stability and data consistency.

The peer servers must be configured as an array. RCA uses the Cache Array Routing Protocol (CARP) to determine which peer in the array should process the incoming request. If the requested file is not in the proxy’s cache, it will query the other peers in the array to determine which proxy might have the object cached. If the object is in another cache, the owning proxy indicates where the file can be found in the shared file system. If the object is not in any cache, the proxy processing the request will get the file and then cache it.

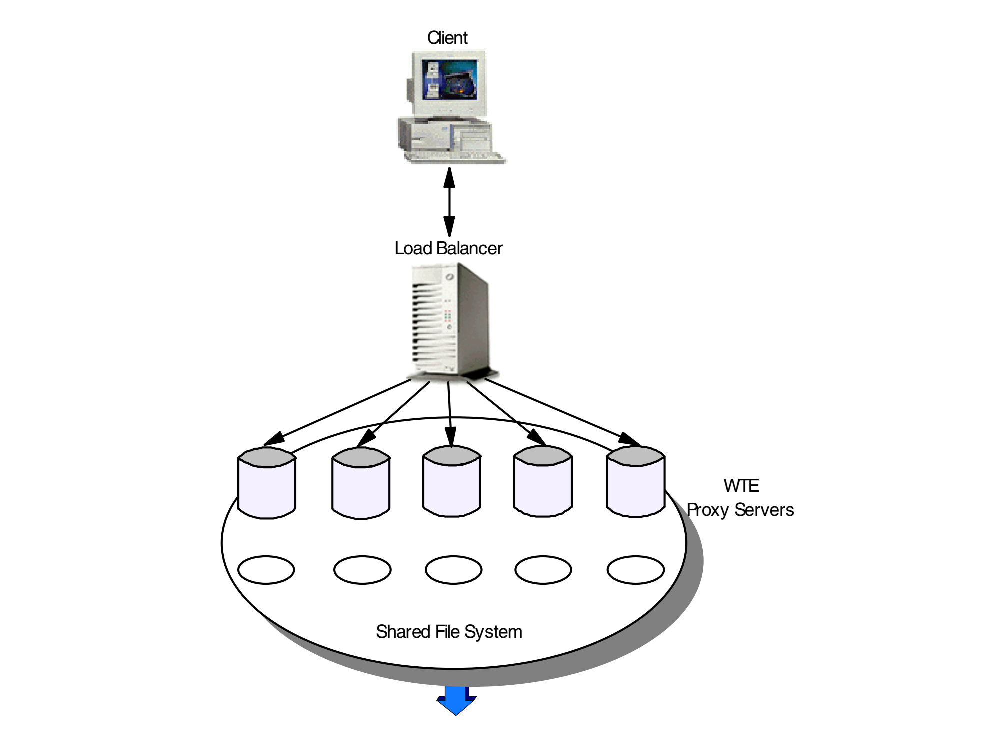

Figure 11. How RCA Works

As you can see from Figure 11, RCA is one of the new features of IBM WebSphere Performance Pack that can work by combining together all the components.

2.5.6.2 Improved Cache Hit Rates with Innovative Algorithms

WTE uses enhanced caching algorithms aimed at increasing the likelihood that a request for an object on the Web will find it in the first cache it searches. The algorithms improve *hit probability* by streamlining the ways this local cache obtains and temporarily stores information from remote caches.

Traditional caching algorithms were designed to manage objects strictly on the basis of frequency of use, independent of their size and movement across a network. In reality, Internet objects vary in size, as do the costs of keeping them in cache and distributing them to local subscribers. The new algorithms used by WTE are optimized for ISP environments, where large-scale performance and bandwidth costs are critical. These innovative algorithms guide the proxy server or servers to manage objects based on ISP requirements for reducing bandwidth costs and improving performance for local subscribers.

2.5.6.3 Flexible Caching Configuration

WTE enables ISPs to distribute the caching function across multiple servers to improve performance and response time. Providing more local caching can also result in savings in line costs. In addition, WTE also improves information flow by keeping each server in the system aware of what is available in the cache of its peers. This cache array enhancement reduces storage costs by minimizing the need for redundant data, and improves performance with the addition of servers in the cache array as required. Because it is scalable - readily expandable in size through the addition of servers as needed - the solution allows ISPs to grow as their subscriber base grows. CARP-like replication further enhances performance by minimizing the number of index hits required for the average search. This combination of central and local caching gives the network administrator the control and flexibility needed to maintain a high performance network.

2.6 Flexible-Client SOCKS

WTE has a particular feature, named *flexible-client SOCKS*, that allows the caching and filtering proxy server to reside behind a firewall or SOCKS server without sharing the same physical machine. Requests going to the proxy can then be routed directly to the destination Web server, instead of sending all requests through the SOCKS server. None of the components of IBM WebSphere Performance Pack include a SOCKS server. We recommend that you install the SOCKS server provided by IBM eNetwork Firewall.

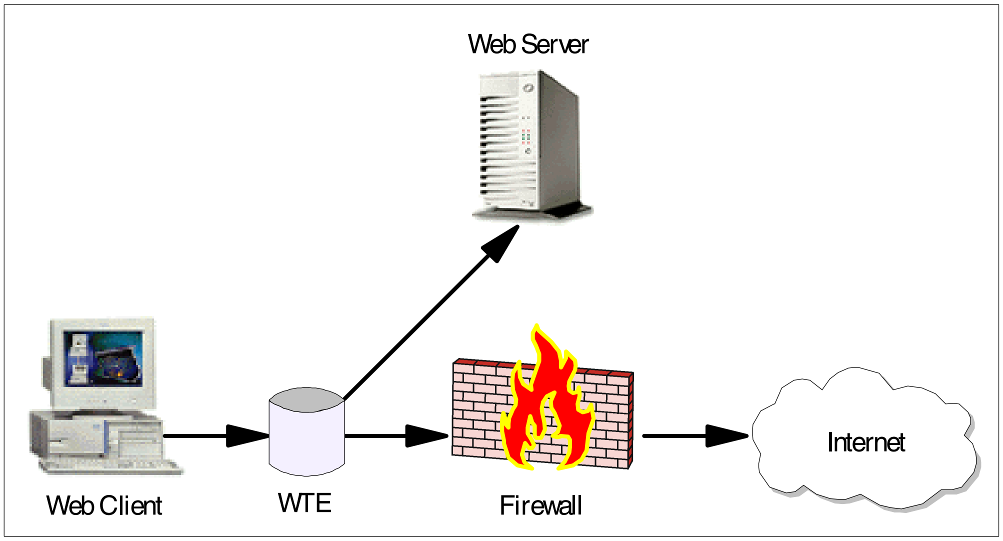

Figure 12. Flexible-Client SOCKS in WTE

The flexible-client SOCKS functionality helps security by allowing a firewall server to be isolated from the proxy server, even if this requires additional hardware and can produce higher latency on requests. The load on the firewall is reduced by having the caching and filtering proxy server handle internal requests. Moreover the administrator can easily specify which requests the caching and filtering proxy server sends to the SOCKS server and which requests it redirects back to the local domain.

A traditional proxy server installed behind a firewall would route all the requests to the firewall itself. On the other hand, the flexible-client SOCKS provided by WTE lets you specify which IP addresses or domains should be contacted directly by the proxy server and which ones should be contacted through the SOCKS server.

2.7 Proxy Chaining

Proxy chaining is a mechanism that allows you to create a hierarchical chain of proxies, each proxy belonging to a certain level in the hierarchy. If a proxy server in the lowest level of the hierarchy cannot serve the requested URL from its cache, it does not forward the client's request to the content Web server, but to the proxy server that has been configured as the proxy server of the immediately higher level in the hierarchy. The higher the level in the chain, the larger the number of users that access that proxy, so the possibility that a proxy server at a higher level in the chain finds the requested document in its cache becomes greater.

The proxy server at the highest level of the chain contacts the Web content server to retrieve the documents if such a document was not in its cache. After that, the proxy server passes the document back down the hierarchy, and all the proxies under it, including the originating proxy server, cache the document. The originating proxy server also serves the document to the requesting client.

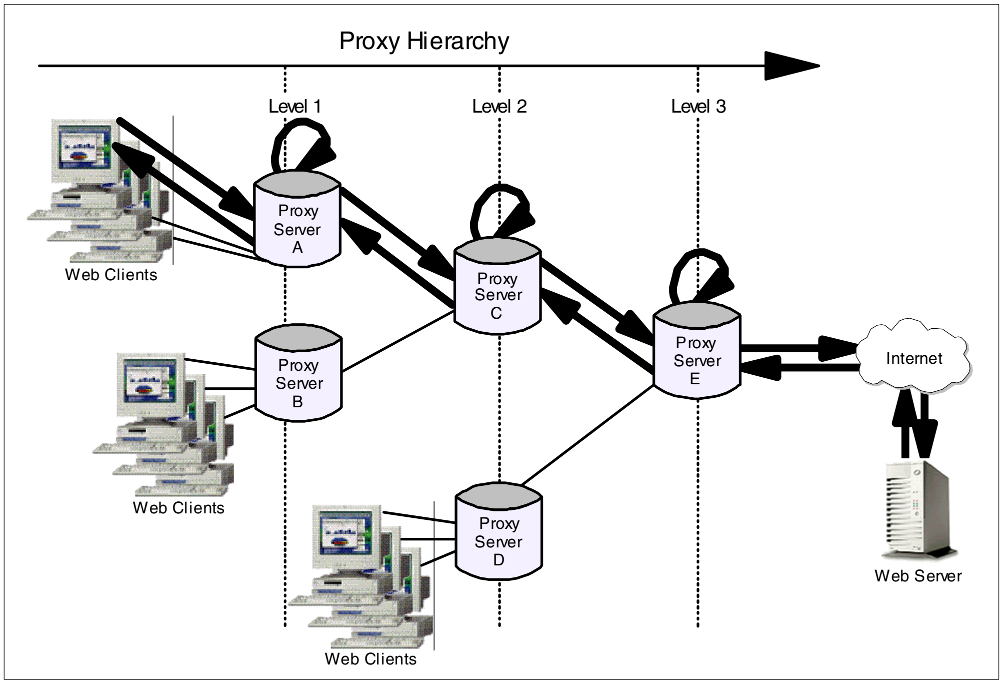

Figure 13. Graphical Representation of Proxy Chaining

Notice that the request goes through the daisy chain of servers and the response passes through the same proxies but in the reverse order. If intermediate proxies have caching enabled, each of them would search its own cache for the requested resource and return a cached copy, without forwarding the request to the originating Web server.

This architecture offers the following advantages:

- Proxies at a lower level, which are closer to the client that originated the request, benefit from the caches of the higher-level proxies.
- Proxy chaining reduces the load on the highest level proxy (typically, the proxy closest to the firewall) and ultimately on the Web server, since lower-level proxies may already have the document cached.
- The larger the number of users, the higher the probability that the proxy server already has the document in its cache. Considering that high-level proxies serve a larger number of clients, many requests that cannot be honored by a

low-level proxy can be resolved by higher-level proxies, since other groups of clients may have already requested the same files.

However, it should be kept in mind that:

- A high-level proxy should have a larger cache, since it has to honor the requests of a large number of users.
- Proxy chaining greatly increases response time for requests, especially for those files that have not been cached yet by any proxies in the hierarchy.
- The risk of failure in a chain increases with each additional node.

WTE allows the creation of proxy chains based on the protocol (HTTP, FTP or Gopher) of the request to be forwarded. In other words, a proxy server can be configured to send all incoming requests with a particular protocol to a higher level proxy server in the chain.

It is important to note that the more proxies you chain together, the greater the latency you introduce. Chaining two proxies is the recommended limit. You may not need to chain proxies unless you are in a large organization where you would place *small* proxies close to the users (in a branch office for example) and then chain to a *large* corporate proxy that connects directly to the Internet.

2.8 PICS Filtering at the Proxy Server Level

Platform for Internet Content Selection (PICS) is a technological standard that builds up a content labeling and filtering system for Web information. It is formed by a set of specifications used to create and manage ratings for the information published on any given Web site. The idea behind PICS is relatively simple. Since people have all kinds of different preferences and values, there should be a ratings and labeling standard that people can use to choose what content they would welcome seeing, or wish to exclude, based on certain parameters.

PICS development began in mid-1995, when the computing and online industries became sensitized to the possibility of online content censorship by the U.S. government and other governments around the world. Working under the aegis of the World Wide Web Consortium (W3C), the interested parties came together to create a technological solution that would support different rating systems. The PICS standard was adopted in May 1996. For the most up-to-date PICS information, see the World Wide Web Consortium's PICS Web site at `http://www.w3.org/PICS/`.

PICS is a standard that lists rules for rating the information contained on any given Web site. Rating decisions are made on the basis of particular categories, usually violence, pornography, language and nudity. Some Web browsers, such as Microsoft Internet Explorer 4.0, are PICS-compliant and are able to filter content information as content is received from the destination server. However, relying on browser settings may not be a safe solution, because such settings can be adjusted at the browser, and may be easily compromised.

With WTE, you can implement PICS filtering at the proxy server level. This removes the responsibility from the client, and your proxy server administrators can directly prevent certain types of information from being served to specific browsers (or to groups of browsers). Browsers with PICS filter settings defined will then be able to perform further filtering of the Web pages they will be served.

This method ensures that clients will get only the level of content specified at the proxy. Interaction between the client and the proxy administrator would be required to change the sensitivity of the filter.

Using this centralized approach, the PICS filtering process is transparent to all users. When users request an HTML page, they will see the requested URL or get an error message similar to the following:

```
Error 403
Blocked by Filtering Rule
```

WTE allows an administrator to specify filtering rules based on PICS labels. When a URL is accessed, the caching and filtering proxy server uses these rules to determine if it passed or failed. The PICS labels can be supplied to the caching and filtering proxy server in several ways. They can be stored locally on the proxy's hard disk, supplied by a label bureau or even provided by the Web server. Some URLs have the label embedded within their HTML files under the <META> tag or in the HTTP response header. Security mechanisms, such as message digests and digital signatures, can be incorporated in the label creation to grant label validity.

The diagram below shows the logical flow of the PICS filtering at a proxy server level, assuming that the PICS labels are provided by a label bureau:

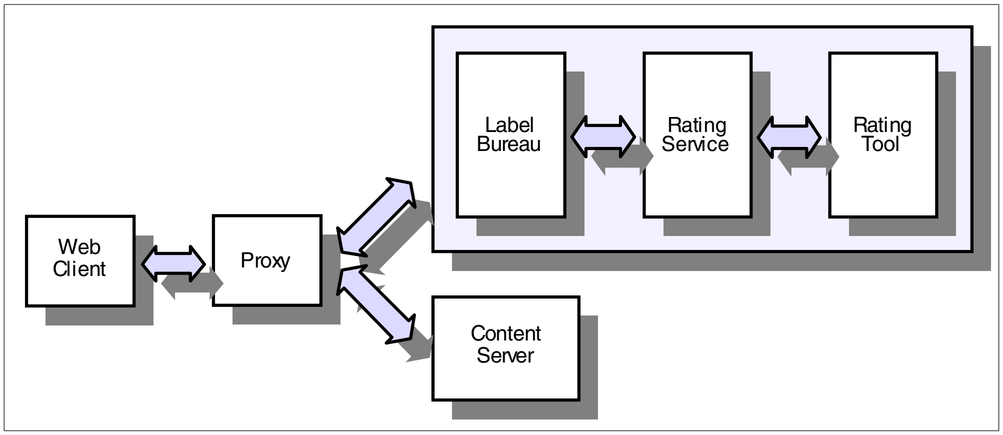

Figure 14. PICS Filtering at the Proxy Server Level

Here, the proxy server administrator has set up and enabled a filtering profile on the proxy server. To decide whether a particular document will be passed or blocked, the proxy filtering profile uses the values contained in the PICS labels supplied by a label bureau server, which is in turn managed by a rating service. The rating service might own a rating tool to examine particular URLs and create labels describing those URLs. This rating tool can implement a procedure to discover a new site as soon as it goes online, examine the site and store a PICS label at the label bureau server.

Later, when content from this site is requested by the user's client software, the PICS label is requested from the label bureau server by the proxy. If the profile establishes that the values in this PICS label mean that the content is not wanted, then the proxy will send an HTML page back to the client, explaining the reason why the content is not being delivered. It may or may not include instructions on how to bypass the blocking or how to initiate the correction of a faulty rating if necessary. However, in most cases, the values in the label mean the content is acceptable, and the content is fetched from the Web server to the proxy and then forwarded to the user.

If the proxy caching is enabled, during further fetches the user sees no delay, since the proxy server caches only pages that have previously been accepted, and then serves them directly from its cache.

The label bureau server might tell the proxy that this site has not yet been rated, and the profile determines whether the user is sent the content anyway or a `Not Yet Rated` page by the proxy is sent. The reviewer's tool is notified by the label bureau that this site should be rated as soon as possible. Some label bureau servers in the future will fetch the content and run a program to create an interim rating that will be returned, pending the reviewer's more accurate site evaluation.

The diagram shown in Figure 14 on page 44, which we have just explained, is simplified if the Web server is enabled to embed PICS labels in the Web documents it serves (either in the HTTP header or in the HTML header), and the proxy server is configured to accept such labels without the need to contact an external label bureau.

2.9 Reverse Proxy

Reverse proxy is a method of making the proxy server transparent to the client. When a proxy server is configured for reverse proxy, it appears to the client to be the origin server. The client is not aware that the request is actually being sent to another server. Use the `PROXY` directive to configure WTE as a reverse proxy server. The first parameter of the `PROXY` directive specifies the template to be used by the client to access the content on the origin server, and the second parameter identifies the template for the actual URL of the content to be retrieved. A reverse proxy server can be protected using the `PROTECT` directive.

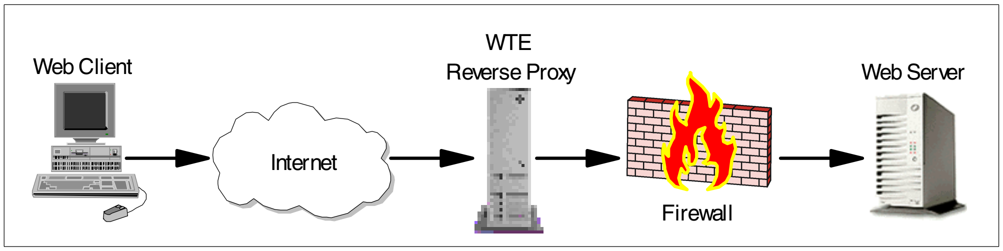

Figure 15. Reverse Proxy Basic Scenario

Suppose a company wants to allow customers to order products over the Internet, for example, using IBM WebSphere Application Server. Connecting IBM WebSphere Application Server directly to the Internet exposes the entire ordering system to potential hacking by outsiders. To eliminate that exposure, this company should put IBM WebSphere Application Server behind a firewall, in order to protect it from hackers. Then, WTE should be placed between the firewall and the Internet, and configured as a reverse proxy server. Only WTE can access the firewall and the customer ordering system is protected from potential hackers by the firewall.

Notice that Figure 15 shows only a basic reverse proxy scenario. In reality IBM WebSphere Performance Pack allows you to implement much more complex architectures, where multiple WTE reverse proxy servers are load-balanced by an ND server and share the same cache using the AFS file system.

2.10 Proxy Activity Monitor

The Proxy Activity Monitor consists of multiple pages that contain information about the activity of the proxy server. It provides summary information and recent entries from the cache and proxy access logs. You can use this data to configure the caching features and to improve your server’s performance.

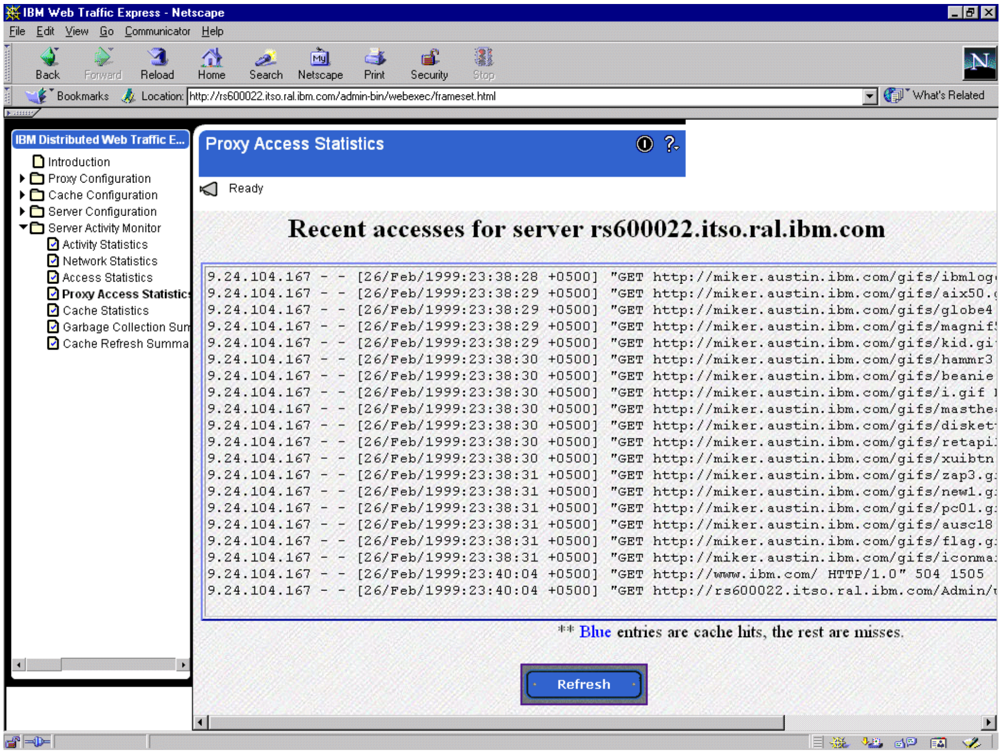

Recent accesses for server rs600022.itso.ral.ibm.com

```
9.24.104.167 - - [26/Feb/1999:23:38:28 +0500] "GET http://miker.austin.ibm.com/gifs/ibmlog
9.24.104.167 - - [26/Feb/1999:23:38:29 +0500] "GET http://miker.austin.ibm.com/gifs/aix50.
9.24.104.167 - - [26/Feb/1999:23:38:29 +0500] "GET http://miker.austin.ibm.com/gifs/globe4
9.24.104.167 - - [26/Feb/1999:23:38:29 +0500] "GET http://miker.austin.ibm.com/gifs/magnif
9.24.104.167 - - [26/Feb/1999:23:38:29 +0500] "GET http://miker.austin.ibm.com/gifs/kid.gi
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/hammr3
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/beanie
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/i.gif
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/masthe
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/disket
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/retapi
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/xuibtn
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/zap3.g
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/new1.g
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/pc01.g
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/ausc18
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/flag.g
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/iconma
9.24.104.167 - - [26/Feb/1999:23:40:04 +0500] "GET http://www.ibm.com/ HTTP/1.0" 504 1505
9.24.104.167 - - [26/Feb/1999:23:40:04 +0500] "GET http://rs600022.itso.ral.ibm.com/Admin/
```

** Blue entries are cache hits, the rest are misses.

Figure 16. Server Activity Monitor - Proxy Access Statistics

2.11 WTE Proxy Server Access Protection

WTE can be configured to protect access to its resources, when it works as a typical Web server. It is also possible to configure it to require user ID and password authentication to all the users that try to access the proxy function. It is possible for you to specify, for example, that only authorized users can access your caching and filtering proxy server, by requiring them to be authenticated via their user ID and password. As another example, you can restrict access to your caching and filtering proxy server, allowing only requests generated by users belonging to specified domains.

2.12 Handling Header Information

WTE allows you to increase client anonymity by configuring the caching and filtering proxy server to strip or modify HTTP header information.

2.13 SSL Tunneling

WTE supports Secure Sockets Layer (SSL) connections. SSL secure connections involve encryption and decryption processes and are established directly between the client browser and the destination Web server. The caching and filtering proxy server does not make any attempt to cache or decrypt the information that the client and the proxy server exchange during an SSL connection, but it establishes a connection to the destination Web server and passes the requests to it without looking at the data.

Chapter 3. WTE Installation

In this chapter, we show you step by step how to perform the installation of IBM Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2, on AIX, Solaris and Windows NT. We also describe how to start, stop and restart WTE on these three platforms, and how to uninstall WTE.

3.1 System Requirements

Before the installation, it is recommended that you check the system hardware and software requirements as described in Table 1 on page 49. It is also recommended that you read the readme.txt file on the installation CD, paying special attention to the section called KNOWN PROBLEMS, RESTRICTIONS, AND CONSIDERATIONS.

Table 1. System Hardware and Software Requirements

	AIX	**Solaris**	**Windows NT**
Hardware Platform	IBM RISC/6000	SUN SPARC or ULTRASPARC Workstation or Server	Intel x86, supported by Windows NT 4.0
Operating System	AIX 4.2.1 with APAR IX73200, AIX 4.3.0 with PTF U452478, or later	Solaris 2.5.1 or 2.6	Windows NT Workstation or Windows NT Server 4.0
Disk Space Requirement	~ 24 MB	~ 32 MB	~ 18 MB NTFS/HPFS
Communication	TCP/IP		
Java[a]	JRE 1.1.6 or later[b]	JDK 1.1.6 or later[c]	JDK 1.1.6 or later[d]
Memory	64 MB or more	64 MB or more	32 MB or more
Paging Space	128 MB or more	128 MB or more	32 MB or more
Caching Requirement	Additional 1 MB per 100 MB of cache		
SNMP Support	SystemView Agent Development Toolkit[e]	SNMP Agent for Solaris[f]	SystemView Agent Development Toolkit[g]

[a] WTE needs Java only for the WebSphere Performance Pack common installation feature.

[b] The Java Runtime Environment (JRE) for AIX can be downloaded from `http://www.ibm.com/java/jdk/download/index.html`.

[c] The Java Development Kit (JDK) for Solaris can be downloaded from `http://www.javasoft.com`.

[d] The Java Development Kit (JDK) for Windows systems can be downloaded from `http://www.javasoft.com`.

[e] Download from `http://www.support.tivoli.com/sva/shasdk.html`.

[f] Download from `http://www.lotus.com`.

[g] Download from `http://www.support.tivoli.com/sva/shasdk.html`.

3.1.1 Browser Requirements on the WTE Configuration Client

To configure WTE V2.0 with the Configuration and Administration Forms, you need a configuration client workstation physically connected to the network. On this workstation, you need to install a browser that:

- Can display frames
- Supports Java Development Kit (JDK) 1.1

- Is enabled for both JavaScript and Java, as shown in the following figure:

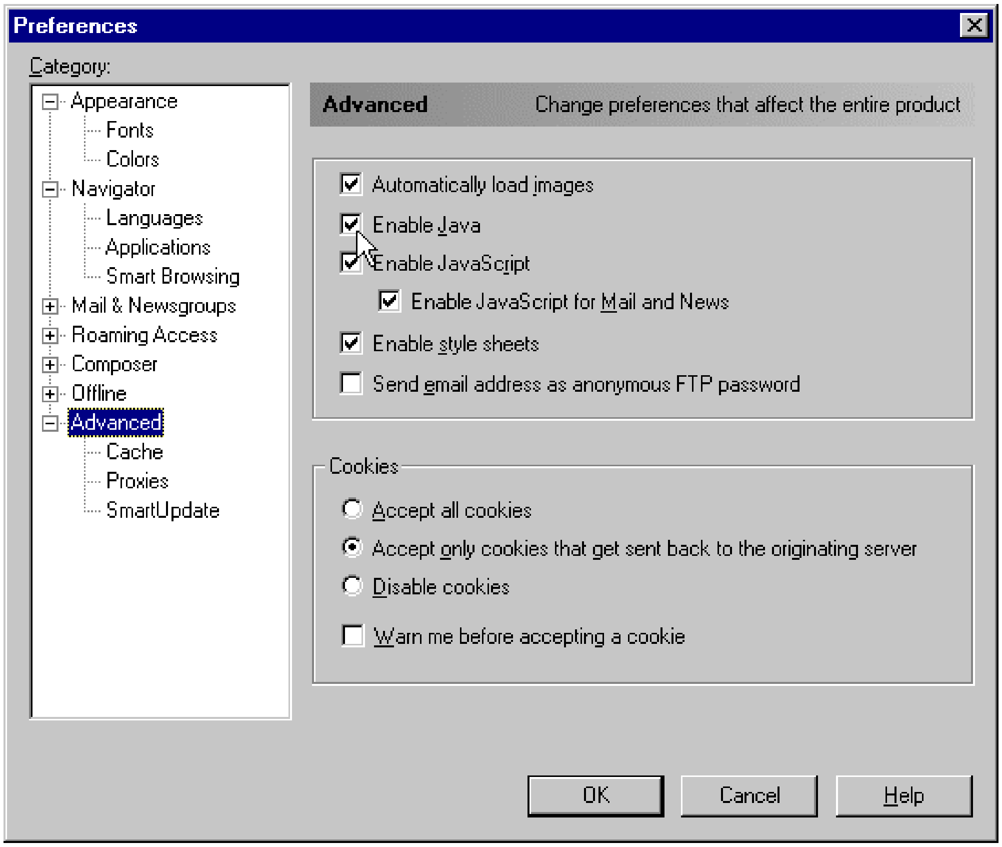

Preferences

Category:

Advanced Change preferences that affect the entire product

☑ Automatically load images

☑ Enable Java

☑ Enable JavaScript

☑ Enable JavaScript for Mail and News

☑ Enable style sheets

□ Send email address as anonymous FTP password

Cookies

○ Accept all cookies

● Accept only cookies that get sent back to the originating server

○ Disable cookies

□ Warn me before accepting a cookie

Figure 17. Enabling Java and JavaScript in the Browser

- Has color resolution set to at least 256 colors (operating system setting)
- Is set to cache documents and compare the cached document with the network document *every time*, as shown in the following figure:

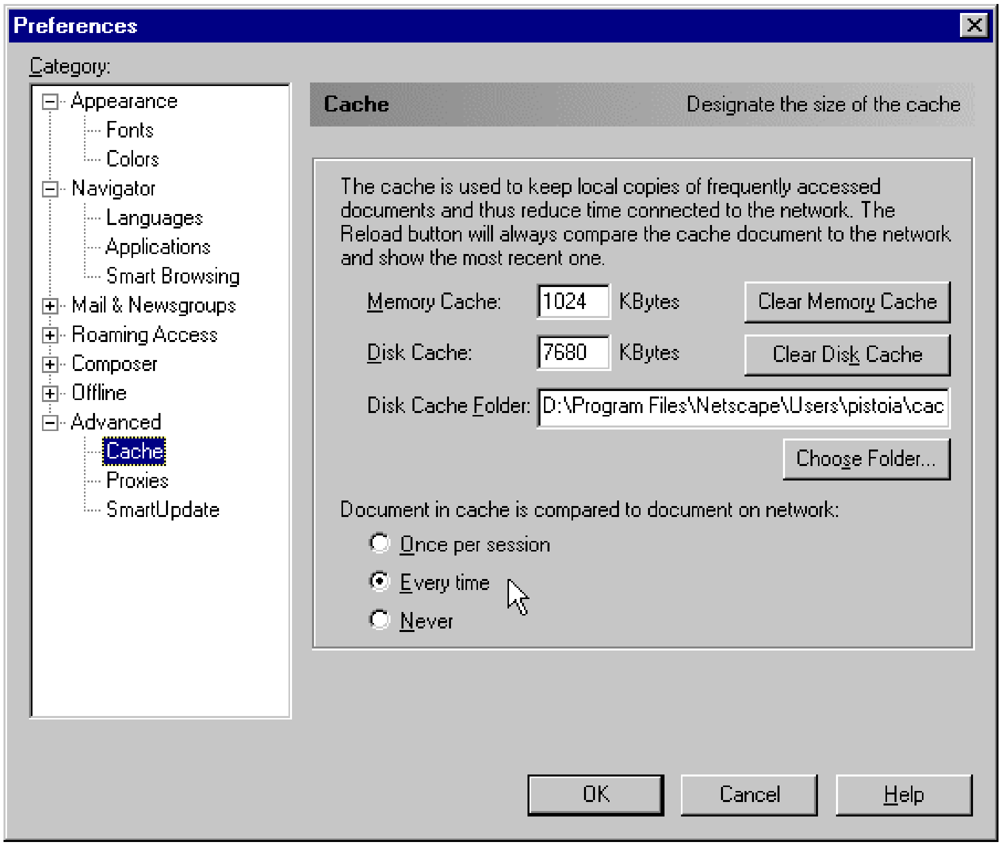

The cache is used to keep local copies of frequently accessed documents and thus reduce time connected to the network. The Reload button will always compare the cache document to the network and show the most recent one.

Memory Cache: 1024 KBytes

Disk Cache: 7680 KBytes

Disk Cache Folder: D:\Program Files\Netscape\Users\pistoia\cac

Document in cache is compared to document on network:

○ Once per session

● Every time

○ Never

Figure 18. Configuring the Browser Cache

Java and JavaScript must be enabled in your browser because in this version, the Configuration and Administration Forms graphical user interface (GUI) makes use of both Java and JavaScript technologies.

3.1.1.1 Recommended Browsers

The WTE V2.0 Configuration and Administration Forms GUI has been successfully tested with the following browsers:

Table 2. Recommended Browsers

Operating System	Minimum Browser Required
Windows NT 4.0	Netscape Navigator V4.04 (with the JDK 1.1 SmartUpdate patch)
	Microsoft Internet Explorer V4.0
Windows 95	Netscape Navigator V4.04 (with the JDK 1.1 SmartUpdate patch)
	Microsoft Internet Explorer V4.0
AIX	Netscape Navigator V4.04
OS/2	Netscape Navigator V2.02 (service level 7) and Java 1.1.2 or higher for OS/2 (only the runtime code is required)

Current versions of Microsoft Internet Explorer and Netscape Navigator can be downloaded from `http://www.microsoft.com/downloads` and `http://www.netscape.com/computing/download/index.html` respectively.

Notice the following:

- **Netscape Navigator on Windows systems**

  Netscape Navigator V4.04 and the JDK 1.1 SmartUpdate patch are two separate installations. Netscape Navigator V4.04 includes the Java Development Kit (JDK) 1.1.2. The JDK 1.1 SmartUpdate patch updates the JDK level to 1.1.4. To confirm that you have the correct level, look at the header in the Netscape Navigator Java Console available from the Communicator menu after clicking **Tools**. The header should say that the version of Java on your system is at least 1.1.4:

  ```
  Netscape Communications Corporation -- Java 1.1.4
  ```

  You can download the JDK 1.1 SmartUpdate patch from `http://developer.netscape.com/software/jdk/download.html`.

- **Windows NT 4.0 display settings**

  Windows NT 4.0 should be configured to run with more than 256 colors to enable GIFs to display properly.

- **Netscape Navigator on the OS/2 platform**

  To confirm that you have Netscape Navigator service level 7, open the Installation utility in the Netscape folder. Select the **Netscape 2.02.00** item and click **Details**, then **Product Status**. Select **Netscape Navigator** then click **Service Level**.

  To confirm that you have Java 1.1.2 or higher, enter

  ```
  java -version
  ```

  on the command line.

  If you have not installed Netscape Navigator, install Java first. When you install Navigator, you will be prompted for the Java level.

  If you have already installed Navigator, install Java and then click **Java Version Selection** in the Netscape folder to select **Java 1.1**.

  To download Navigator and Java, point your browser to `http://service.boulder.ibm.com/asd-bin/doc/en_us/catalog.htm`. If you are installing Java 1.1.4, you need OS/2 Feature Installer Version 1.1 or higher. You can download the latest version of OS/2 Feature Installer from this site.

3.2 Installation on AIX

We performed the installation of WTE on a uniprocessor IBM RS/6000 43P having 192 MB of RAM, 2.2 GB of hard disk and one token-ring interface. This RS/6000 had AIX Version 4.3.1 and Java Runtime Environment (JRE) 1.1.6 pre-installed. We installed WTE in the default location /usr/lpp/WSPP and the cache and logs in a file system, named /wte, that we created.

The /usr and /wte File Systems

Note that we will install the product in the default /usr/lpp/WSPP directory and place the caching and log files in the /wte file system.

The installation program will not increase the /usr file system size automatically during the install. Hence, if space is not sufficient, we have to increase it manually, as explained in the following sequence of instructions:

1. From a command line, enter:

   ```
   df -k /usr
   ```

2. Check column `Free` for space available. Perform the following steps if additional space is required:

   1. From a command line, enter:

      ```
      smitty jfs
      ```

   2. Select **Change / Show Characteristics of a Journaled File System** and then press Enter.
   3. Select **/usr** in the File System Name screen.
   4. Set `SIZE of file system` to the required size and then press Enter.

We have also decided to place the cache root directory and all the log files in the file system /wte. We will create the /wte file system using the AIX `smit` utility as follows:

1. From a command line, enter:

   ```
   smitty crjfs
   ```

2. Select **rootvg** then press Enter.
3. Set `SIZE of filesystem (in 512-byte blocks)` to `1204224`.
4. Set `MOUNT POINT` to `/wte`.
5. Set `Mount Automatically at system restart` to **yes**.
6. Leave the rest as default and press Enter.
7. Mount the /wte filesystem manually using the following command:

   ```
   mount /wte
   ```

Since the WTE server will be caching and logging into the /wte file system, we have to set up the permissions for the /wte file system to enable the WTE to write to this new file system. We would recommend that the permissions for this file system should be read and write by the WTE administrator only.

A Note on Performance

Do note that the reason we are placing the cache and the logs in a single file system is because of the limitation of the single disk drive that we have. However, for performance reasons, it is recommended that cache and cache access logs be placed in separate disk drives.

3.2.1 Installation Preparation

The AIX installation program for each component of IBM WebSphere Performance Pack makes use of Java InstallShield's setup class. For this reason, you will be prompted to install Sun JDK 1.1.6 or higher, if it is not already installed.

Use the following command to determine whether or not Java is installed on your machine:

```
java -version
```

If Java is installed, the command above displays the version of the Java Virtual Machine (JVM). If Java is not installed on your machine or installed at a lower level than 1.1.6, you can find the install image for JDK 1.1.6 for AIX on the IBM Websphere Performance Pack V2 installation CD (in the jdk/aix directory) or at `http://www.ibm.com/java/jdk/download/index.html`. The installation of JDK on AIX is described in the IBM redbook *Network Computing Framework Component Guide,* SG24-2119.

Paging Space

One of the installation prerequisites is that your system should have a minimum of 128 MB of paging space. To check the amount of paging space on your system, enter the following command as a root user on a command line:

```
lsps -a
```

We found that in our system the paging space size was already 192 MB. If you do not have the required paging space you can use the `smit` utility to increase or add a paging space. This operation involves the following steps:

1. From a command line, enter:

   ```
   smitty pgsp
   ```

2. To change the size of an existing paging space, select **Change / Show Characteristics of a Paging Space**
3. To add a new paging space, select **Add Another Paging Space**

3.2.2 Installing WTE on AIX

To start the installation on AIX, follow the steps listed below:

1. Insert the IBM WebSphere Performance Pack Version 2 CD in the CD-ROM drive.
2. From a command line, enter the following commands:

   ```
   mkdir /cdrom
   mount -rv cdrfs /dev/cd0 /cdrom
   cd /cdrom/aix
   ```

   - If you already have JRE installed as described above, start the Java InstallShield installation program, by entering the command:

     ```
     java setup
     ```

   - If you do not have JRE installed as described above, start the installation script install.sh. This script will check the JRE version and install the necessary JRE from the installation CD if necessary. In this case, you should enter the following command:

```
./install.sh
```

Enter the information requested by the installation script. It will ask you if you would want to install the JDK that is shipped with the installation CD-ROM.

The first screen you will see is the Welcome window. After clicking the **Next** button, you are prompted to enter the destination location. Note that this is a working directory for Webshere Performance Pack. The default location is /usr/lpp/WSPP:

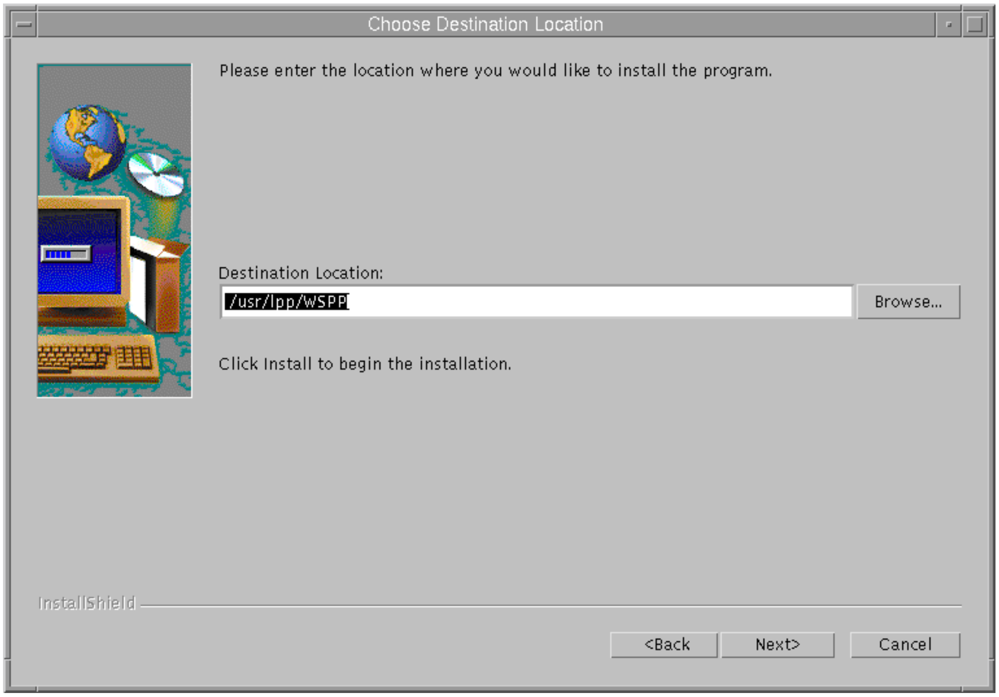

Please enter the location where you would like to install the program.

Destination Location:

/usr/lpp/WSPP

Click Install to begin the installation.

Figure 19. Choose Destination Location

In the above window, you are prompted to click **Install** to begin the installation. Such a button does not exist, due to a known InstallShield for Java problem. Click the button labeled **Next** to continue instead. If the destination directory does not exist on your system, you will be presented with a question dialog asking if you would like the location to be created. Then you will be presented with a window allowing you to choose which IBM WebSphere Performance Pack V2 components you want to install. You can select one or more of the following four options:

- **File Sharing**
- **Load Balancing**
- **Caching and Filtering**
- **Common Configuration**

This is shown in the following figure:

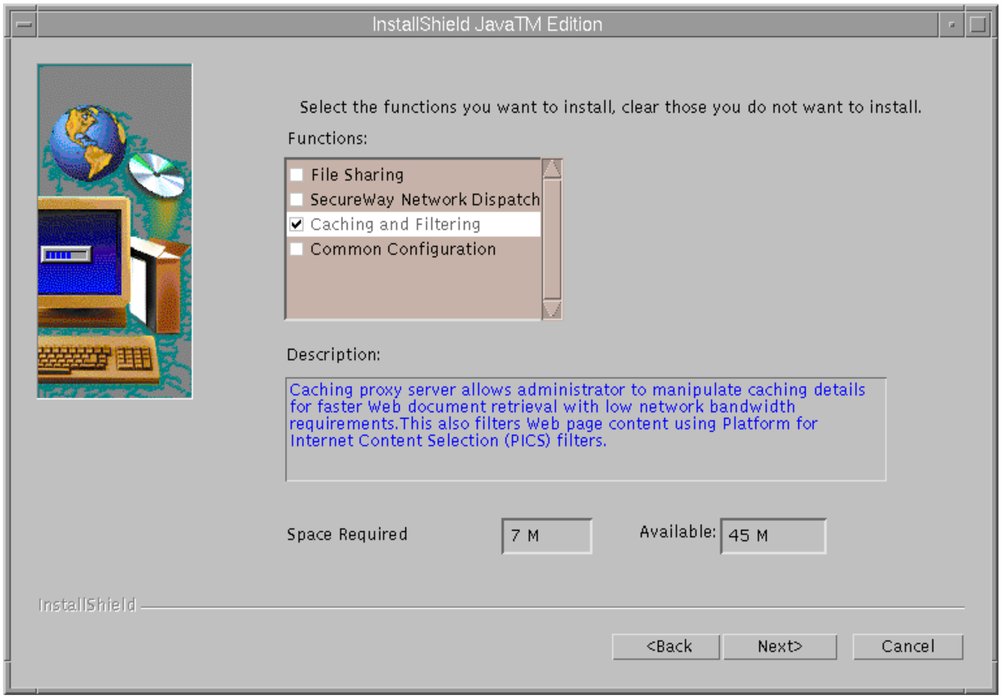

Select the functions you want to install, clear those you do not want to install.

Functions:

- □ File Sharing
- □ SecureWay Network Dispatch
- ☑ Caching and Filtering
- □ Common Configuration

Description:

Caching proxy server allows administrator to manipulate caching details for faster Web document retrieval with low network bandwidth requirements.This also filters Web page content using Platform for Internet Content Selection (PICS) filters.

Space Required 7 M

Available: 45 M

Figure 20. Choose the Components to Install

When we selected **Caching and Filtering**, a description of that component was displayed as well as the space required to install it. This window also shows you how much space is available in the destination location file system. If you do not have enough free space in your destination location file system and you click **Next**, you will see a warning window letting you know that there is not enough space to perform the installation.

Notice that on AIX, WTE comes as a single package, with no subcomponents.

Common Configuration and WTE Simultaneous Installation

You may also want to install the Common Configuration option on a different machine to use it for PICS filtering configuration.

Common Configuration is a Web-based configuration tool for IBM WebSphere Performance Pack that requires the simultaneous presence of a Web server to work properly. A selection of Web servers is allowed. All of these Web servers, by default, listen on port 80, which is also the default port that the WTE server listens on. Hence, if the Common Configuration is to be installed on the same machine as the WTE servers, it is recommended that you change the Common Configuration’s Web server to listen on a port other than 80.

After you select **Next**, you will get a window that asks if you want the installation program to enable caching. if you select this option, you are required to enter the caching parameters. We click **Caching** to enable caching. Then, we enter the Cache Root Directory as `/wte/cache` and Cache Access Log file as `/wte/logs/http-cache`. Also, we set the Cache Size to `500 MB` as shown in the following figure:

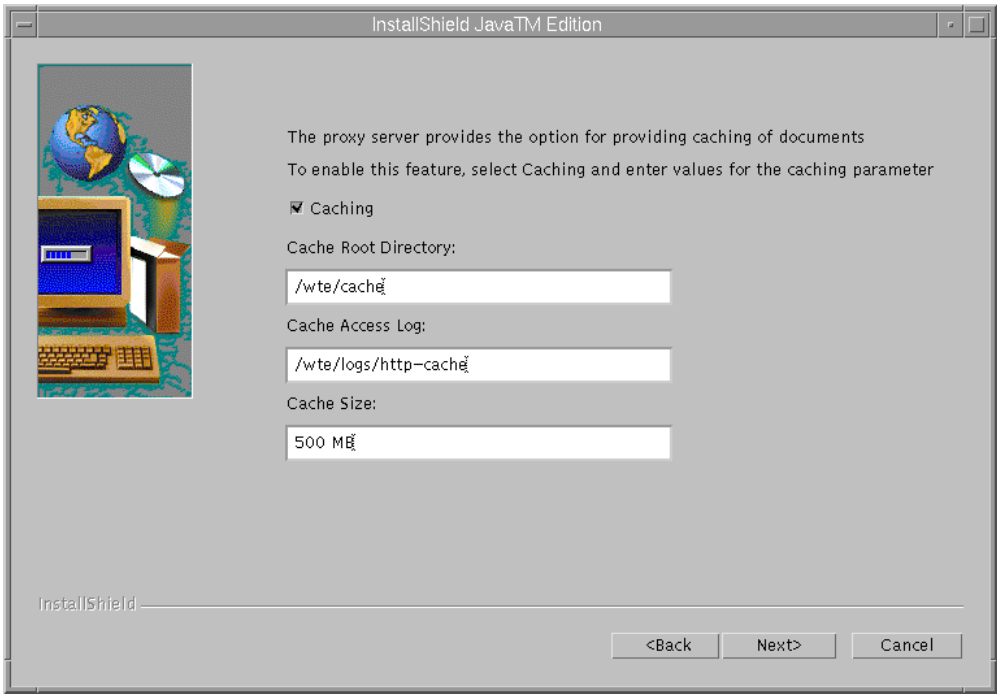

The proxy server provides the option for providing caching of documents

To enable this feature, select Caching and enter values for the caching parameter

☑ Caching

Cache Root Directory: /wte/cache

Cache Access Log: /wte/logs/http-cache

Cache Size: 500 MB

Figure 21. Selecting Caching and Entering Caching Parameters

Notice that the directories /wte/cache and /wte/logs need to be explicitly created. The WebSphere Performance Pack common installation procedure does not create any directories in the wte file system for you.

After you select **Next**, you will get a window that asks if you want the installation program to replace other programs already installed on your system. We select **No** in this case, even though we have not installed any programs on our system yet.

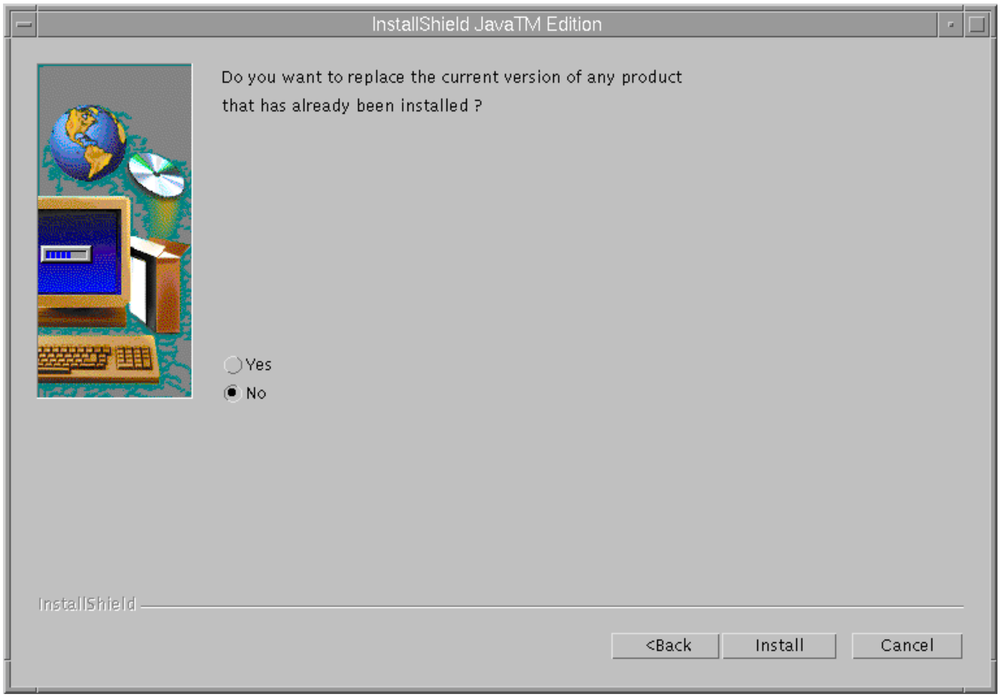

Do you want to replace the current version of any product that has already been installed ?

○ Yes
● No

Figure 22. Choose to Replace Version

Click the **Install** button and after a while the Installation Complete screen will be displayed:

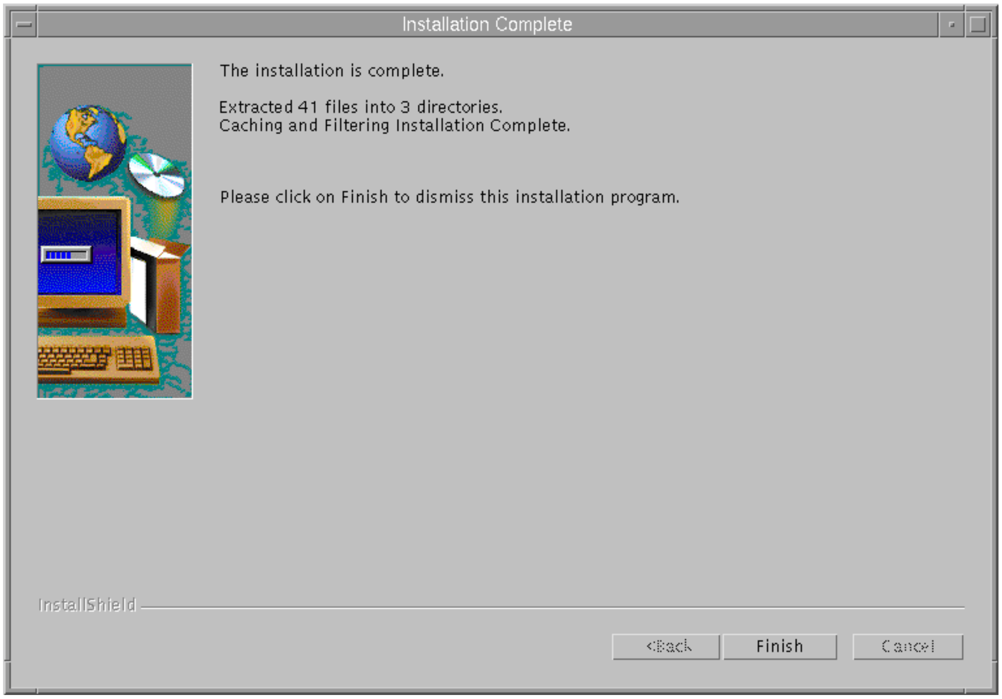

The installation is complete.

Extracted 41 files into 3 directories.
Caching and Filtering Installation Complete.

Please click on Finish to dismiss this installation program.

Figure 23. Installation is Complete

After you finish the installation, WTE starts automatically with the default configuration settings. Notice that the filesets that are installed when the Caching and Filtering component is selected are:

- internet_server.base.admin
- internet_server.base.doc
- internet_server.base.httpd
- internet_server.msg.en_US.httpd
- internet_server.proxy.admin
- internet_server.proxy.docs
- internet_server.proxy.exe
- intnd.nd.driver

After clicking **Finish**, we verified that the product had been installed by entering the following command, which verifies that the httpd daemon is running:

```
lssrc -s httpd
```

WTE is a caching proxy server. It can handle Web requests from multiple clients, forward such requests to remote Web servers, cache the retrieved contents if these are cacheable and serve them to the clients. Subsequent requests for the same contents will be served directly from the cache. WTE uses a process called httpd daemon for such functions. If the product has been correctly installed, the httpd daemon should run at this point, and the above command should produce an output similar to the following figure:

```
Subsystem         Group            PID     Status
 httpd            tcpip            12852   active
```

Late-Breaking Information

After the installation, it is recommended that you read the read.me file for information about installing and using the WTE. This information also corrects and supplements the information in the *User’s Guide* and the *Webmaster’s Guide*. The read.me file is /usr/lpp/internet_server.proxy/read.me.

3.3 Installation on Solaris

We performed the installation of WTE on a uniprocessor SPARC station-20 having 64 MB of RAM, 2.0 GB of hard disk and one token-ring interface. This SPARC station-20 had Solaris 2.6 and JRE 1.1.6 pre-installed. We installed WTE in the default location /opt/WSPP.

The /opt/WSPP Directory and /wte File System

Note that we will install the product in the default /opt/WSPP directory and place the caching and log files in /wte file system. The installation program will not increase the /usr file system size automatically during the install. Hence, if space is not sufficient, we have to increase it manually. Refer to “The /usr and /wte File Systems” on page 53 for more details.

Notice also that the directories for the caching and log files must be explicitly created. The WebSphere Performance Pack common installation routine does not create any directories in the /wte file system for you.

Since the WTE server will be caching and logging into the /wte file system, we have to set up the permissions for the /wte file system to enable WTE to write to this new file system. We would recommend that the permissions for this file system should be read and write by the WTE administrator only.

3.3.1 Installation Preparation

The Solaris installation program for each component of IBM WebSphere Performance Pack is very similar to AIX. It makes use of Java InstallShield’s setup class. You will be prompted to install SUN JDK 1.1.6 or higher if it is not already installed.

The CD-ROM of IBM WebSphere Performance Pack V2 ships with the JDK 1.1.6 for Solaris, found in the directory jdk/sun. However, we chose to pre-install the JDK 1.1.7 for Solaris, since this was the latest non-beta version that was available at the time we wrote this redbook.

Use the following command to determine whether or not Java is installed on your machine:

```
java -version
```

If Java is installed, the command above displays the version of Java. If Java is not installed on your machine or installed at a lower level than 1.1.6, you can find the install image for JDK 1.1.6 for Solaris on the IBM Websphere Performance Pack V2.0 installation CD-ROM (in the jdk/sun directory) or at `http://www.sun.com/solaris/jdk/download.1.1.7/`. For information on installing the JDK on Solaris, read the README.sparc file that comes with the product.

Paging Space

One of the installation prerequisites is that your system should have a minimum of 128 MB of paging space. To check the amount of paging space on your system, enter the following as a root user on a command line:

```
swap -s
```

We found that in our system the paging space size was already 235 MB.

3.3.2 Installing WTE on Solaris

Installating WTE on Solaris is very similar to AIX. All screens are exactly the same, except the screen where you choose the Destination Location to install the software. In Solaris, the default is /opt/WSPP, where we chose to install. This screen is as shown below:

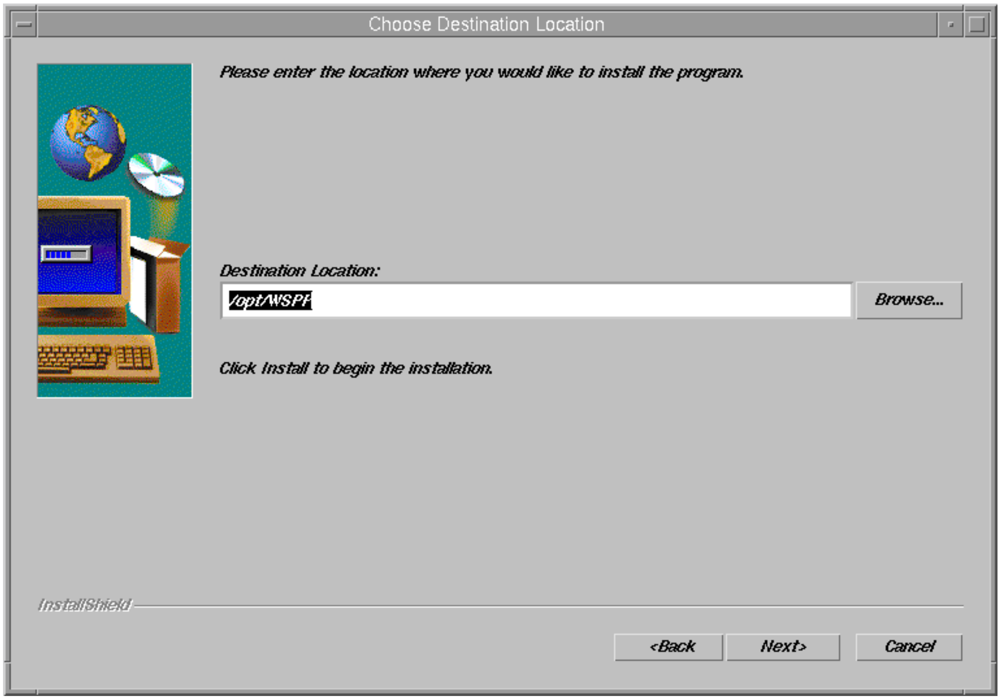
Choose Destination Location

Please enter the location where you would like to install the program.

Destination Location: /opt/WSPP

Click Install to begin the installation.

Figure 24. Choose Destination Location on Solaris

Late-Breaking Information

After the installation, it is recommended that you read the read.me file for information about installing and using WTE. This information also corrects and supplements the information in the *User's Guide* and the *Webmaster's Guide*. The read.me file is /opt/IBMWTE/usr/internet/server_root/read.me.

3.4 Installation on Windows NT

In this section we describe all the steps that are necessary to install the WTE on the Windows NT platform.

First of all, we describe the hardware and software environment on which we performed this installation. The machine was an IBM PC 750 with 166 MHz of CPU, 96 MB of RAM, 1.5 GB of hard disk and one token-ring adapter. This machine had been previously installed with Windows NT Server 4.0 and Service Pack 4 had been applied.

A Note on Performance

Do note that we are placing the cache and the logs in a single partition because of the limitation of the single disk drive that we have. However, for performance reasons, it is recommended that cache and cache access logs be placed in separate disk drives.

Since the WTE server will be caching and logging into the designated partition, we have to set up the permissions for this partition to enable the WTE to write to this new file system. We would recommend that the permissions for this file system should be read and write by the WTE administrator only.

The Windows NT installation program for each component of the IBM WebSphere Performance Pack makes use of Java InstallShield's setup class. You will be prompted to install SUN JDK 1.1.6 or a higher version if it is not already installed. The CD-ROM of IBM WebSphere Performance Pack V2.0 ships with the JDK 1.1.6 for Windows NT, found in the directory jdk\nt. We chose to pre-install JDK 1.1.6 for Windows NT.

Use the following command to determine which Java version is installed on your machine:

```
java -version
```

If Java is installed on your machine, the command above displays the version of the installation code. If Java is not installed, you can install it from the CD-ROM of IBM WebSphere Performance Pack. Alternatively, the latest version of the JDK for Windows systems can be downloaded for free from the JavaSoft Web site `http://www.javasoft.com/products/jdk/1.1/download-jdk-windows.html`. The JDK 1.1.6 installation for Windows NT is very easy and so we skip its description. However, for a detailed description, you can see the IBM redbook *Internet Security in the Network Computing Framework*, SG24-5220.

3.4.1 Installing WTE on Windows NT

To start the installation on NT, follow the steps listed below:

1. Insert the IBM WebSphere Performance Pack Version 2 CD in the CD-ROM drive.
2. There are two ways you can perform the installation on NT:
   - Change to the NT subdirectory on the drive containing the CD-ROM. Then enter:

     ```
     java setup
     ```

   - You can also install from the Start menu. Select **Run...**, then click on **Browse...** and open the Setup.exe program located in the NT directory of the CD-ROM drive, as shown in the following figure:

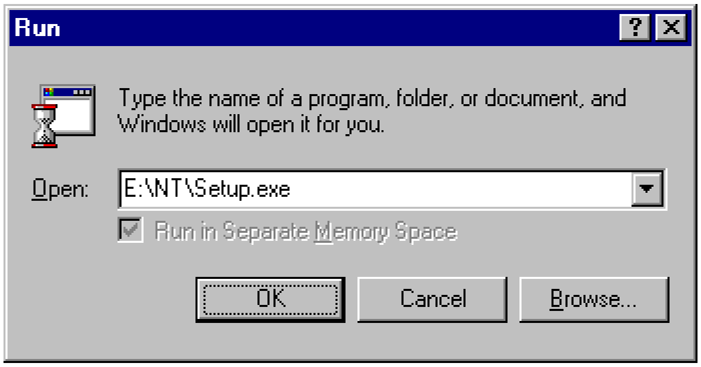

Run
?
Type the name of a program, folder, or document, and Windows will open it for you.
Open:
E:\NT\Setup.exe
☑ Run in Separate Memory Space
OK
Cancel
Browse...

Figure 25. Installation Program's Name and Path

Then press **OK**. Notice that E in our case was the letter assigned to the CD-ROM drive.

Error Message?

If you see an error message saying:

```
Can't find class setup
```

when you execute the command:

```
java setup
```

or nothing happens when executing from the Start menu, the reason can be either one of the following:

- The setup.class file was not found because you are in the wrong directory. The correct directory is the NT directory under the CD-ROM drive.
- The CLASSPATH was not set correctly. Correct its value as explained in the README file associated with the JDK you installed on your system.

The first screen you will see is the Welcome window. After clicking the **Next** button, you will be able to view the README Information screen:

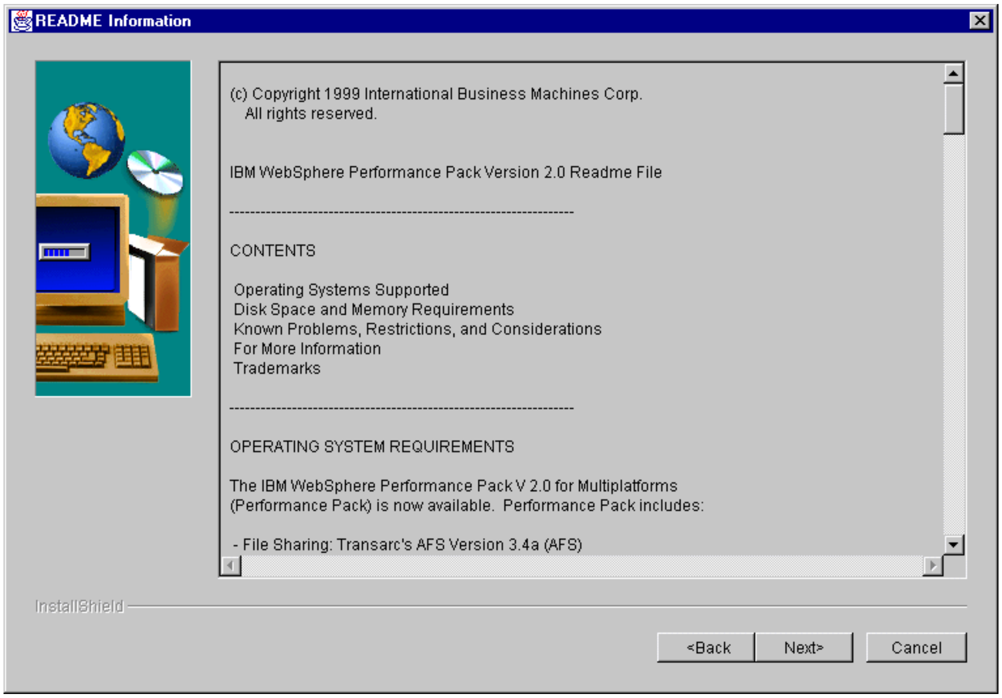

README Information

(c) Copyright 1999 International Business Machines Corp.
All rights reserved.

IBM WebSphere Performance Pack Version 2.0 Readme File

CONTENTS

Operating Systems Supported
Disk Space and Memory Requirements
Known Problems, Restrictions, and Considerations
For More Information
Trademarks

OPERATING SYSTEM REQUIREMENTS

The IBM WebSphere Performance Pack V 2.0 for Multiplatforms
(Performance Pack) is now available. Performance Pack includes:

- File Sharing: Transarc's AFS Version 3.4a (AFS)

Figure 26. README Information

You should view the README as it contains a very useful section called KNOWN PROBLEMS, RESTRICTIONS, AND CONSIDERATIONS. This section contains late-breaking news of some known problems. After viewing the README information screen and clicking **Next**, you will be prompted to enter the destination location as shown in the following screen:

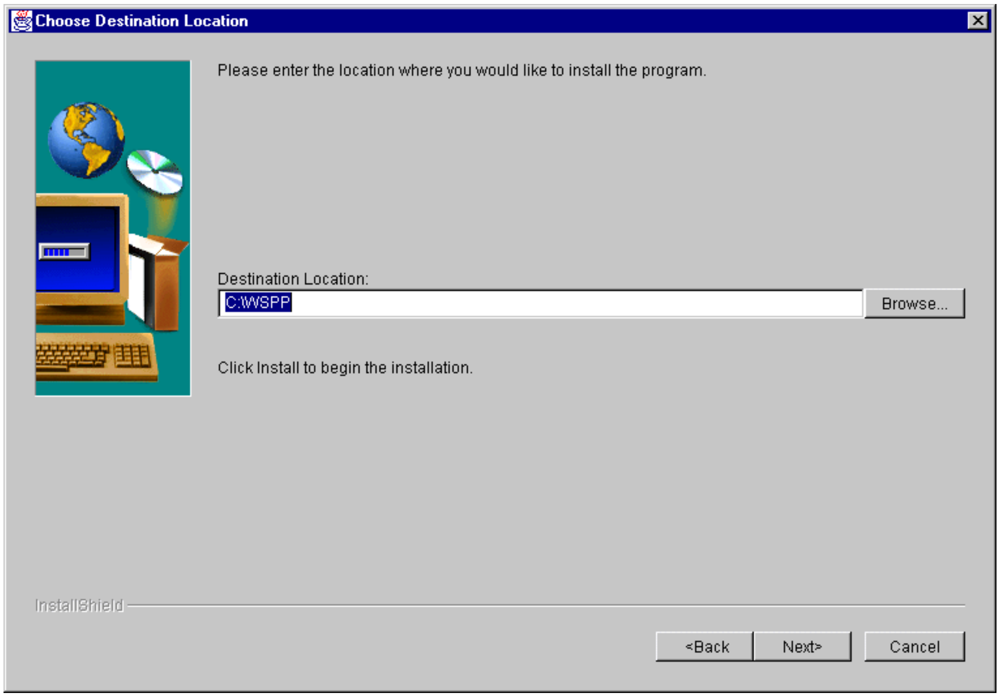

Choose Destination Location

Please enter the location where you would like to install the program.

Destination Location: C:WSPP

Click Install to begin the installation.

Figure 27. Choose Destination Location

Note that this is a working directory for WebSphere Performance Pack. The default location is C:\WSPP. You may click **Browse** to change the destination location. We chose to perform the installation in the default directory, C:\WSPP.

In the above window, you are prompted to click **Install** to begin the installation. Such a button does not exist, due to a known InstallShield for Java problem. Click the button labeled **Next** to continue instead.

If the destination directory does not exist on your system, you will be presented with a question dialog asking if you would like the location to be created.

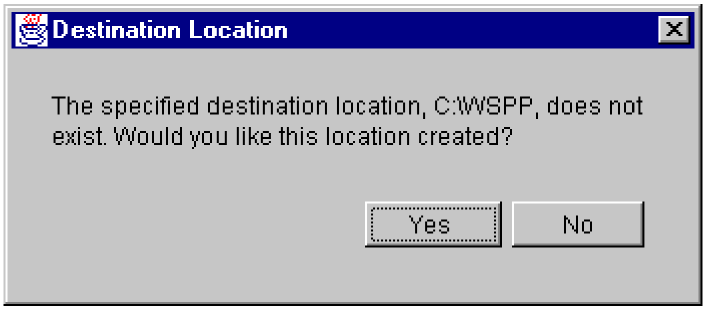

Destination Location

The specified destination location, C:\WSPP, does not exist. Would you like this location created?

Figure 28. Destination Location

In our case, the directory C:\WSPP did not exist yet, so we selected **Yes** to create it. Then you will be presented with a window allowing you to choose which IBM WebSphere Performance Pack V2 components you want to install:

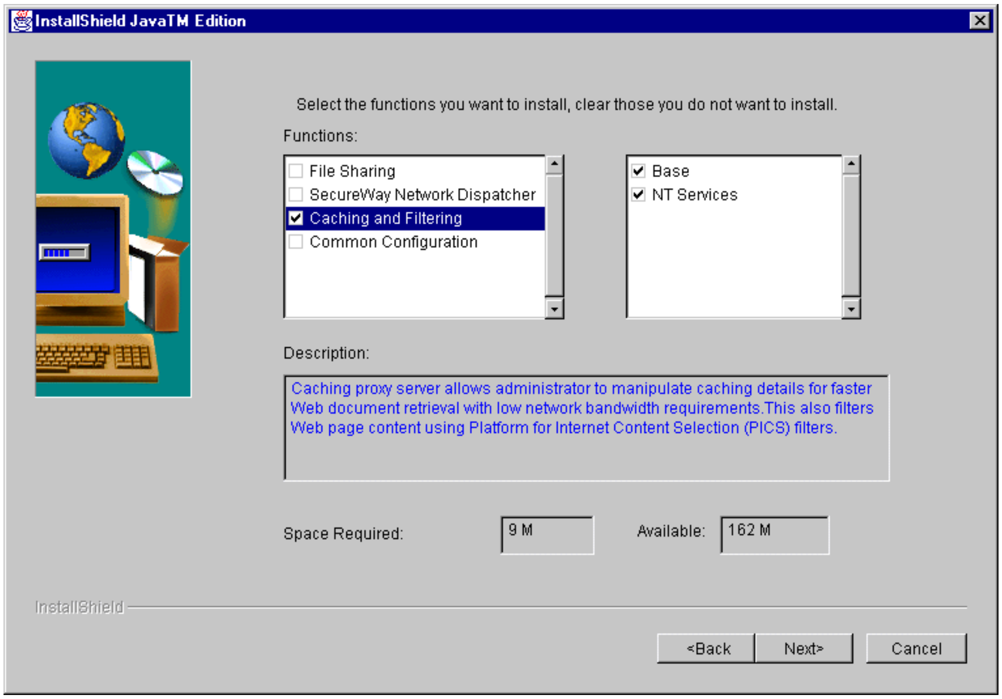

InstallShield JavaTM Edition

Select the functions you want to install, clear those you do not want to install.

Functions:

- □ File Sharing
- □ SecureWay Network Dispatcher
- ☑ Caching and Filtering
- □ Common Configuration

- ☑ Base
- ☑ NT Services

Description:

Caching proxy server allows administrator to manipulate caching details for faster Web document retrieval with low network bandwidth requirements.This also filters Web page content using Platform for Internet Content Selection (PICS) filters.

Space Required: 9 M

Available: 162 M

Figure 29. Selecting the WebSphere Performance Pack Components

We selected **Caching and Filtering**, with both **Base** and **NT Services** options.

The former is the WTE software itself. The latter is an additional set of files that allow the Caching and Filtering component of IBM WebSphere Performance Pack to run as a Windows NT service rather than an application.

- Installing only **Base** has the following consequences:
  - No services are added to Services panel, accessible from the Control Panel.
  - WTE can be started:
    - From the IBM Web Traffic Express 2.0 program group, by selecting **IBM Web Traffic Express**
    - From the command prompt, by entering the `whttpg` command
- Installing **Base** and **NT Services** together has the following consequences:
  - WTE can be activated:
    - As a service from the Services panel, accessible from the Control Panel
    - From a command prompt, by entering `whttpg`
  - The IBM Web Traffic Express 2.0 program group will contain only the Uninstall IBM Web Traffic Express and ReadMe icons, as shown in the following screen:

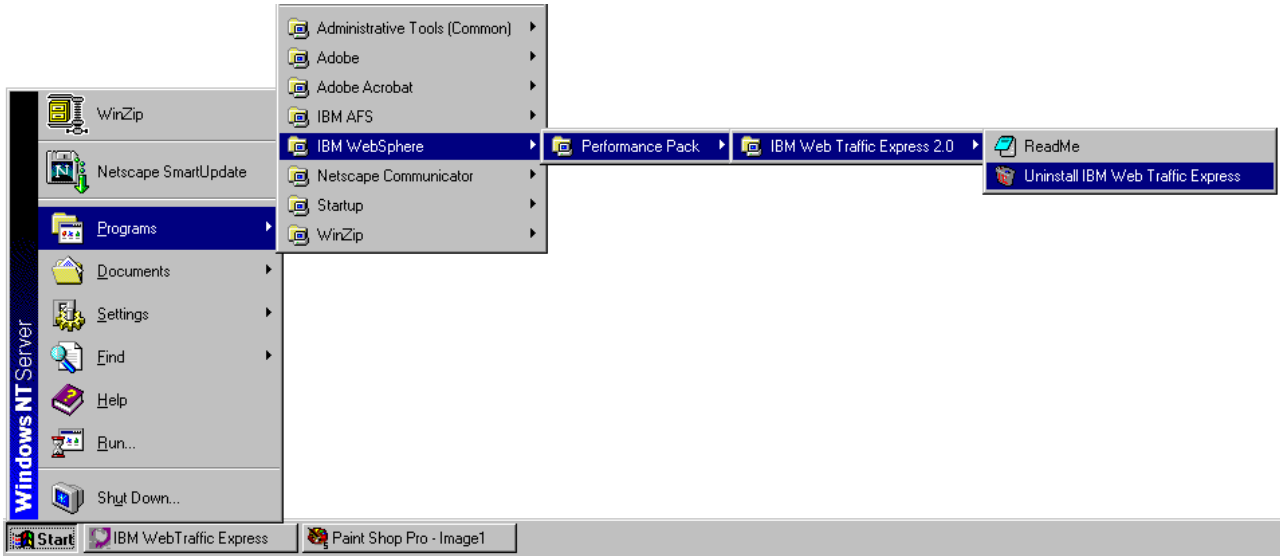

Figure 30. IBM Web Traffic Express 2.0 Program Group when NT Service Is Installed

- You will be allowed to select the **NT Services** option only. However, if you do not have the **Base** installed previously, it will prompt you to install IBM Web Traffic Express before continuing, as shown in the following screen:

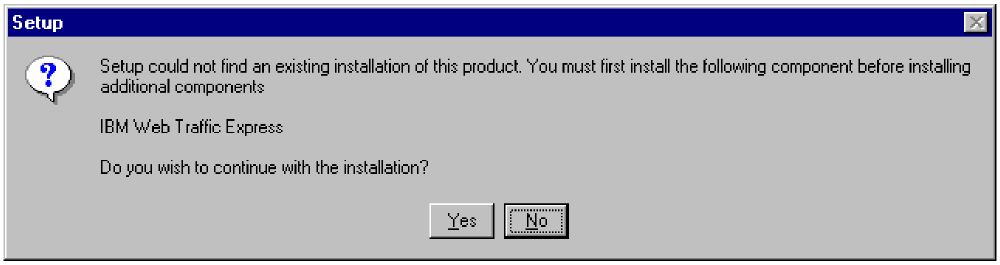
Setup

Setup could not find an existing installation of this product. You must first install the following component before installing additional components

IBM Web Traffic Express

Do you wish to continue with the installation?

Figure 31. Installation of Web Traffic Express Base Component Required

Notice that the product documentation is automatically installed. On our platform, it went in the directory C:\WSPP\WWW\DOCS.

Common Configuration and WTE Simultaneous Installation

You may also want to install the Common Configuration option on a different machine to use it for PICS Filtering configuration.

Common Configuration is a Web-based configuration tool for IBM WebSphere Performance Pack that requires the simultaneous presence of a Web server to work properly. A selection of Web servers is allowed. All of these Web servers, by default, listen on port 80, which is also the default port that the WTE server listens on. Hence, if the Common Configuration is to be installed on the same machine as the WTE servers, it is recommended that you change the Common Configuration's Web server to listen on a port other than 80.

After you click **Next**, you will get another screen that asks you if you want the installation program to replace other programs already installed on your system:

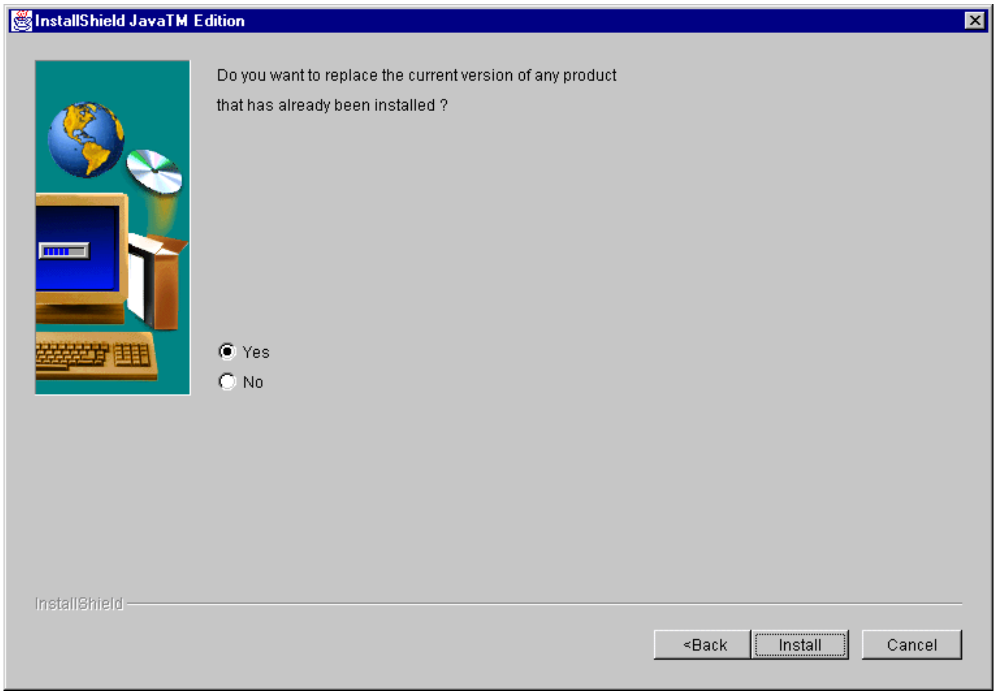

Do you want to replace the current version of any product that has already been installed ?

● Yes
○ No

Figure 32. Replace Current Version

Since we have not installed any other IBM WebSphere Performance Pack component on that particular Windows NT machine, we accepted the deafult option, **Yes**, and clicked **Install.** The installation routine started extracting the installation files for WTE onto the disk.

In the middle of the process, a WTE panel, as shown in Figure 33 on page 68, will be displayed momentarily:

Figure 33. IBM Web Traffic Express Icon

This will be followed by a Reboot Message window, as shown in Figure 34, which asks you to select **OK** to reboot. However, we found that this window appears only momentarily, wiithout allowing the user to select **OK** to reboot.

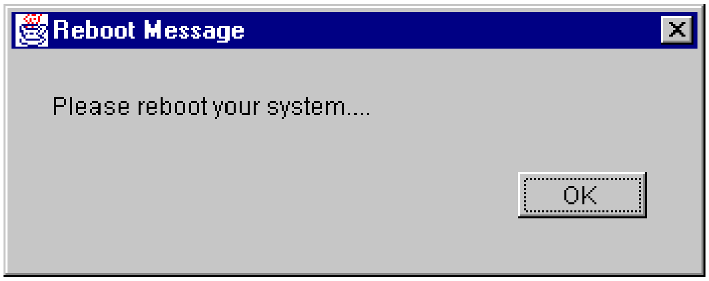

Please reboot your system....

Figure 34. Reboot Message

The final window will inform you that the installation is complete and suggest that you reboot the system, after clicking **Finish** to dismiss the installation program:

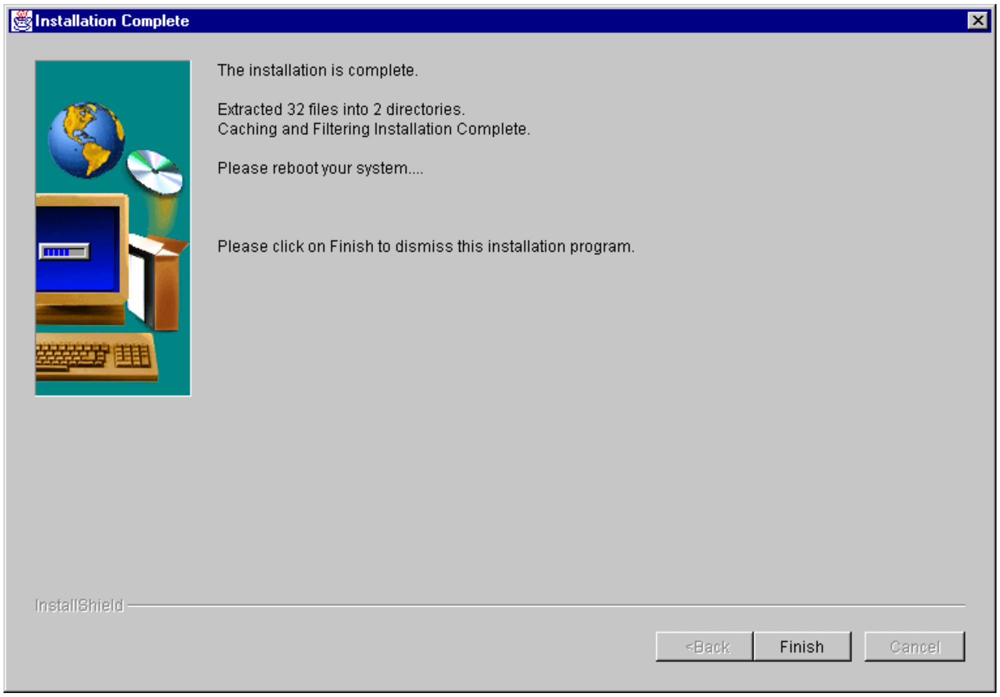

Installation Complete

The installation is complete.

Extracted 32 files into 2 directories.
Caching and Filtering Installation Complete.

Please reboot your system....

Please click on Finish to dismiss this installation program.

Figure 35. Instllation Complete

As recommended by the window above, we restarted the Windows NT Server machine where we had performed the installation. After the reboot, we checked the services available on our computer. To do this, we opened the Services dialog box of the Control Panel folder. We noticed that the IBM Web Traffic Express 2.0 service was now in the list, even though it was not started yet, as shown in the following window:

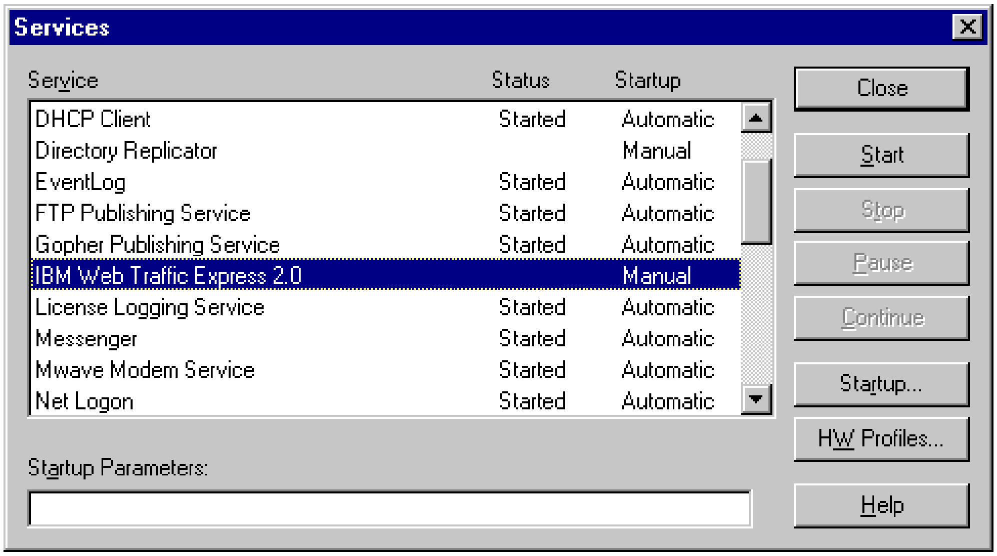

Services

Service	Status	Startup
DHCP Client	Started	Automatic
Directory Replicator		Manual
EventLog	Started	Automatic
FTP Publishing Service	Started	Automatic
Gopher Publishing Service	Started	Automatic
IBM Web Traffic Express 2.0		Manual
License Logging Service	Started	Automatic
Messenger	Started	Automatic
Mwave Modem Service	Started	Automatic
Net Logon	Started	Automatic

Startup Parameters:

Figure 36. IBM WebTraffic Express 2.0 Service

To see how to start and stop WTE on Windows NT systems, refer to 3.5.3, “Starting and Stopping WTE on Windows NT” on page 72.

3.5 Starting and Stopping WTE

The WTE will need to be started on system reboot, either manually or automatically. It will also need to be restarted or refreshed when certain changes are made to activate those changes.

3.5.1 Starting and Stopping WTE on AIX

After the installation, WTE starts automatically with the default configuration settings defined in the /etc/ibmproxy.conf configuration file.

Each time you reboot the system, the /etc/rc.httpd script is run from the /etc/inittab initialization file and the httpd daemon is started automatically under the System Resource Controller (SRC).

In AIX, the SRC is a subsystem controller that facilitates the management and control of complex subsystems. A *subsystem* is any program or process or set of programs or processes that is usually capable of operating independently or with a controlling system. A subsystem is designed as a unit to provide a designated function.

You can start multiple instances of the server, but each instance must listen on a separate port. For a server to start using a different configuration file, you can edit the /etc/rc.httpd script and change the httpd command. To start the server using a different proxy configuration file, say /new/newproxy.conf, instead of the default /etc/ibmproxy.conf, add:

```
-r /new/newproxy.conf
```

to the `start` command in the /etc/rc.httpd file, as shown below:

```
# Start up httpd (HTTP) daemon as a regular server
start /usr/sbin/httpd "$src_running" -r /new/newproxy.conf
```

To stop the server from using the default configuration file, log in as root and from a command line, enter the command:

```
stopsrc -s httpd
```

You should get output similar to the following:

```
0513-044 The stop of the /usr/sbin/httpd Subsystem was completed successfully.
```

To start the server again, log in as root and enter the following command:

```
startsrc -s httpd
```

In this case the output will be similar to the following:

```
0513-059 The httpd Susbsystem has been started. Subsystem PID is 12302.
```

Note

Note that the command `startsrc` will always start the WTE server using the default configuration file. If you have changed the default proxy configuration file to a new file, for example, /new/newproxy.conf, execute the following as a root user:

```
stopsrc -s httpd
httpd -r /new/newproxy.conf
```

After modifying the configuration, you can restart the server without stopping it first, unless you changed one of the *directives not changed on restart* (for more information, see 4.3.3, “Directives not Changed on Restart” on page 89). After editing the configuration file, you can restart the server using the command:

```
refresh -s httpd
```

3.5.1.1 Debugging the WTE Server

We can start the WTE Server in the debugging mode for problem determination purposes using the following command:

```
httpd -r /etc/ibmproxy.conf -vv -debug
```

3.5.2 Starting and Stopping WTE on Solaris

After the installation, WTE starts automatically with the default configuration settings defined in the /etc/ibmproxy.conf configuration file.

Each time you reboot your system, /etc/rc2.d/s88go_httpd is run from the initd script. This script starts the WTE server automatically. By default, the WTE server uses the current configuration settings as defined in the /etc/ibmproxy.conf configuration file. If you want the server to start with a different configuration file or different run-time settings, you can edit /etc/rc2.d/s88go_httpd and change the httpd command.

You can start multiple instances of the server, but each instance must listen on a separate port. To start the server using a different proxy configuration file, for example /new/newproxy.conf, instead of the default /etc/ibmproxy.conf, enter the following command from a command prompt:

```
httpd -r /new/newproxy.conf
```

If you do not want the server to start automatically every time you start or reboot your system, remove the httpd start command and its two kill scripts:

```
Delete s88go_httpd from the /etc/rc2.d directory
Delete k54go_httpd from the /etc/rc0.d directory
Delete k54go_httpd from the /etc/rc1.d directory
```

To start the server with the default run-time settings, enter the command `httpd`.

To start the server as a foreground process, enter the command:

```
httpd -nobg
```

To stop the server, you must either be root, or the owner of the process. Use the `ps` command to find the httpd process ID:

```
ps -ef | grep "httpd"
```

Use the kill command to kill the httpd process.

After modifying your configuration, you can restart the server without stopping it first, unless you changed one of the *directives not changed on restart* (for more information on this, see 4.3.3, “Directives not Changed on Restart” on page 89). After editing the configuration file, you can restart the server using the command:

```
httpd -restart
```

3.5.3 Starting and Stopping WTE on Windows NT

After you finish installing WTE on Windows NT, the IBM Web Traffic Express 2.0 process will not start automatically (see Figure 36 on page 69).

The way you start and stop WTE on Windows NT may depend on the type of installation you have performed.

3.5.3.1 Starting WTE from the Command Prompt

The WTE Server can always be started from the command prompt by entering the `whttpg` command. This command starts the server with the default configuration file \WINNT\ibmproxy.conf.

The option of using the command line to start WTE applies whether you installed WTE by selecting only the **Base** option, or both the **Base** and **NT Services** options (see Figure 29 on page 65). Notice that the whttpg.exe executable file comes by default in the directory C:\WSPP\WWW\BIN.

WTE will pick up the ibmproxy.conf and socks.cnf from the Windows NT installation directory, which by default is C:\WINNT. You can use different proxy and SOCKS configuration files (for example, C:\new\newproxy.conf and C:\new\newsocks.cnf) or start up multiple instances of WTE by using the `-r` option of the `whttpg` command, as follows:

```
C:\WSPP\WWW\BIN\whttpg -r C:\new\newproxy.conf -r C:\new\newsocks.cnf
```

Running the `whttpg` command will bring up the WTE GUI, shown in the following figure:

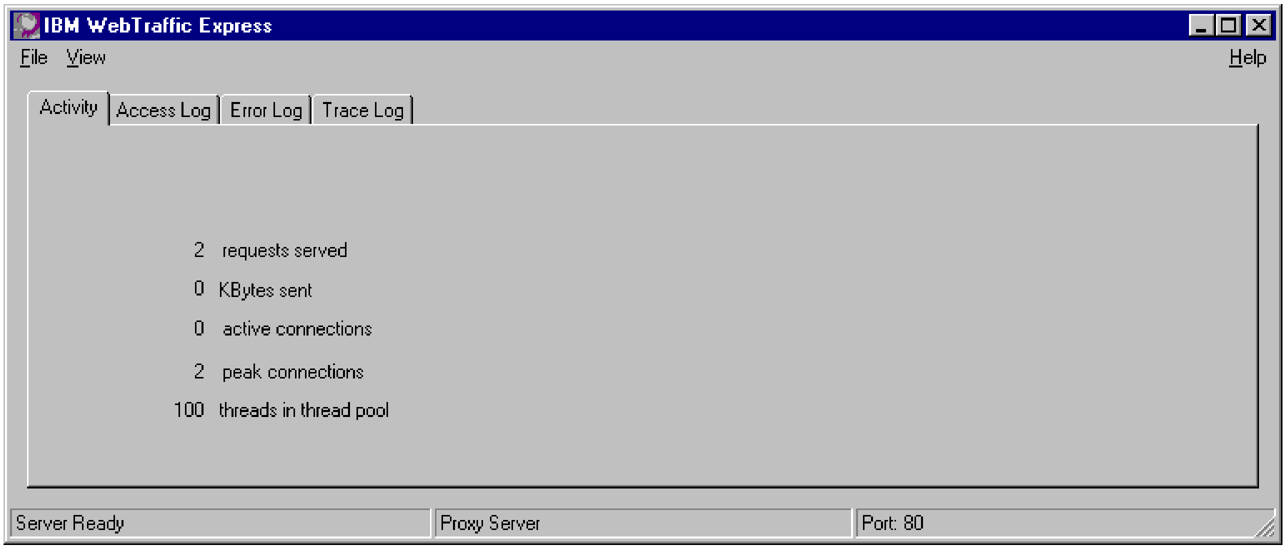

2 requests served
0 KBytes sent
0 active connections
2 peak connections
100 threads in thread pool

Server Ready | Proxy Server | Port: 80

Figure 37. IBM Web Traffic Express - whttpg

The WTE GUI is very useful for monitoring the status of your WTE server. You can also use it to stop and restart WTE:

- You can stop WTE by selecting **Exit** from the File menu.
- You can restart WTE by selecting **Restart** from the File menu.

3.5.3.2 Starting and Stopping WTE from the Services Panel

If you have installed WTE by selecting both the **Base** and **NT Services** options (see Figure 29 on page 65), WTE will run on your system as a Windows NT service, and you can start and stop it from the Services panel, by highlighting **IBM Web Traffic Express 2.0** and pressing the **Start** button (see Figure 36 on page 69). When the service has started, the **Stop** button becomes available and you can click on it to stop WTE.

You can also set the WTE Startup Type to **Automatic**, so that the IBM Web Traffic Express 2.0 process will start automatically every time you reboot the system. To do this, you should press the **Startup...** button in the Services panel when **IBM Web Traffic Express 2.0** is highlighted (see again Figure 36 on page 69). The IBM Web Traffic Express Service window will be brought up. Set the Startup Type to **Automatic** and press **OK**, as shown in the following window:

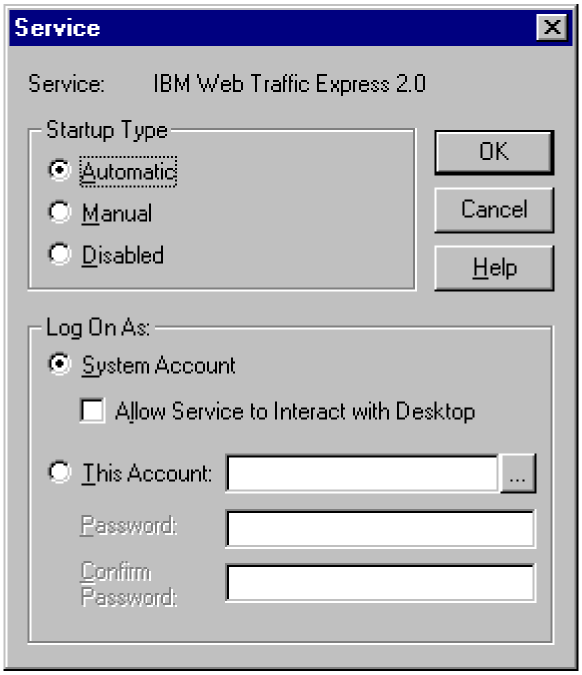

Service

Service: IBM Web Traffic Express 2.0

Startup Type

● Automatic

○ Manual

○ Disabled

Log On As:

● System Account

□ Allow Service to Interact with Desktop

○ This Account:

Password:

Confirm Password:

Figure 38. Setting Automatic Startup Type on Windows NT

Notice that if you also click on the square button **Allow Service to Interact with Desktop** (see Figure 38), WTE will be started, and activatate the WTE GUI automatically. The WTE GUI is as shown in Figure 37 on page 72.

When the GUI is active, you can use it to monitor the WTE activities and to stop or restart WTE.

3.5.3.3 Starting WTE from the Programs Menu

If you have installed WTE as only a Windows NT application by selecting only the **Base** option (see Figure 29 on page 65), you can start WTE from the **IBM Web**

Traffic Express icon, accessible from the Programs menu. This icon will be found in the IBM Web Traffic Express 2.0 program group. Note that this icon will be disabled if, when you installed WTE, you selected both the **Base** and **NT Services** options.

Starting WTE from the Programs menu as indicated, will also activate the GUI (see Figure 37 on page 72), which you can use to monitor WTE activities and to stop or restart WTE.

3.5.3.4 Viewing Events Logged by the WTE Server

If the WTE is started as a Windows NT service (see 3.5.3.2, “Starting and Stopping WTE from the Services Panel” on page 73), you can use the Windows NT Event Viewer to check if it has started successfully. Select **Event Viewer** from the Administrative Tools (Common) program folder. In the Event Viewer window, select the **Application** option from the Log menu, as shown in the following figure:

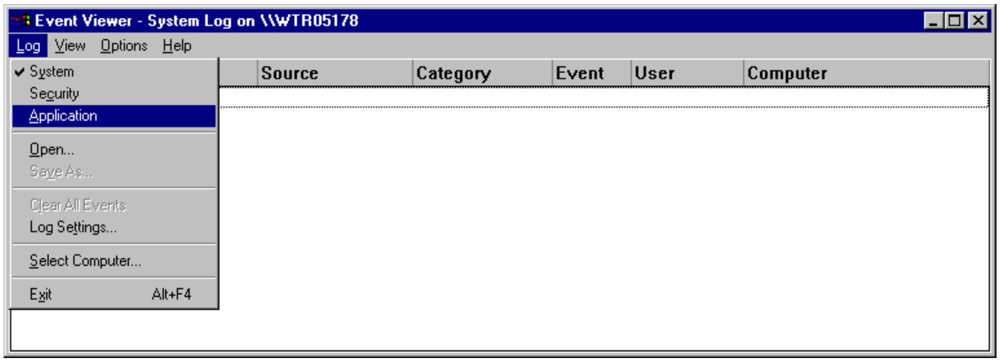

Figure 39. Event Viewer

The Event Viewer window registers a list of the application events. The following figure shows the entry related to IBM Web Traffic Express:

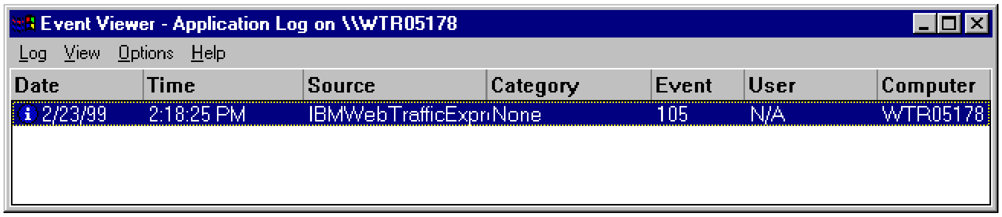

Date	Time	Source	Category	Event	User	Computer
i 2/23/99	2:18:25 PM	IBMWebTrafficExpr	None	105	N/A	WTR05178

Figure 40. Event Viewer - Application Log

In the Event Viewer window, you can highlight and double-click a particular event, in order to see more details about it. The following screen shows that WTE is running normally:

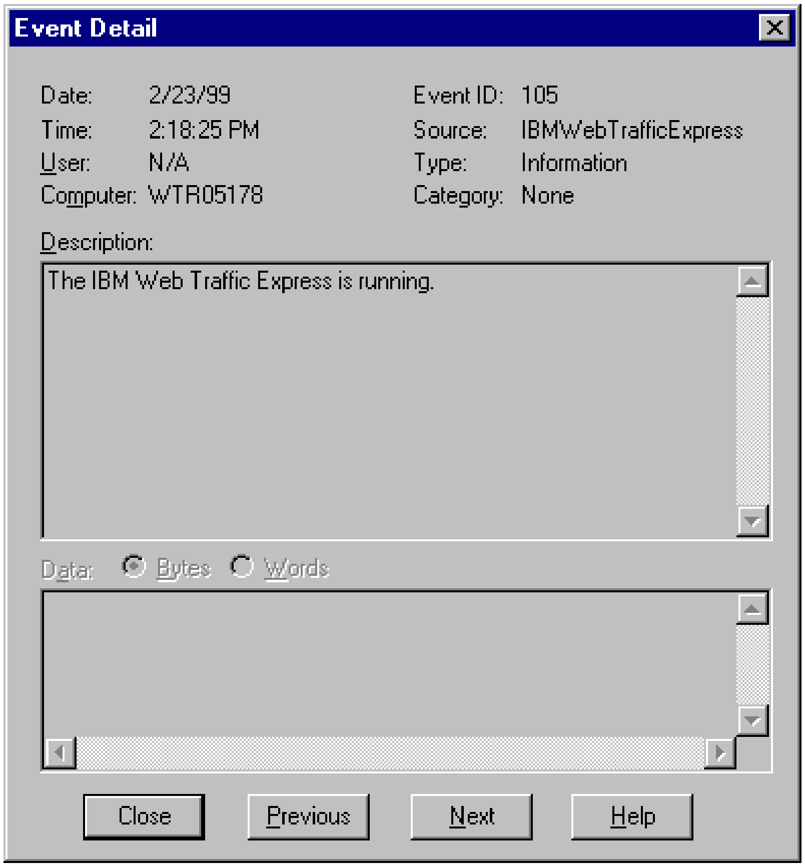

Date: 2/23/99 Event ID: 105

Time: 2:18:25 PM Source: IBMWebTrafficExpress

User: N/A Type: Information

Computer: WTR05178 Category: None

Description:

The IBM Web Traffic Express is running.

Figure 41. Event Detail - IBM WTE Is Running

3.6 Deinstallation of WTE on AIX, Solaris and Windows NT

IBM WebSphere Performance Pack provides an undo.sh script for AIX and Solaris and an undo.cmd script for Windows NT. If for any reason a re-installation of WTE is required, we recommend that all the WTE installed filesets be removed prior to the re-installation. For example, on AIX, if the AIX ODM licensed program product entries are not removed, the install.sh script or Java process will not re-install the necessary files. The filesets belonging to the IBM WTE are as described in 3.2.2, “Installing WTE on AIX” on page 54.

The undo commands are located in the directory where WebSphere Performance Pack was installed. Behind the scenes, the undo script runs the following command:

```
java wsppUninstall
```

Hence, it is required that Java is still installed on the system.

After the undo script is executed, you will be presented with the following screen:

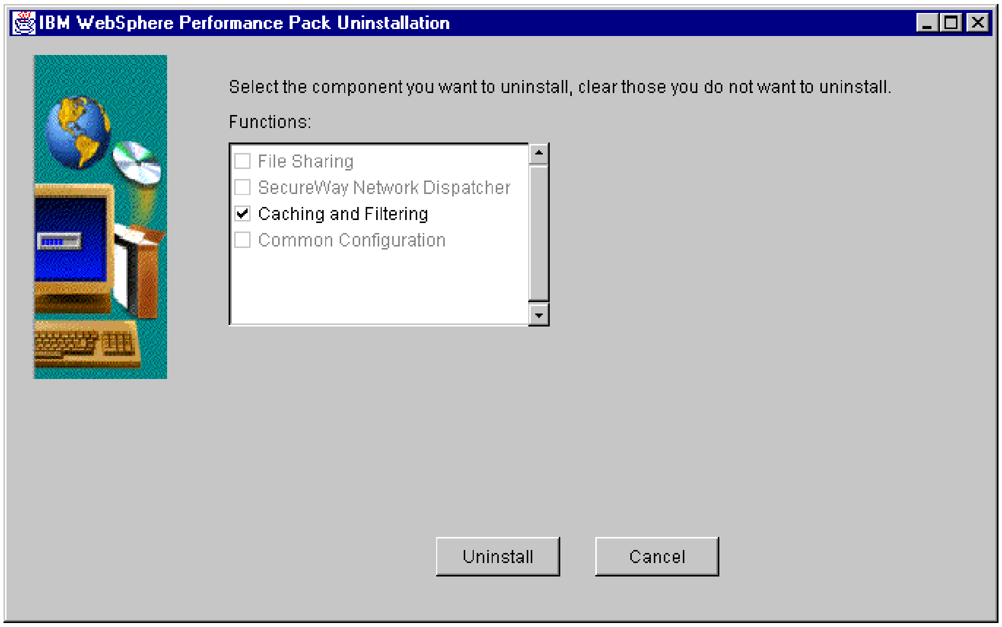

IBM WebSphere Performance Pack Uninstallation

Select the component you want to uninstall, clear those you do not want to uninstall.

Functions:

- □ File Sharing
- □ SecureWay Network Dispatcher
- ☑ Caching and Filtering
- □ Common Configuration

Figure 42. Uninstalling WTE - First Window

Click on the **Caching and Filtering** option followed by the **Uninstall** button to begin.

The next screen is a confirmation screen, where you click **OK** to proceed:

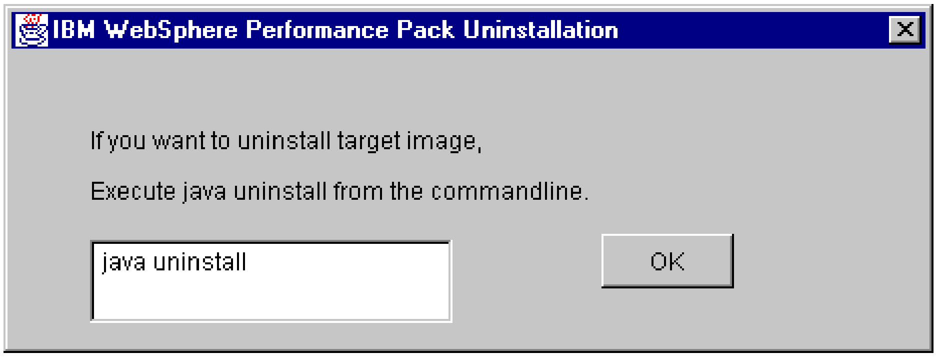

IBM WebSphere Performance Pack Uninstallation

If you want to uninstall target image,

Execute java uninstall from the commandline.

```
java uninstall
```

Figure 43. Uninstalling WTE- Second Window

Chapter 4. WTE Basic Configuration

In this chapter we describe how to perform a basic configuration of IBM Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2. The configuration issues we describe in this chapter are enough to get started with a fully functional WTE environment.

In particular, we go through the basic configuration options offered by the WTE Configuration and Administration Forms, and we will show how the same modifications can be obtained by manually editing the WTE configuration file.

For WTE advanced configuration refer to Chapter 5, “WTE Performance and Status” on page 99 through Chapter 15, “WTE Cancel Control” on page 237.

4.1 Connecting to the WTE Server

You can use your favorite Web browser to connect to the WTE server’s front page by setting the URL to . Follow the recommendations provided in 3.1.1, “Browser Requirements on the WTE Configuration Client” on page 49 to set up your browser for the configuration process.

From the WTE front page, you can access the following:

- Configuration and Administration forms
- IBM WTE browser tuning requirements
- IBM WTE Web site at `http://www.software.ibm.com/webservers/wte/`
- IBM WTE documentation

In our case, we pointed the browser to the URL `http://rs600022.itso.ral.ibm.com`, which was the URL of the AIX machine where we had installed WTE (see 3.2, “Installation on AIX” on page 52) and this is the window we saw:

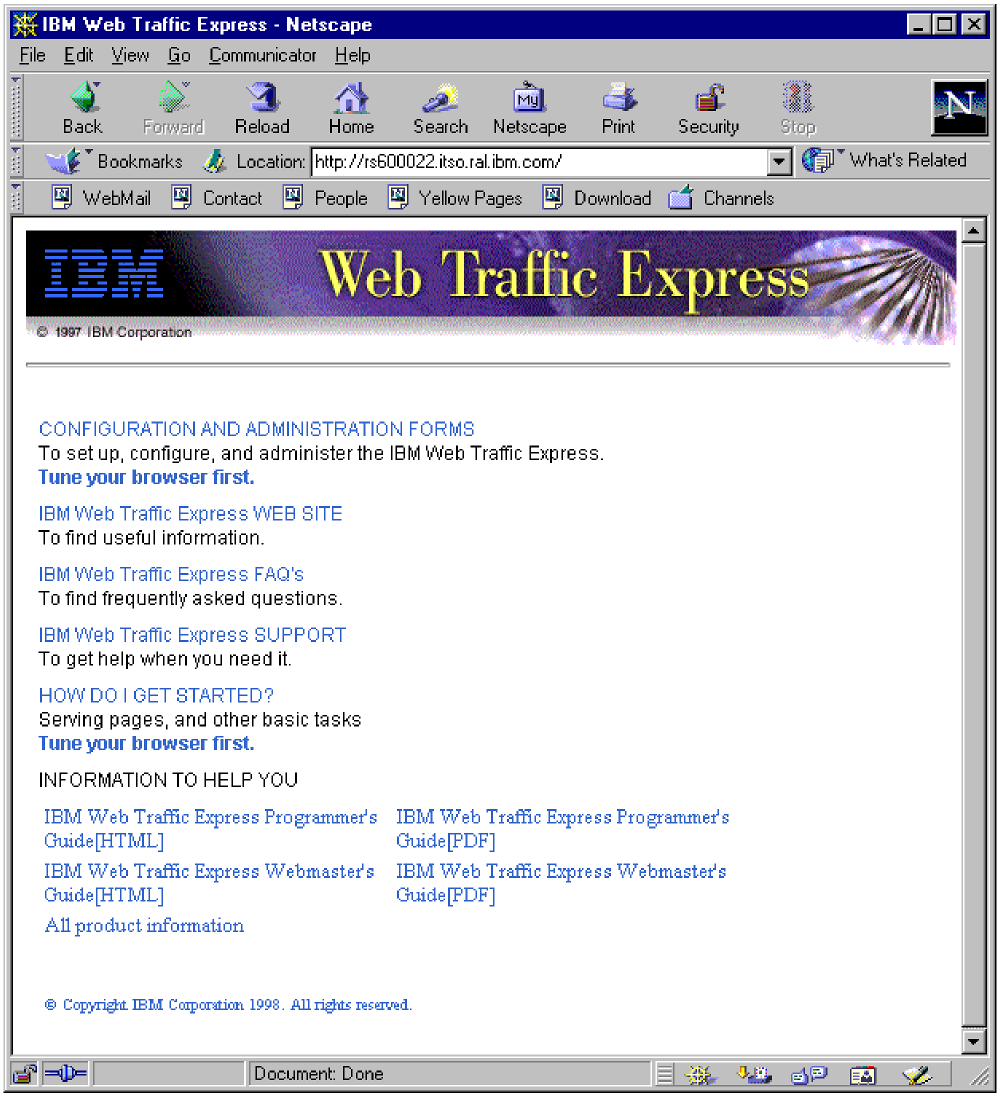

IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/
What's Related
WebMail Contact People Yellow Pages Download Channels
IBM Web Traffic Express
© 1997 IBM Corporation
Document: Done

CONFIGURATION AND ADMINISTRATION FORMS
To set up, configure, and administer the IBM Web Traffic Express.
Tune your browser first.

IBM Web Traffic Express WEB SITE
To find useful information.

IBM Web Traffic Express FAQ's
To find frequently asked questions.

IBM Web Traffic Express SUPPORT
To get help when you need it.

HOW DO I GET STARTED?
Serving pages, and other basic tasks
Tune your browser first.

INFORMATION TO HELP YOU

IBM Web Traffic Express Programmer's Guide[HTML]	IBM Web Traffic Express Programmer's Guide[PDF]
IBM Web Traffic Express Webmaster's Guide[HTML]	IBM Web Traffic Express Webmaster's Guide[PDF]
All product information	

© Copyright IBM Corporation 1998. All rights reserved.

Figure 44. IBM Web Traffic Express Front Page

4.2 WTE Administration Settings

Once the WTE server is running, you can start to configure it. To do this, you can access the Configuration and Administration forms from a Web browser. A link to the Configuration and Administration forms is found on the WTE front page (see Figure 44 on page 78). The browser can be on the same machine as the server or on any remote client that has access to the server. Remember that, to use the Configuration and Administration forms, you must first authenticate yourself by entering a recognized user name and password.

This section describes how to define a user ID and password to administer WTE. We show this process on AIX, Windows NT and Solaris, the three platforms where IBM WebSphere Performance Pack is supported, as there are some differences.

User Names, Passwords and Groups

The proxy server maintains its own list of user names, passwords, and group names. These entities are specific to the IBM WTE installation performed on your system, and are not related to system users and groups of the underlying operating system.

4.2.1 Administrator User ID and Password on AIX

For security reasons, the installation process on AIX does not create any default WTE administrator user names or passwords. Once defined, WTE administrator user names and passwords are stored in the WTE password file. The default is webadmin.passwd.

This password file is located in the /usr/lpp/internet/server_root/protect directory. It is empty by default after the installation, because, as we said, the installation process does not create any default administrator user names or passwords.

To define an administrator user name and password to access the Administration and Configuration forms from a Web browser, you can use the `htadm` command as follows:

```
htadm -adduser password_file user-name password real-name
```

The *real-name* field is optional. It should contain the name you want to use to identify the user name you are adding. WTE does not do anything with that field, which contains a description for the administrator, but if you don't put it in, the system will prompt you for it, at which time you can just press Enter and it will finish.

Notice that the htadm executable file is automatically installed in the /usr/sbin directory. For this reason, you will not need to specify the full path to this file when you launch the `htadm` command. Nonetheless, you have to specify the full path to the password file or execute `htadm` in the same directory where the password file is located, as shown:

```
cd /usr/lpp/internet/server_root/protect
htadm -adduser webadmin.passwd wteadmin wte1234 "WTEadmin: Poh Yee Tiong"
```

This will set wteadmin as the user name and `wte1234` as the password, and will remind us that wteadmin is the WTE server administrator. In other words, wteadmin will be the administrator of WTE. The information that you enter with the `htadm` command will be written into the password file. This is what we found in the file webadmin.passwd on our system after entering the above command:

```
wteadmin:VTX.OTki5iCYK:WTEadmin: Poh Yee Tiong
```

As you can see, the password is stored encrypted in the password file. Interestingly, WTE seems to be encrypting each new password with a different key. You can see this by defining two users with the same password and noticing that the same password is encrypted in two different ways.

You can also make the `htadm` command prompt you for the password, if you don't type that on the `htadm` line. In this case, the advantage is that you are prompted twice, and also the password never appears in the clear while you are typing it, because, at least on AIX, the `htadm` program does not show the characters you type in. This way, the password does not even show up in the history list.

4.2.2 Administrator User ID and Password on Solaris

The considerations for the Solaris platform are very similar to what we said about the AIX platform. For this reason, we recommend that you refer to 4.2.1, "Administrator User ID and Password on AIX" on page 79.

4.2.3 Administrator User ID and Password on Windows NT

On Windows NT systems, user names and passwords are stored in the caching proxy server default password file, which is admin.pwd. This file is located in the Windows NT installation directory, which by default is C:\WINNT.

On Windows NT systems, unlike on AIX systems, after the installation of WTE this file contains a default user known as admin with a default password, `admin`. For security reasons, you should remove this default user ID and password immediately after the installation of WTE, as discussed in "One Note about the Default User" on page 81.

To define a different administrator user name and password to use the Administration and Configuration forms from a Web browser, we use the `htadm` command as follows:

```
htadm -adduser password_file user-name password real-name
```

Notice that, by default, HTAdm.exe is installed in the C:\WSPP\WWW\bin directory, and this path is automatically added to the Path system environment variable when WTE is installed. Therefore, you can launch the `htadm` command from the command line without the need to specify the full path to HTAdm.exe. Nonetheless, for *`password_file`,* you have to specify the full path or execute `htadm` in the directory where the password file resides.

The *`real-name`* field is not really required. It should contain the name you want to use to identify the user name. WTE does not do anything with that field, but if you don't specify it, the system will prompt you for it, at which time you can just press Enter for no identity or enter the real user name.

We entered the following three commands:

```
C:
cd \WINNT
htadm -adduser admin.pwd wteadmin wte1234 "WTEadmin: Poh Yee Tiong"
```

This will set wteadmin as the user name, wte1234 as the password. This is the administrator of the WTE. The password that you enter will be written into the password file. This is what we found in the file admin.pwd on our system after entering the above command:

```
admin:ZsveyKXG2NfGc:Administrator
wte.admin:MBqPp41Q6WdFU:WTEadmin: Poh Yee Tiong
```

You can also make the `htadm` command prompt you for the password, if you don't type it on the `htadm` command line. This has the advantage that you are prompted to enter the same password twice, for verification reasons. However, take note that on Windows NT, passwords show up in clear when you are entering and verifying them.

Notice also that the password is stored encrypted in the password file. Interestingly, WTE seems to be encrypting each new password with a different key. You can see this by defining two users with the same password and noticing that the same password is encrypted in two different ways.

One Note about the Default User

For security reasons, it is recommended that you remove the admin user that was installed by default during the installation or at least that you change its password.

1. To remove the default user, enter the following command from the directory where admin.pwd is located:

   ```
   htadm -deluser admin.pwd admin
   ```

2. To change only the password for the user admin, enter the following command from the directory where admin.pwd is located:

   ```
   htadm -passwd admin.pwd password
   ```

   where *password* is your new password

Restricting the Access to HTAdm.exe and admin.pwd

At the time this book was written, we found that the command file \WSPP\WWW\bin\HTAdm.exe and the password file \WINNT\admin.pwd were accessible to every user of the system.

It is recommended that you define the WTE administrator as a Windows NT user on the machine where WTE is installed, and then set these two files to be accessible to the WTE administrator only. Notice that you can do this only if you are using NTFS, which is a prerequisite file system for the WTE installation on Windows NT.

4.3 Basic Configuration

This section describes the basic configuration steps for WTE and shows how the configuration can be issued through directives in the configuration files or through the graphical user interface (GUI).

4.3.1 The WTE Configuration File

The WTE server functionality is regulated according to the directives contained in its configuration file. In WTE V2, which comes with IBM WebSphere Performance Pack Version 2, there exists only one configuration file for all its functions. This file is ibmproxy.conf and its path for the three supported platforms is described in Table 3 on page 82. In WTE V1, which was the Caching and Filtering component of IBM WebSphere Performance Pack V1, there were multiple configuration files, each with a specific configuration function. Table 3 on page 82 shows how the configuration files of the first version of the product relate to the single configuration file used in Version 2:

Table 3. The IBM Web Traffic Express Configuration Files

	IBM WTE Version 1	IBM WTE Version 2
AIX	/etc/httpd.conf /etc/javelin.conf /etc/socks.conf	/etc/ibmproxy.conf[a]
Solaris		
NT	C:\WINNT\httpd.cnf C:\WINNT\javelin.cnf C:\WINNT\socks.cnf	C:\WINNT\ibmproxy.conf[b]

a. Version 1 files are merged and saved in /etc/ibmproxy.conf.old if migrated to Version 2.
b. Version 1 files are merged and saved in \WWW\httpdjavelin.merged if migrated to Version 2.

All the directives supported by the ibmproxy.conf file are shown in the following list:

- Basic directives
- Process control
- Logging directives
  - Log file directives
  - Log archive directives
  - Log filtering directives
- Method directives
- Content presentation directives
  - Welcome pages directives
  - Directory browsing directives
  - CGI program directives
  - Content type directives
- Error message directives
- API directives
- User authentication and document protection
- Mapping rules
- Performance directives
- Timeout directives
- Proxy directives
- Proxy caching directives
- Proxy cache garbage collection directives
- Advanced proxy and caching directives
- Remote Cache Access (RCA) directives
- SNMP directives
- Icon directives
- Cache agent directives
- PICS filtering directives
- Miscellaneous directives

There are two ways you can set or modify the ibmproxy.conf of the WTE server:

1. Edit the ibmproxy.conf file and modify the configuration directives.
2. Access the Configuration and Administration forms from a Web browser after authenticating as the WTE server administrator.

The Common Configuration utility is actually another option you can select to configure WTE although its functionality in WTE V2 is limited to the PICS filtering configuration only. See Chapter 20, “Common Configuration” on page 293 for further details.

In this section, we first show you how to configure WTE using the Configuration and Administration forms, as they provide a combination of HTML forms and CGI-BIN programs that build a user interface that is very simple to interact with. Then, we describe what the corresponding directives would be in the configuration files.

4.3.2 Using the Configuration and Administration Forms

We performed the configuration of WTE on the same AIX machine we used to show the installation process (see 3.2, “Installation on AIX” on page 52). The configuration of this component on Solaris and Windows NT is very similar, and any difference will be explicitly specified.

As we explained in 4.1, “Connecting to the WTE Server” on page 77, once the WTE server has started, you can access its front page from a Web browser located on the same machine or on another machine connected to the network. To do this, enter the URL of the server’s front page. The URL we connected to was `http://rs600022.itso.ral.ibm.com`, since rs600022.itso.ral.ibm.com was the fully qualified host name of the machine where WTE was installed.

The WTE server’s response to this request was similar to the screen shown in Figure 44 on page 78. This is the home page of the WTE server. To open the Configuration and Administration Forms, click the **CONFIGURATION AND ADMINISTRATION FORMS** Web link.

Disabling the Cache on the Browser

It is recommended that you *disable* caching on your browser or alternately, that you configure the browser to compare the cached document with the network document *every time*. This is to prevent reading in stale forms information from the browser cache, which could create confusion.

Disabling the cache can be achieved by setting Disk Cache and Memory Cache to `0 KBytes`. Configuring the browser to compare the cached documents with the network documents every time can be achieved as indicated in Figure 18 on page 51.

We found that after the cache is disabled, and when resizing the window, we received an error from the browser:

```
Your browser does not have proper Java or Javascript support.
```

You can recover from this error by clicking the **Reload** button of the WTE GUI.

If you have not used the Configuration and Administration forms since starting your browser, you will be prompted for the IBM WTE server administrator’s user name and password, as shown next:

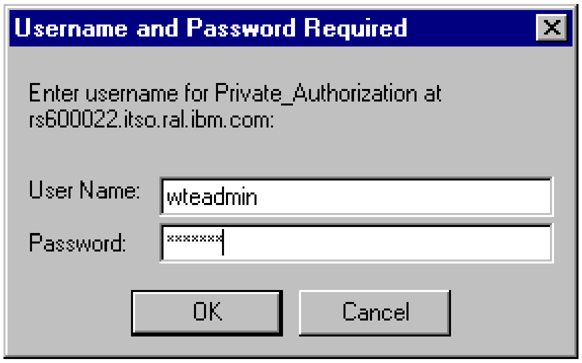

Enter username for Private_Authorization at rs600022.itso.ral.ibm.com:

User Name: wteadmin

Password: *******

Figure 45. Enter IBM WTE Server User Name and Password

Here you should enter the administrator user ID and password you defined earlier (see 4.2, “WTE Administration Settings” on page 78). We entered `wteadmin` in the User Name field and `wte1234` in the Password field, and then we clicked the **OK** button. If the authentication is successful, you will be taken to the WTE GUI Introduction configuration page, as shown in the following figure:

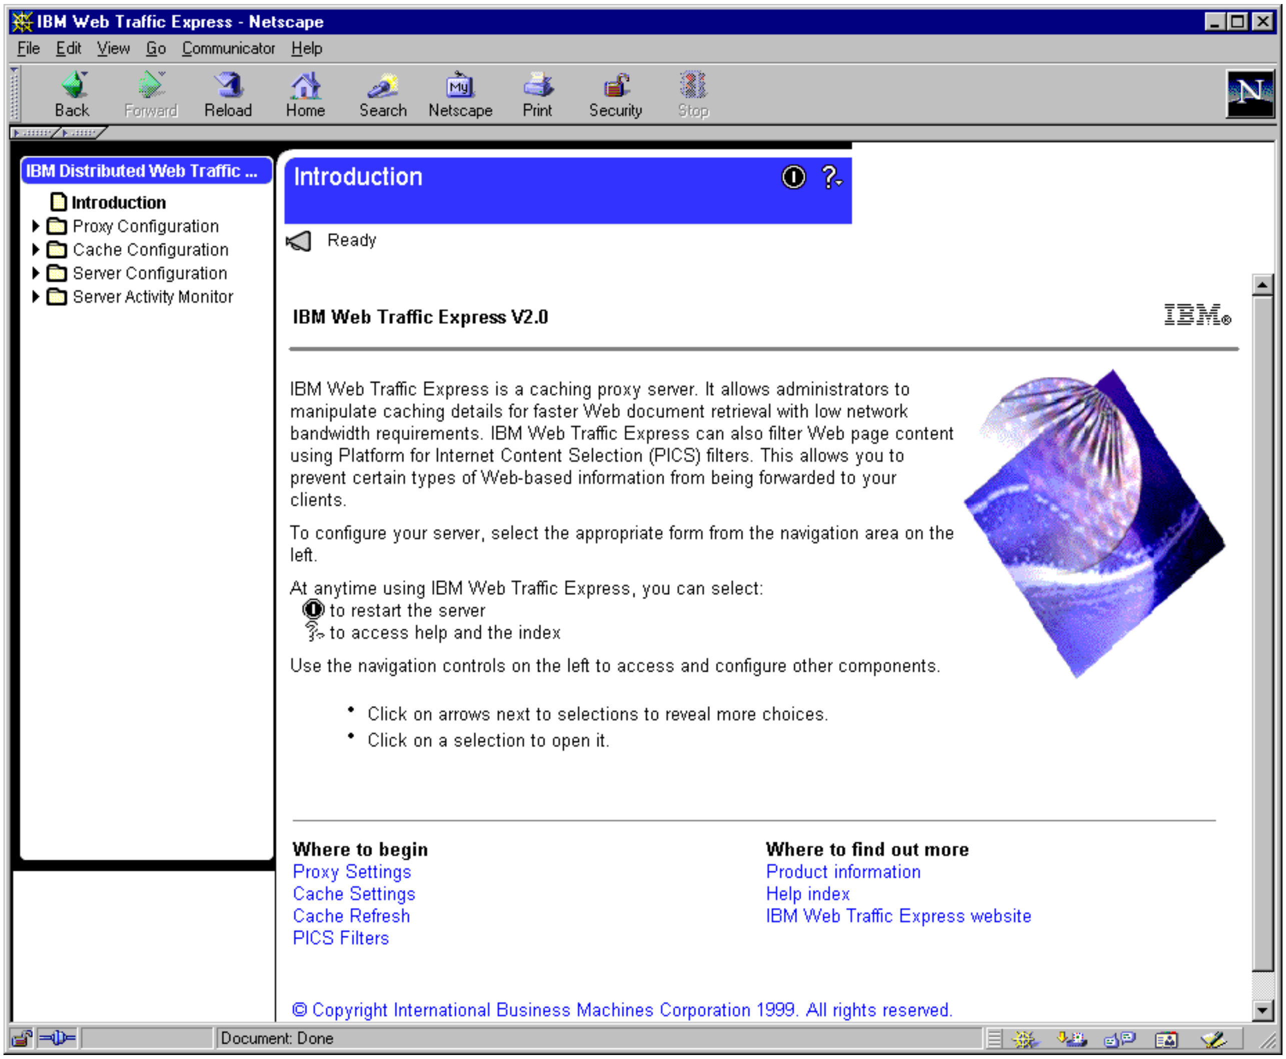

IBM Web Traffic Express V2.0

IBM Web Traffic Express is a caching proxy server. It allows administrators to manipulate caching details for faster Web document retrieval with low network bandwidth requirements. IBM Web Traffic Express can also filter Web page content using Platform for Internet Content Selection (PICS) filters. This allows you to prevent certain types of Web-based information from being forwarded to your clients.

To configure your server, select the appropriate form from the navigation area on the left.

At anytime using IBM Web Traffic Express, you can select:
① to restart the server
? to access help and the index

Use the navigation controls on the left to access and configure other components.

- Click on arrows next to selections to reveal more choices.
- Click on a selection to open it.

Where to begin
Proxy Settings
Cache Settings
Cache Refresh
PICS Filters

Where to find out more
Product information
Help index
IBM Web Traffic Express website

© Copyright International Business Machines Corporation 1999. All rights reserved.

Figure 46. The IBM Web Traffic Express Introduction

For the rest of the book, we will refer to this as the WTE GUI and we will refer to the left frame as the *navigation frame*.

You can click one of the folders in the navigation frame on the left. The items are **Proxy Configuration**, **Cache Configuration**, **Server Configuration**, and **Server Activity Monitor**.

If you click the **Proxy Configuration** folder in the navigation frame, you will be presented with the forms shown in Figure 47:

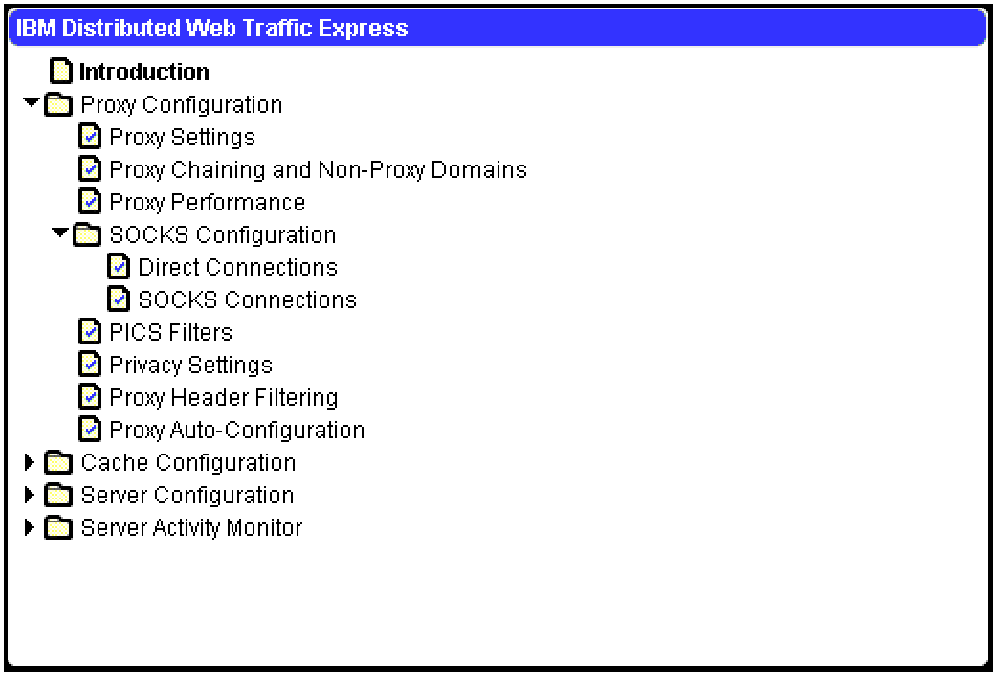

Figure 47. IBM WTE GUI Navigation Frame - Proxy Configuration

If you click the **Cache Configuration** folder in the navigation frame, you will be presented with the forms shown in Figure 48:

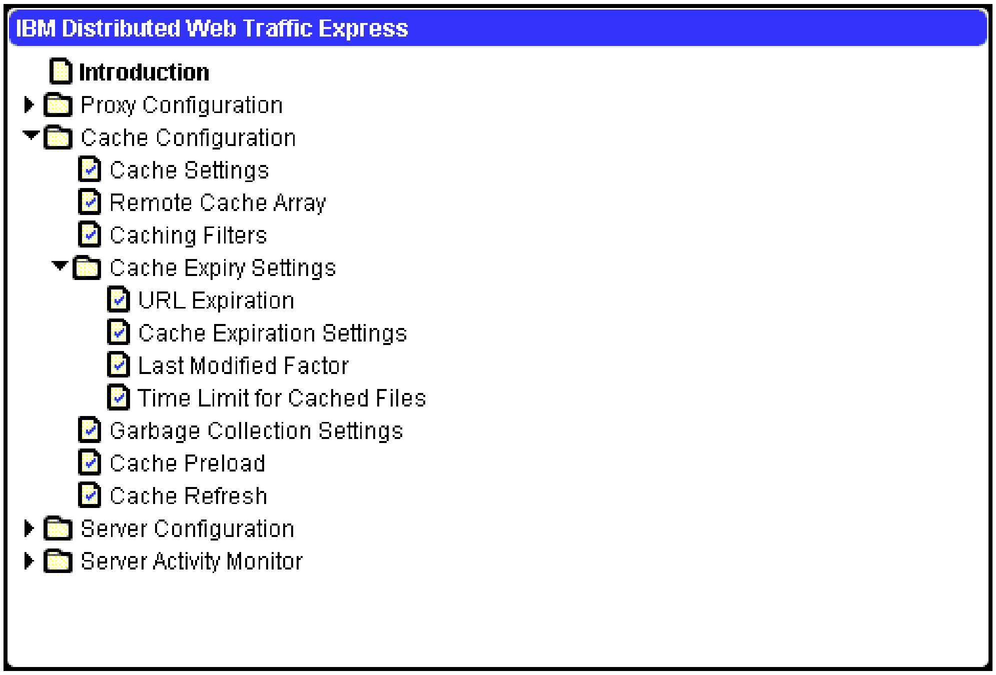

Figure 48. IBM WTE GUI Navigation Frame - Cache Configuration

Click the **Server Configuration** folder in the navigation frame, and you will be presented with the forms shown in Figure 49:

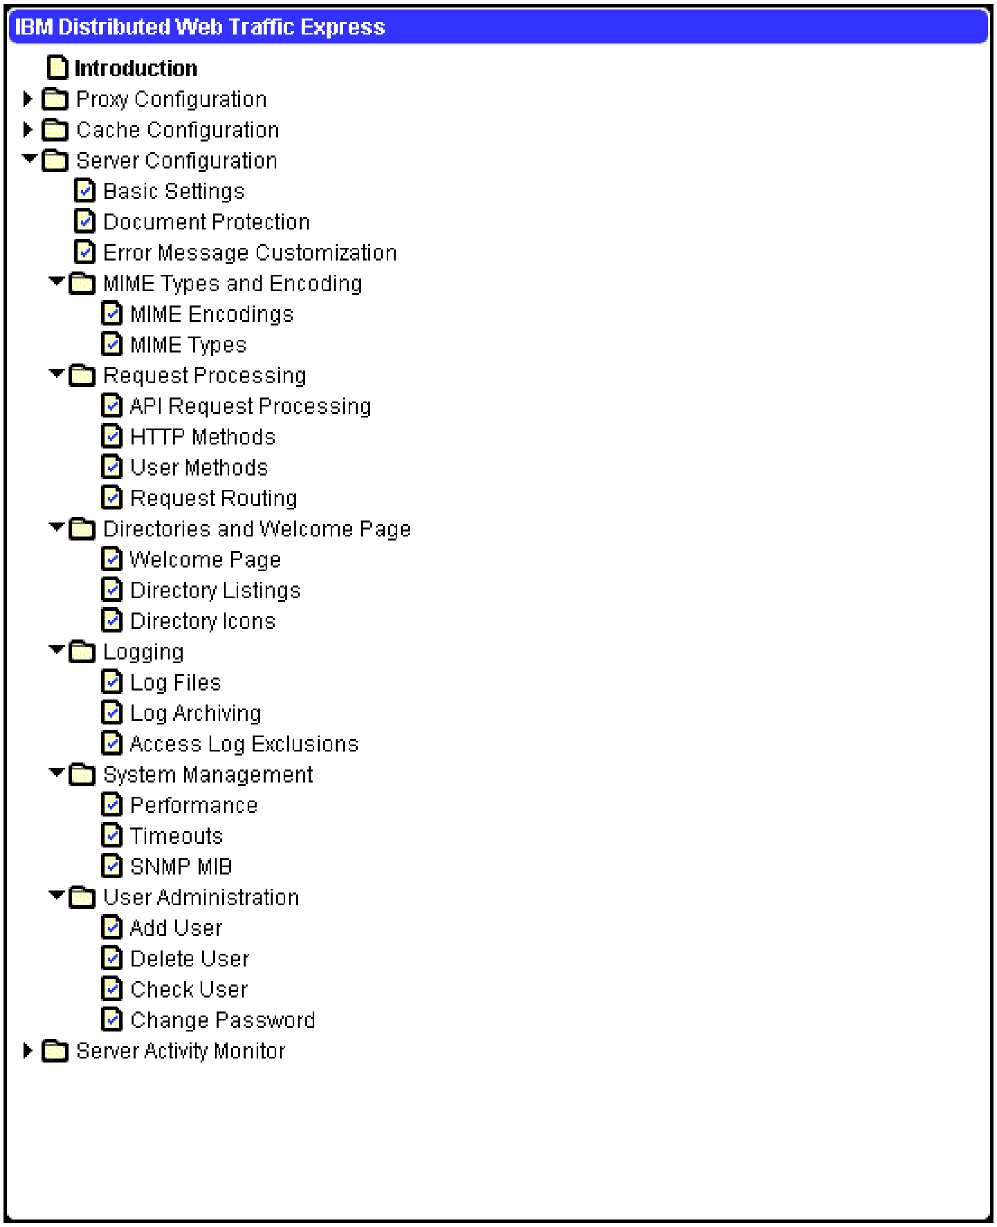

Figure 49. IBM WTE GUI Navigation Frame - Server Configuration

Click the **Server Activity Monitor** folder in the navigation frame, and you will be presented with the forms shown in Figure 50:

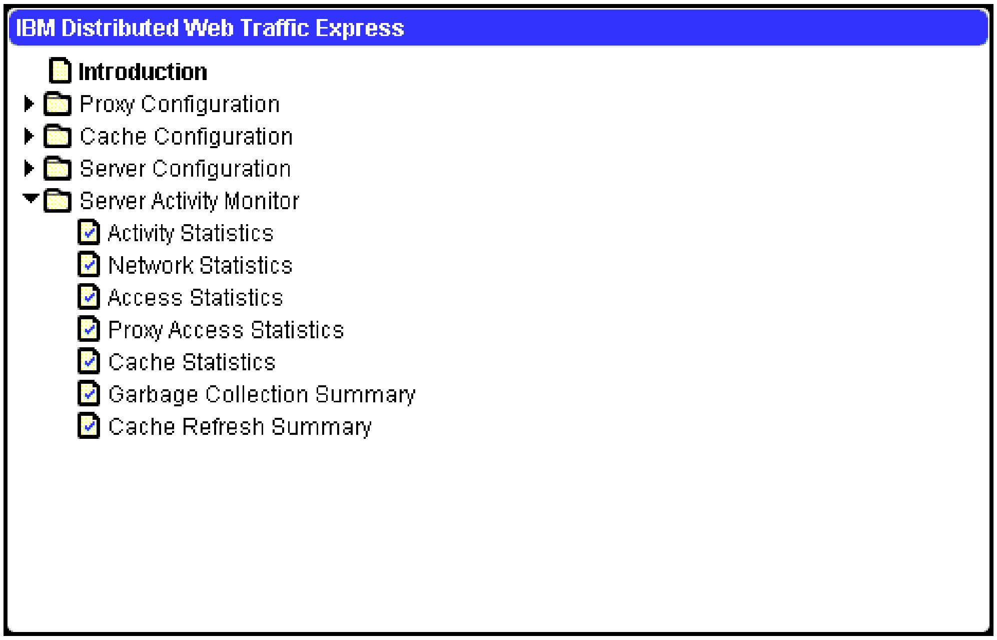

Figure 50. IBM WTE GUI Navigation Frame - Server Activity Monitor

Click the form name in the navigation frame and you will be presented with the relevant forms. The GUI will display the current configuration values in its input fields. If you have not changed your configuration since installation, the values you see are the defaults.

From any form, you can enter information about how you want to configure that particular part of your server. Each form provides instructions to assist you in deciding what changes to make. For further information, you can click the help icon **?** at the top of each form. The **?** help icon provides links to:

- **Field help**

  Descriptions of the fields on each panel

- **How do I?**

  Detailed steps for using the form to perform particular tasks

- **Index**

  Index of the help information

After you have completed the information for the form, you must click **Submit** to update the server configuration with the changes you made. The **Submit** button is located below the input fields on each form. If you decide you do not want to use the changes you made to the form, click **Reset**. This returns the fields on the form to the values they had when you first came to the form.

After you click on the **Submit** button, the server will show you a message indicating whether your input was accepted or not.

To restart the server from the Configuration and Administration forms, click the restart icon in the top frame. You will see a message saying:

```
Restarting.. httpd
```

Restart completes in-process requests, stops accepting requests, and reloads the changed configuration file. After the changed configuration file has been reloaded, the server accepts requests again.

4.3.3 Directives not Changed on Restart

Some directives are not refreshed on a restart. If these directives are changed while the server is running, you must stop the server, then start it again, as indicated in 3.5, “Starting and Stopping WTE” on page 70. The directives that are not refreshed on a WTE restart are shown in the following list:

Table 4. Directives not Changed on Restart

Directive Group	**Directives**

Caching	`CacheRoot`
	`Caching`
	`CacheSize`
	`CacheFiles`
	`DiskBlockSize`

SNMP	`Snmp`
	`SnmpCommunity`
	`WebmasterEmail`

Logging	`AccessLog`
	`CacheAccessLog`
	`ErrorLog`
	`ProxyAccessLog`
	`ServerRoot`

CGI	`InheritEnv` `DisinheritEnv`
Performance	`MaxActiveThreads`
UNIX Process Control	`UserId` `GroupId`

Network Access	`Hostname`
	`BindSpecific`
	`ListenBacklog`
	`Port`

Miscellaneous	`TransparentProxy`[a]

[a] Used on AIX only.

4.3.4 Enabling Proxy Function

To enable the proxy function and which protocols WTE should process, click **Proxy Settings** in the navigation frame on the WTE GUI and you will get the Proxy Settings form, which is shown in the following figure:

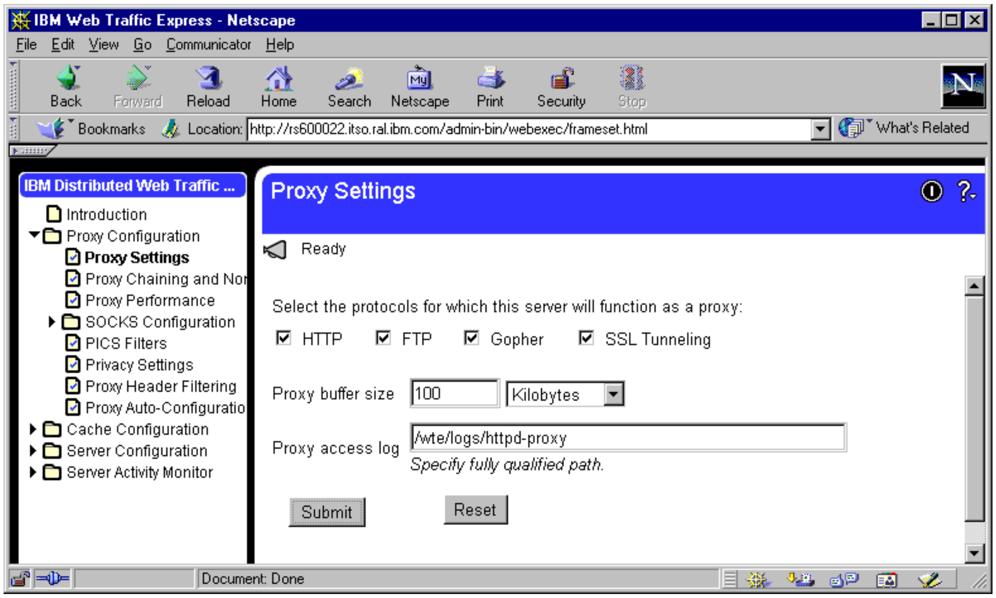

Proxy Settings

Ready

Select the protocols for which this server will function as a proxy:

☑ HTTP ☑ FTP ☑ Gopher ☑ SSL Tunneling

Proxy buffer size 100 Kilobytes

Proxy access log /wte/logs/httpd-proxy
Specify fully qualified path.

Submit Reset

Figure 51. Proxy Settings

WTE is a proxy server when installed. The default setting automatically provides proxy function for the following protocols:

- HTTP
- FTP
- Gopher
- SSL Tunneling

When the proxy returns dynamic data, such as output data from CGI-BIN programs, API programs, server-side includes, and Java servlets, it must buffer the data, and you can set the value of the buffer size in the Proxy buffer size field. We accepted the default value, which is 100 KB on all three supported platforms.

The next entry in the Proxy Settings form is Proxy access log. This field allows you to specify the file name where the WTE server will write the access statistics. The fully qualified path is required. Each day at midnight, the WTE server, if it is running, starts a new log file. It uses the specified file name and appends a date suffix as an extension. These log files are used for logging proxy requests.

Default values are:

- /usr/lpp/internet/server_root/logs/httpd-proxy on AIX
- C:\WSPP\WWW\Logs\httpd-proxy on Windows NT

As this log file can take up a significant amount of space on your file systems, it is recommended that a different file system be created to hold the log files. In our situation, on AIX, a file system called /wte/logs was created to store all the log files.

We accepted the default settings, but we modified the value of the Proxy access log field, and set it to `/wte/logs/http-proxy`, since we wanted to locate the log files in the /wte file system we had created earlier.

Using the Directives

The ibmproxy.conf configuration file uses directives to enable the proxy function. The configuration we have just presented could have been issued by modifying the ibmproxy.conf configuration file using the following directives:

```
Proxy http:*
Proxy ftp:*
Proxy gopher:*
SSL Tunneling On
MaxContentLengthBuffer 100 K
ProxyAccessLog /wte/logs/httpd-proxy
```

Note that permissions are important when changing the default directory where the log files will reside. The WTE server will write to that directory as the server's user ID and group ID specified in the ibmproxy.conf configuration file (nobody/nobody by default). Therefore, if you have created a new directory for the log files, you must ensure that the WTE server's user ID can write to that directory.

The process of changing the default proxy access log path and file name on Windows NT is very similar.

For the configuration changes to be written in the configuration files, click **Submit** on the WTE GUI, and you will get a reply reporting that the requested configuration changes have been completed successfully.

For the changes to take effect, you have to restart the WTE server and click the **Reset** button on the WTE GUI to restart the httpd daemon, which will re-read the ibmproxy.conf file. The result is as shown in Figure 52.

Note on Clicking the Restart Icon

Clicking the **Restart** icon will stop and restart the httpd daemon or whttpg.exe. However, be aware that no confirmation screen is provided to inform you that the restart operation took place successfully.

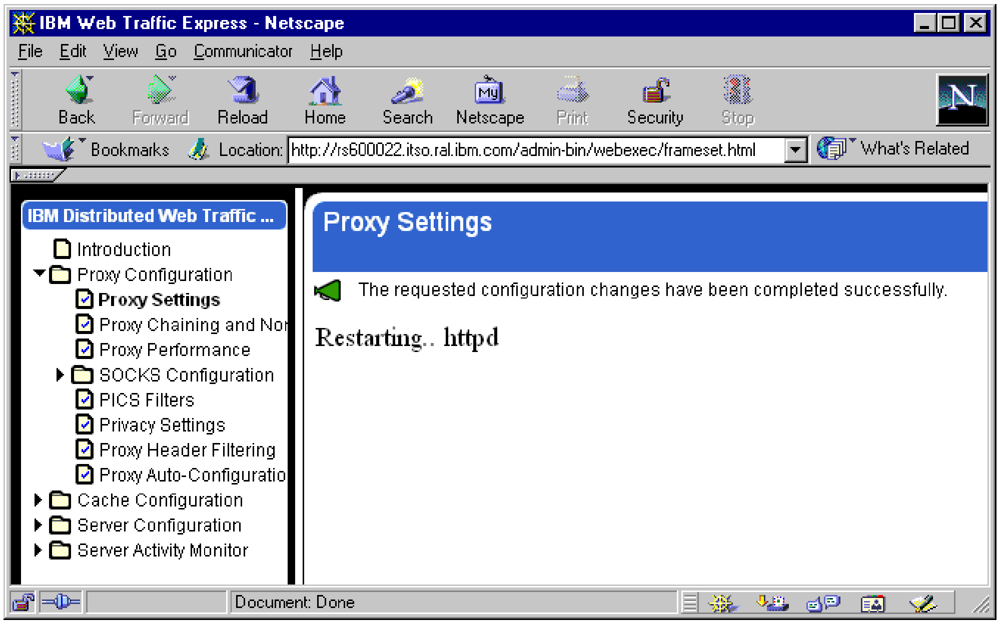

The requested configuration changes have been completed successfully.

Restarting.. httpd

Figure 52. IBM WTE GUI - Restarting WTE

4.3.5 Enabling the Pure Proxy Function

The WTE server's function is to be a caching and filtering proxy server. In theory, WTE could even work as a content server. However, for best performance, it is recommended to use it only as a proxy server. In fact, by default, the initial settings instruct this component to act as a pure proxy server. You can see this if you select the **Proxy Performance** form from the navigation menu, shown in Figure 53:

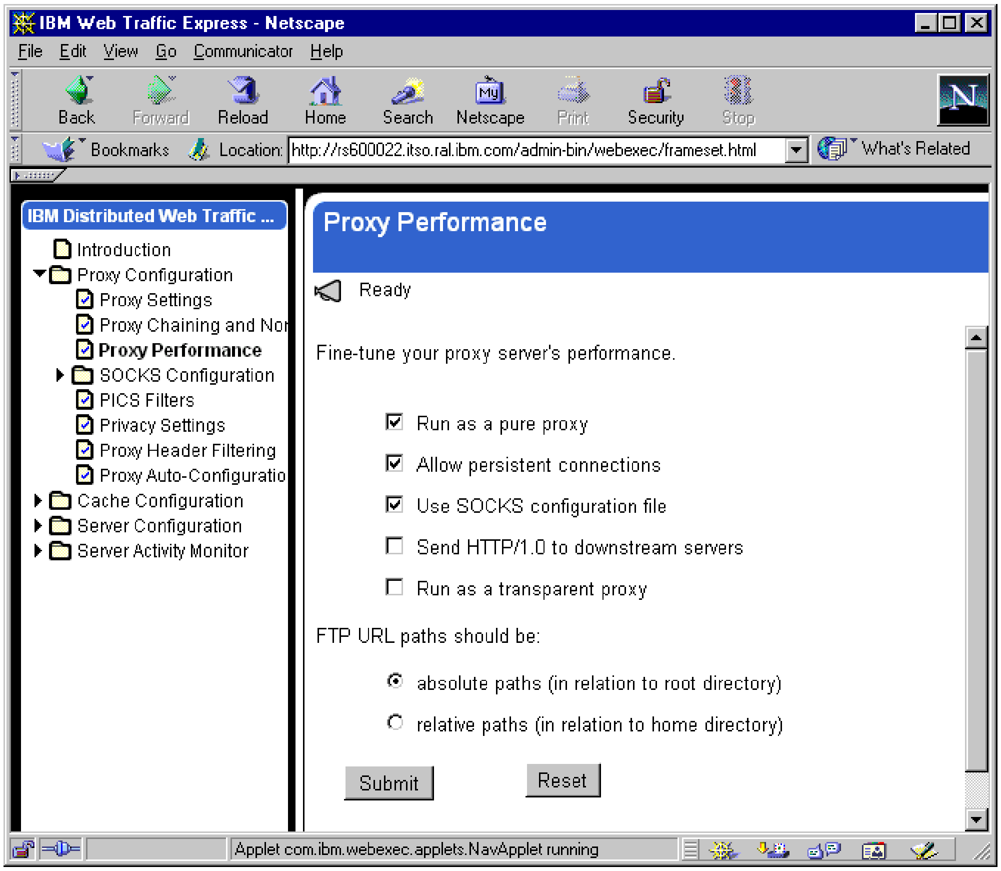
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Distributed Web Traffic ...
Introduction
Proxy Configuration
Proxy Settings
Proxy Chaining and Nor
Proxy Performance
SOCKS Configuration
PICS Filters
Privacy Settings
Proxy Header Filtering
Proxy Auto-Configuratio
Cache Configuration
Server Configuration
Server Activity Monitor
Submit
Reset
Applet com.ibm.webexec.applets.NavApplet running

Proxy Performance

Ready

Fine-tune your proxy server's performance.

☑ Run as a pure proxy

☑ Allow persistent connections

☑ Use SOCKS configuration file

□ Send HTTP/1.0 to downstream servers

□ Run as a transparent proxy

FTP URL paths should be:

● absolute paths (in relation to root directory)

○ relative paths (in relation to home directory)

Figure 53. Performance Settings

The parameter we were interested in was **Run as a pure proxy**. We kept the default configuration, with the **Run as a pure proxy** check box marked, because we wanted to use the WTE server as a pure caching and filtering proxy server.

Using the Directives

The ibmproxy.conf configuration file uses the following directive to enable the WTE to function as a pure proxy:

```
PureProxy On
```

By default, the WTE server is configured to function as a pure proxy; hence, we did not need to restart the server for this setting.

4.3.6 Enabling Basic Caching

In this section, we will enable the basic caching capabilities of the WTE. Caching is done by locally saving copies of the files that the clients request. This way, WTE can use the cached files, when they are subsequently requested, to serve its clients quickly because it does not need to fetch the files from a remote Web server again. Note that not all files can be cached. A list of non-cacheable files can be found in 2.5.1, “Controlling Which Documents Are Kept in the Cache” on page 36.

Caching and its parameters can be configured at install time. However, the following Caching Settings form can be used to modify what is set or to enable caching.

To enable the caching capabilities of the WTE, *caching must be enabled* (default is off!) and you *must* define the root directory for cached files (by default, this is not defined) and cache access log file name.

Select **Cache Settings** from Cache Configuration on the navigation frame, to retrieve the Cache Settings form, as shown in the following figure:

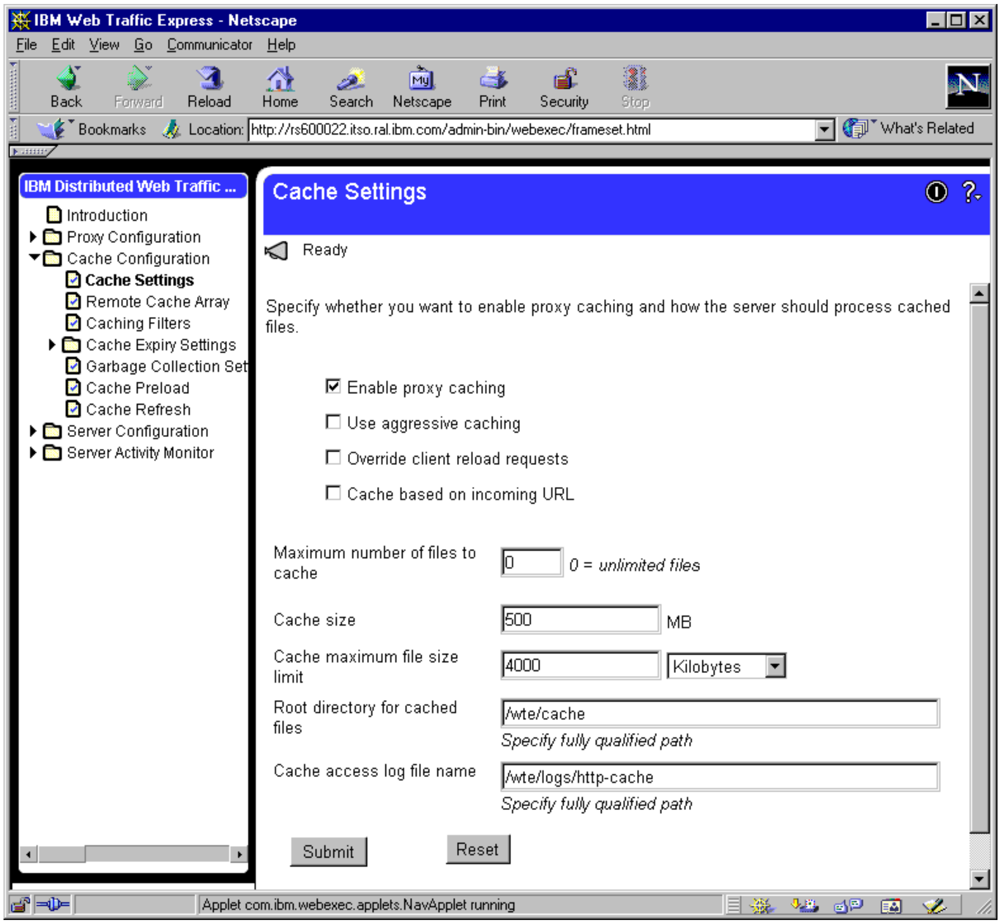
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Distributed Web Traffic ...
Introduction
Proxy Configuration
Cache Configuration
Cache Settings
Remote Cache Array
Caching Filters
Cache Expiry Settings
Garbage Collection Set
Cache Preload
Cache Refresh
Server Configuration
Server Activity Monitor
Ready
Submit
Reset
Applet com.ibm.webexec.applets.NavApplet running

Cache Settings

Specify whether you want to enable proxy caching and how the server should process cached files.

☑ Enable proxy caching

☐ Use aggressive caching

☐ Override client reload requests

☐ Cache based on incoming URL

Maximum number of files to cache	0	*0 = unlimited files*
Cache size	500	MB
Cache maximum file size limit	4000	Kilobytes
Root directory for cached files	/wte/cache *Specify fully qualified path*	
Cache access log file name	/wte/logs/http-cache *Specify fully qualified path*	

Figure 54. IBM WTE GUI - Cache Settings

Check on the **Enable proxy caching** check box to enable the caching capabilities.

WTE allows you to limit the amount of disk space used for the entire cache. You can use the Cache size field to enter the maximum amount (in MB) of disk space you want to reserve for the proxy cache. We accepted the default value of 500 MB, which was defined at install time.

In the text field named Root directory for cached files, enter the directory under which all files will be cached. Notice that no value is defined by default if caching is not defined initially. On the UNIX platform, we stored the cache files in the

/wte/cache directory which we had created earlier. Therefore, we entered `/wte/cache`.

Cache access log file is used for logging hits on the proxy cache. This is only valid if the WTE server is running as a proxy. We will store the cache access log file in the /wte/logs/httpd-cache file.

Notice the requirement to specify fully qualified paths.

Using the Directives

The ibmproxy.conf configuration file should use the following directives:

```
CacheSize 500
CacheLimit_2 400 K
CacheRoot /wte/cache # UNIX Servers
CacheAccessLog /wte/logs/httpd-cache # UNIX servers
```

Again, note that permissions are important when changing the default directory where the log files will reside. The WTE server will write to that directory as the user ID/group ID specified in the ibmproxy.conf configuration file (nobody/nobody by default). So if you have created a new directory for the logs, you must ensure that the WTE server’s user ID can write to that directory.

For the changes to take effect, you have to restart the WTE server. Click the **Restart** icon on the WTE GUI to restart the httpd daemon. This operation will force the httpd daemon to re-read the ibmproxy.conf file.

Note on Clicking the Restart Icon

Clicking the **Restart** icon will stop and restart the httpd daemon. However, be aware that no confirmation screen is provided to inform you that the restart operation took place successfully.

4.3.7 Garbage Collection

Once caching is enabled, disk space and file maintenance are required. WTE provides a storage reclamation process, known as *garbage collection*, for disk space and file management. This process should be enabled so that you can prevent the cache of your proxy server from growing beyond the maximum size that you set.

Garbage collection is a process that deletes all expired files as defined by the Cache Expiration Settings form. Expired files should be the cached files that are no longer used. We configured the Cache Expiration Settings form to:

- Expire a cached HTTP file which was unused for 2 days
- Expire a cached FTP file which was unused for 3 days
- Expire a cached Gopher file which was unused for 12 hours
- Set FTP Default Expiration time to 1 day
- Set Gopher Default Expiration time to 2 days

- Enable cached file expiration checking
- Disable caching of files due to expire within 10 minutes

The WTE cache expiration is set by clicking the **Cache Configuration** folder, followed by **Cache Expiry Settings** and **Cache Expiration Settings**. The WTE Cache Expiration Settings form is shown in Figure 55:

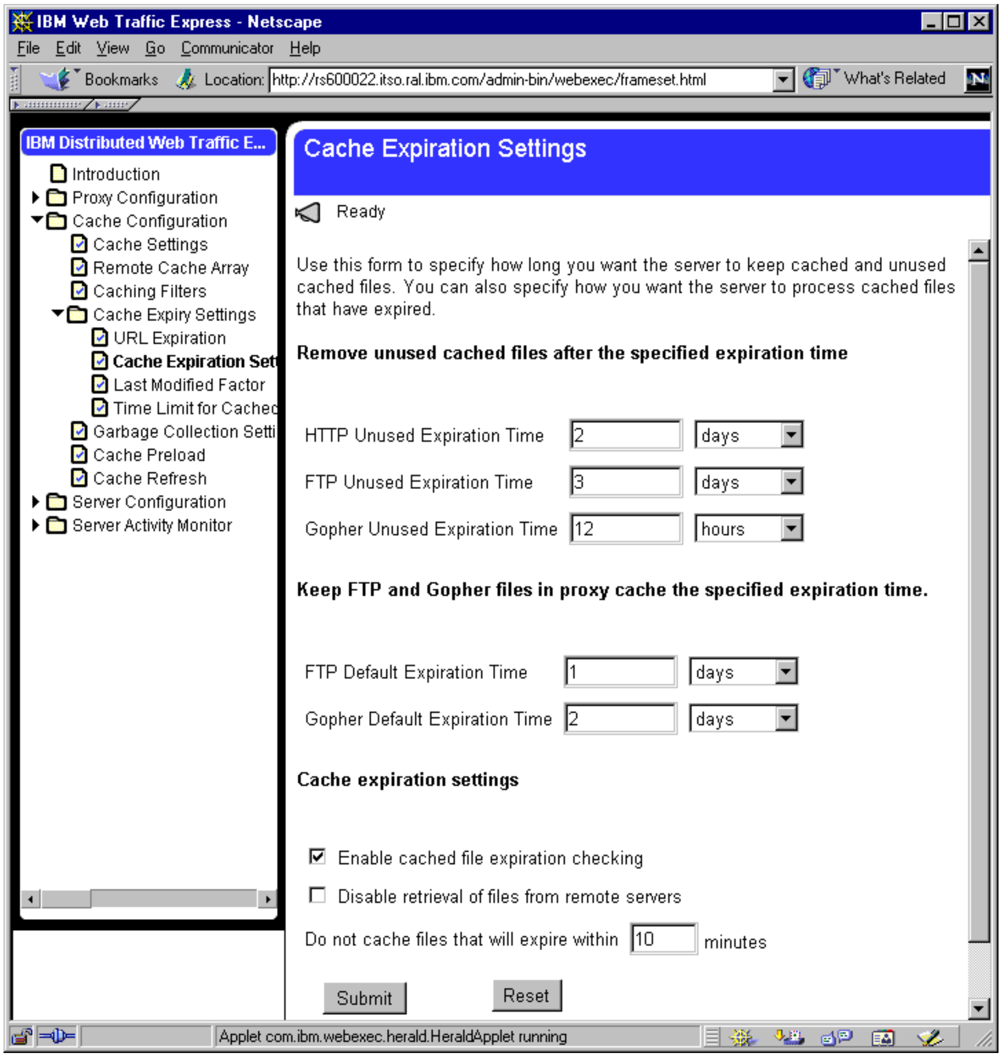

Cache Expiration Settings

Ready

Use this form to specify how long you want the server to keep cached and unused cached files. You can also specify how you want the server to process cached files that have expired.

Remove unused cached files after the specified expiration time

HTTP Unused Expiration Time	2	days
FTP Unused Expiration Time	3	days
Gopher Unused Expiration Time	12	hours

Keep FTP and Gopher files in proxy cache the specified expiration time.

FTP Default Expiration Time	1	days
Gopher Default Expiration Time	2	days

Cache expiration settings

☑ Enable cached file expiration checking

☐ Disable retrieval of files from remote servers

Do not cache files that will expire within 10 minutes

Figure 55. IBM WTE GUI - Cache Expiration Settings

By default, the garbage collection process is enabled, and is performed once a day, at 3:00 a.m. The storage reclamation process allocates 1000 KB of RAM memory. The amount of cache to remain after this process is 75% and the cache algorithm to use is bandwidth. We explain now the meaning of these default settings.

Garbage collection provides three algorithms for choosing how files are removed from the cache. They are:

- **bandwidth** – the algorithm that optimizes network bandwidth
- **responsetime** – the algorithm that optimizes user response time

- **blend** – the algorithm that blends the two and gives you a balance of network bandwidth and user response time

When you are tuning the cache to minimize network bandwidth, larger files are given a lower priority for deletion. They are less likely to get removed during the storage reclamation. When you are tuning the cache to minimize response time, larger files are given a higher priority for deletion and, therefore, are more likely to be removed during garbage collection.

We changed the Maximum memory allocated to `2000 KB` as the garbage collection process is best performed if it can read all cache information into memory. We also set the Cache algorithm to use to **blend** because it gives you a balance of network bandwidth and user response time.

To access the garbage collection settings form, click the **Cache Configuration** folder in the navigation frame. Then, select **Garbage Collection Settings**, as shown in the following figure:

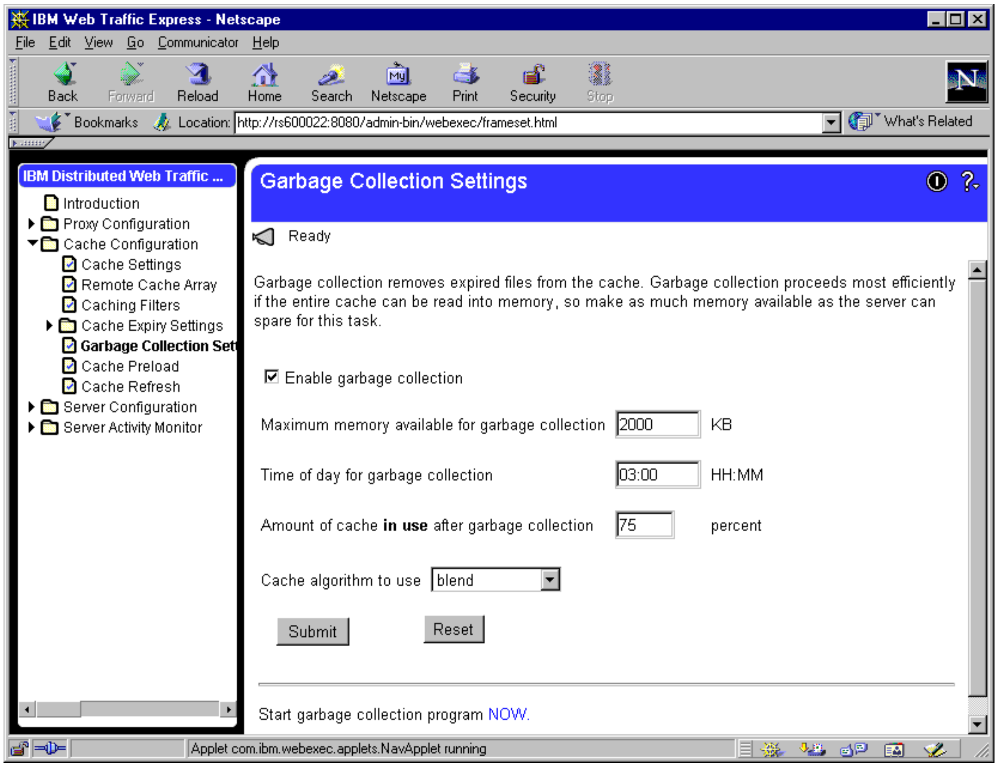
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022:8080/admin-bin/webexec/frameset.html
What's Related
IBM Distributed Web Traffic ...
Introduction
Proxy Configuration
Cache Configuration
Cache Settings
Remote Cache Array
Caching Filters
Cache Expiry Settings
Garbage Collection Sett
Cache Preload
Cache Refresh
Server Configuration
Server Activity Monitor
Applet com.ibm.webexec.applets.NavApplet running

Garbage Collection Settings

Ready

Garbage collection removes expired files from the cache. Garbage collection proceeds most efficiently if the entire cache can be read into memory, so make as much memory available as the server can spare for this task.

☑ Enable garbage collection

Maximum memory available for garbage collection 2000 KB

Time of day for garbage collection 03:00 HH:MM

Amount of cache **in use** after garbage collection 75 percent

Cache algorithm to use blend

Submit Reset

Start garbage collection program NOW.

Figure 56. IBM WTE GUI - Garbage Collection Settings

The garbage collection is a resource intensive process. Hence, it is recommended that it be set to run when the load of requests is estimated to be the lowest. In our case, we kept its default value, 3:00 a.m.

Using the Directives

The same configuration could have been issued by directly editing the ibmproxy.conf configuration file. In particular, we would have used the following directives:

```
Gc On
GcMemUsage 2000
GcDailyGC 03:00
GcMaxInUse 75
CacheAlogorithm blend
```

After making changes via the WTE GUI, in order for the configuration changes to be written in the configuration file, click **Submit** on the Garbage Collection Settings form. You will get the Confirmation page. Click **Continue** to proceed.

For the changes to take effect, you have to restart the WTE server. Click the **Restart** icon on the WTE GUI to restart the httpd daemon. This operation will re-read the ibmproxy.conf file.

Note on Clicking the Restart Icon

Clicking the **Restart** icon will stop and restart the httpd daemon. However, be aware that no confirmation screen is provided to inform you that the restart operation took place successfully.

Chapter 5. WTE Performance and Status

In this chapter, we show you how to use the Server Activity Monitor of IBM Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2. This tool monitors statistics and access log entries, and uses the statistics to assist the WTE administrator in the tuning of WTE for performance improvements.

We also describe how the data collected from the Server Activity Monitor can be used to tune WTE for better performance.

5.1 WTE Server Activity Monitor

The Server Activity Monitor enables you to display:

- Activity statistics
- Network statistics
- Access statistics
- Proxy access statistics
- Cache satistics
- Garbage collection summary
- Cache refresh summary

5.1.1 Activity Statistics

The Activity Statistics page provides information on the number of threads (connections) used and available, server response time, throughput, number of requests processed and also number of errors. To view the Activity Statistics page, click the **Server Activity Statistics** folder and select **Activity Statistics**, as shown in Figure 65 on page 108. Click **Refresh** to update the information.

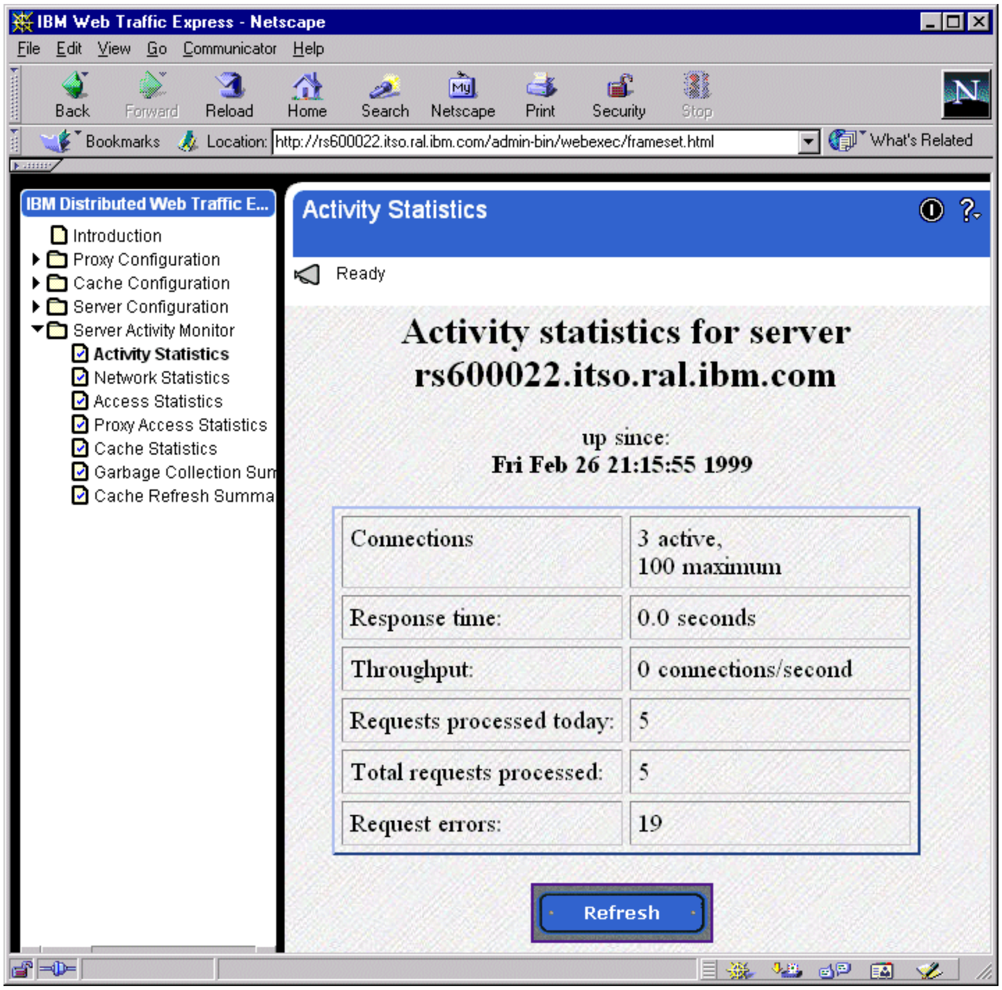

Activity statistics for server rs600022.itso.ral.ibm.com

up since:
Fri Feb 26 21:15:55 1999

Connections	3 active, 100 maximum
Response time:	0.0 seconds
Throughput:	0 connections/second
Requests processed today:	5
Total requests processed:	5
Request errors:	19

Figure 57. IBM WTE GUI - Activity Statistics

5.1.2 Network Statistics

The Network Statistics page provides information about the network on which the proxy is running, such as data rate and bytes sent and received. To view the Network Statistics, click the **Server Activity Monitor** folder and select **Network Statistics**, as shown in Figure 58. Click **Refresh** to update the information.

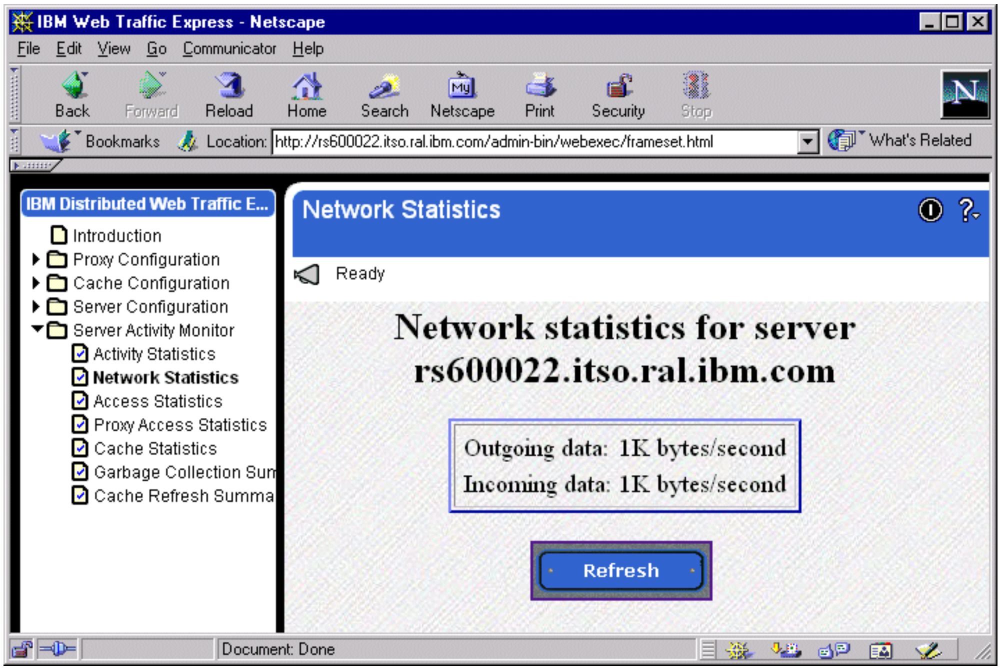

Network statistics for server rs600022.itso.ral.ibm.com

Outgoing data: 1K bytes/second
Incoming data: 1K bytes/second

Figure 58. IBM WTE GUI -- Network Statistics

5.1.3 Access Statistics

The Access Statistics page provides information found in the access log files, including users accessing your proxy, IP addresses, user IDs (for protected pages), date and time of access, and the HTTP method performed.[1] To view the Access Statistics, click the **Server Activity Monitor** folder and select **Access Statistics**, as shown in Figure 59. Click **Refresh** to update the information.

[1] HTTP methods are GET, PUT, POST, and DELETE.

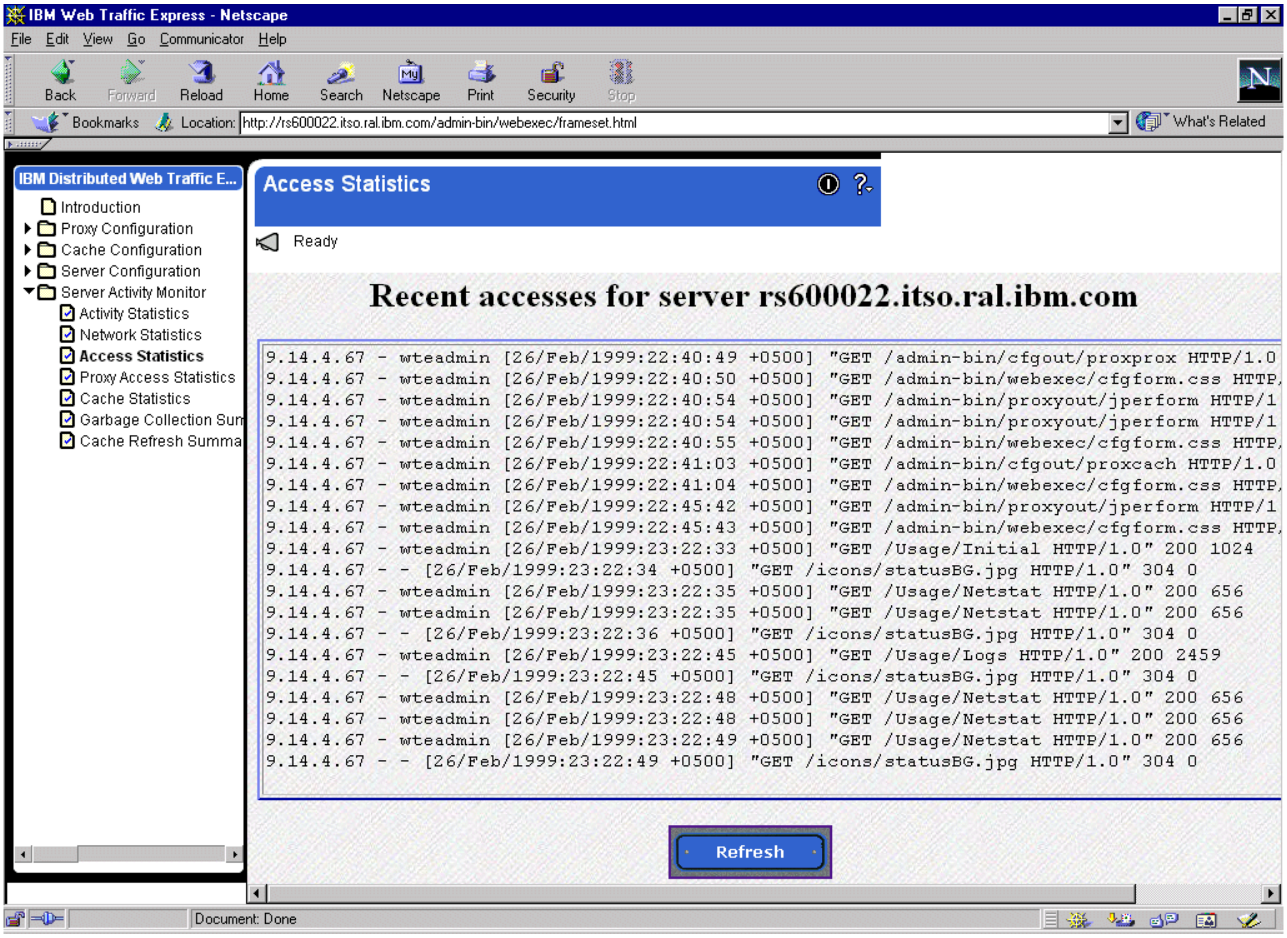

Recent accesses for server rs600022.itso.ral.ibm.com

```
9.14.4.67 - wteadmin [26/Feb/1999:22:40:49 +0500] "GET /admin-bin/cfgout/proxprox HTTP/1.0
9.14.4.67 - wteadmin [26/Feb/1999:22:40:50 +0500] "GET /admin-bin/webexec/cfgform.css HTTP/
9.14.4.67 - wteadmin [26/Feb/1999:22:40:54 +0500] "GET /admin-bin/proxyout/jperform HTTP/1
9.14.4.67 - wteadmin [26/Feb/1999:22:40:54 +0500] "GET /admin-bin/proxyout/jperform HTTP/1
9.14.4.67 - wteadmin [26/Feb/1999:22:40:55 +0500] "GET /admin-bin/webexec/cfgform.css HTTP/
9.14.4.67 - wteadmin [26/Feb/1999:22:41:03 +0500] "GET /admin-bin/cfgout/proxcach HTTP/1.0
9.14.4.67 - wteadmin [26/Feb/1999:22:41:04 +0500] "GET /admin-bin/webexec/cfgform.css HTTP/
9.14.4.67 - wteadmin [26/Feb/1999:22:45:42 +0500] "GET /admin-bin/proxyout/jperform HTTP/1
9.14.4.67 - wteadmin [26/Feb/1999:22:45:43 +0500] "GET /admin-bin/webexec/cfgform.css HTTP/
9.14.4.67 - wteadmin [26/Feb/1999:23:22:33 +0500] "GET /Usage/Initial HTTP/1.0" 200 1024
9.14.4.67 - - [26/Feb/1999:23:22:34 +0500] "GET /icons/statusBG.jpg HTTP/1.0" 304 0
9.14.4.67 - wteadmin [26/Feb/1999:23:22:35 +0500] "GET /Usage/Netstat HTTP/1.0" 200 656
9.14.4.67 - wteadmin [26/Feb/1999:23:22:35 +0500] "GET /Usage/Netstat HTTP/1.0" 200 656
9.14.4.67 - - [26/Feb/1999:23:22:36 +0500] "GET /icons/statusBG.jpg HTTP/1.0" 304 0
9.14.4.67 - wteadmin [26/Feb/1999:23:22:45 +0500] "GET /Usage/Logs HTTP/1.0" 200 2459
9.14.4.67 - - [26/Feb/1999:23:22:45 +0500] "GET /icons/statusBG.jpg HTTP/1.0" 304 0
9.14.4.67 - wteadmin [26/Feb/1999:23:22:48 +0500] "GET /Usage/Netstat HTTP/1.0" 200 656
9.14.4.67 - wteadmin [26/Feb/1999:23:22:48 +0500] "GET /Usage/Netstat HTTP/1.0" 200 656
9.14.4.67 - wteadmin [26/Feb/1999:23:22:49 +0500] "GET /Usage/Netstat HTTP/1.0" 200 656
9.14.4.67 - - [26/Feb/1999:23:22:49 +0500] "GET /icons/statusBG.jpg HTTP/1.0" 304 0
```

Figure 59. IBM WTE GUI - Access Statistics

5.1.4 Proxy Access Statistics

The Proxy Access Statistics page provides information on the proxy activity such as which URLs were requested and if they were retrieved from the cache. Following the URLs are the return codes given to the client, and file size in bytes. To view the Proxy Access Statistics, click the **Server Activity Monitor** folder and select **Proxy Access Statistics**, as shown in Figure 60. Click **Refresh** to update the information.

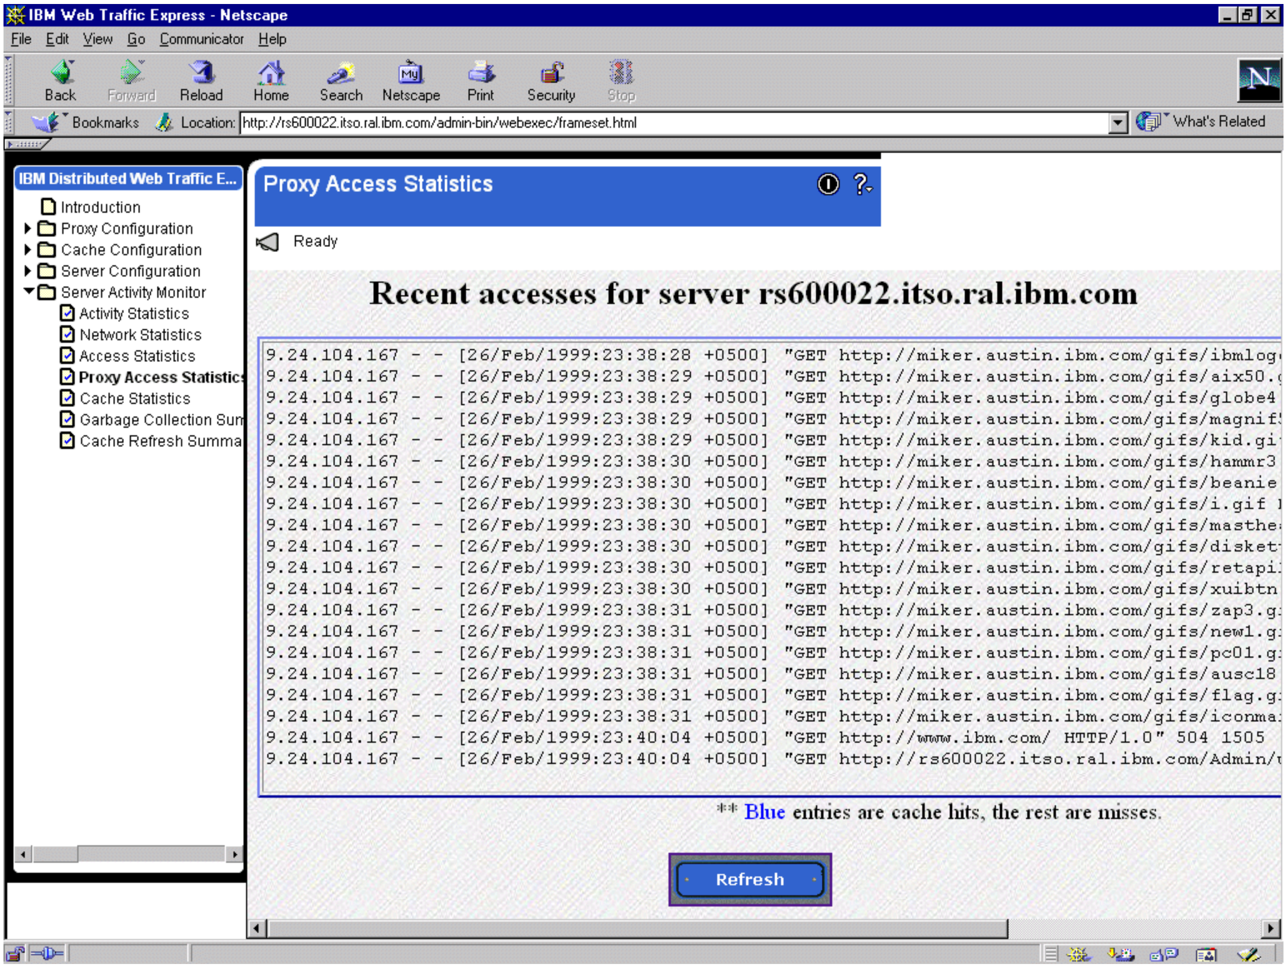

Recent accesses for server rs600022.itso.ral.ibm.com

```
9.24.104.167 - - [26/Feb/1999:23:38:28 +0500] "GET http://miker.austin.ibm.com/gifs/ibmlog
9.24.104.167 - - [26/Feb/1999:23:38:29 +0500] "GET http://miker.austin.ibm.com/gifs/aix50.
9.24.104.167 - - [26/Feb/1999:23:38:29 +0500] "GET http://miker.austin.ibm.com/gifs/globe4
9.24.104.167 - - [26/Feb/1999:23:38:29 +0500] "GET http://miker.austin.ibm.com/gifs/magnif
9.24.104.167 - - [26/Feb/1999:23:38:29 +0500] "GET http://miker.austin.ibm.com/gifs/kid.gi
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/hammr3
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/beanie
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/i.gif
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/masthe
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/disket
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/retapi
9.24.104.167 - - [26/Feb/1999:23:38:30 +0500] "GET http://miker.austin.ibm.com/gifs/xuibtn
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/zap3.g
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/new1.g
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/pc01.g
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/ausc18
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/flag.g
9.24.104.167 - - [26/Feb/1999:23:38:31 +0500] "GET http://miker.austin.ibm.com/gifs/iconma
9.24.104.167 - - [26/Feb/1999:23:40:04 +0500] "GET http://www.ibm.com/ HTTP/1.0" 504 1505
9.24.104.167 - - [26/Feb/1999:23:40:04 +0500] "GET http://rs600022.itso.ral.ibm.com/Admin/
```

** Blue entries are cache hits, the rest are misses.

Figure 60. IBM WTE GUI - Proxy Access Statistics

5.1.5 Cache Statistics

The Cache Statistics page provides the status of the proxy cache, such as whether the cache is currently operational, still being reindexed from a server start, or if garbage collection is running. For example, if the cache is currently operational, the Cache Statistics page will display the message:

```
Cache operational
```

This page also offers a table showing the separate subcaches, their current size, the number of files in each subcache, and the size of the index for the subcache. It also informs you whether or not any of the subcaches are full. For example, if none of the subcaches are full, you will see the following message:

```
None of the subcaches are full
```

To view the Cache Statistics page, click the **Server Activity Monitor** folder and select **Cache Statistics**, as shown in Figure 61. Click **Refresh** to update the information.

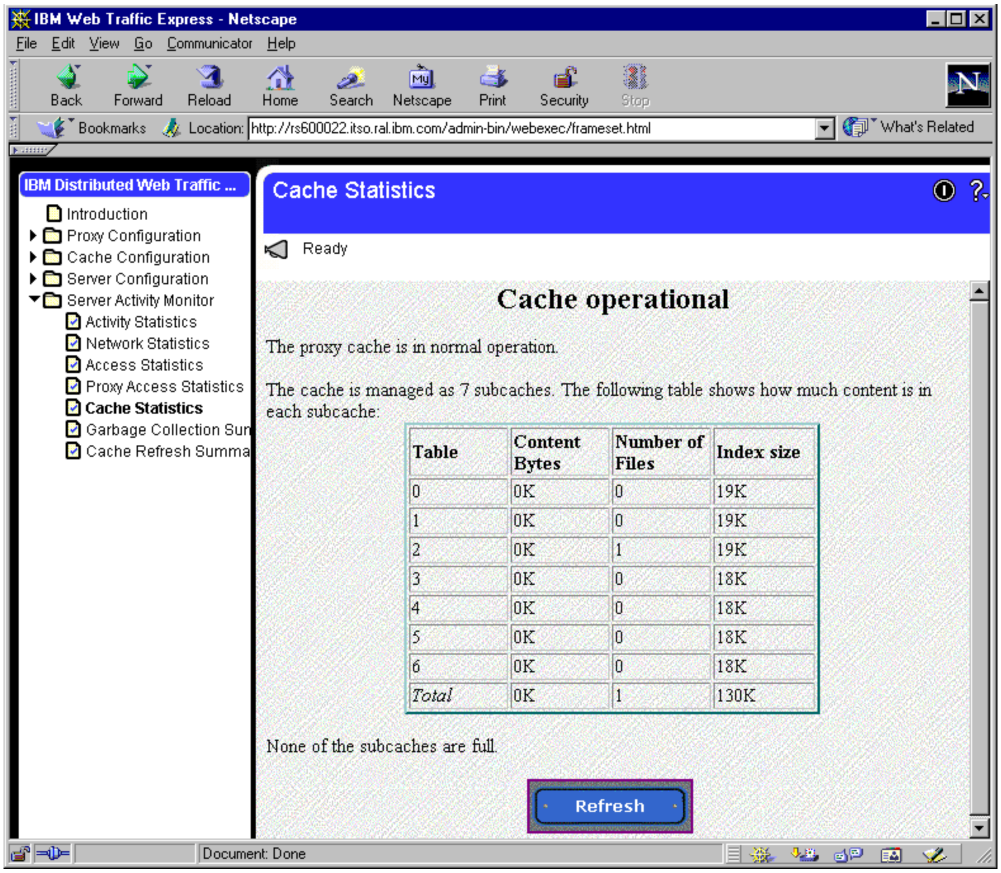

Cache operational

The proxy cache is in normal operation.

The cache is managed as 7 subcaches. The following table shows how much content is in each subcache:

Table	Content Bytes	Number of Files	Index size
0	0K	0	19K
1	0K	0	19K
2	0K	1	19K
3	0K	0	18K
4	0K	0	18K
5	0K	0	18K
6	0K	0	18K
Total	0K	1	130K

None of the subcaches are full.

Figure 61. IBM WTE GUI - Cache Statistics

5.1.6 Garbage Collection Summary

The Garbage Collection Summary page provides information on:

- Starting time of the last garbage collection
- Ending time of the last garbage collection
- Number of files, directories, and bytes in the cache after garbage collection last ran
- The size of the cache as a percentage of the maximum cache size
- Number of files, directories, and bytes removed during garbage collection
- Memory used during garbage collection

Garbage collection must have been run for this page to display any information.

To view the Garbage Collection Summary page, click the **Server Activity Monitor** folder and select **Garbage Collection Summary**, as shown in Figure 62. Click **Refresh** to update the information.

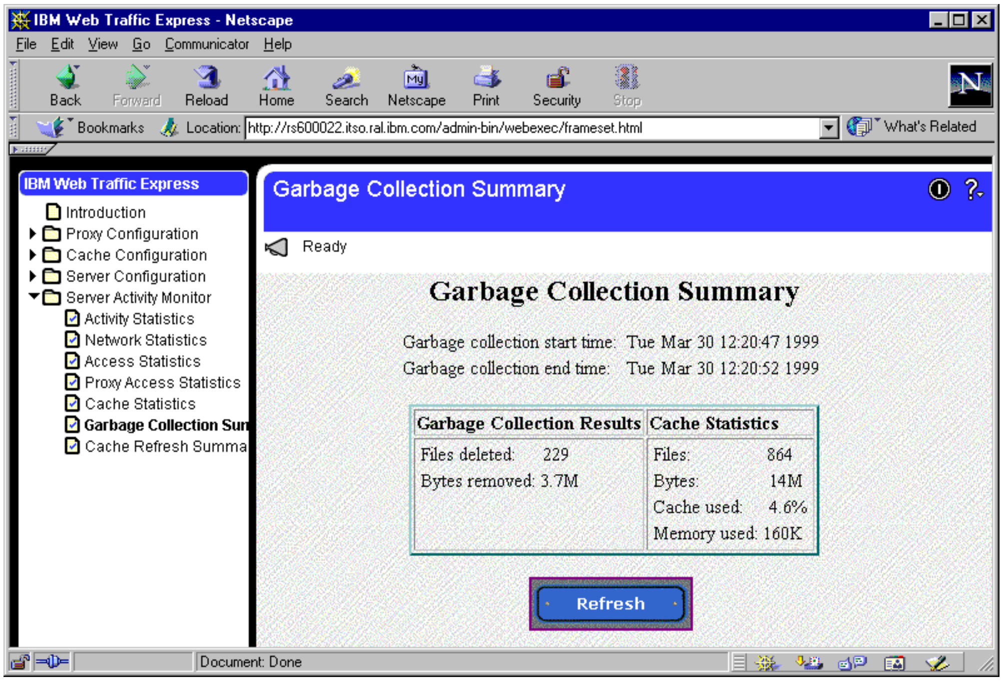

Garbage Collection Summary

Garbage collection start time: Tue Mar 30 12:20:47 1999

Garbage collection end time: Tue Mar 30 12:20:52 1999

Garbage Collection Results	Cache Statistics
Files deleted: 229 Bytes removed: 3.7M	Files: 864 Bytes: 14M Cache used: 4.6% Memory used: 160K

Figure 62. IBM WTE GUI - Garbage Collection Summary

5.1.7 Cache Refresh Summary

The Cache Refresh Summary provides information on:

- Starting time of the last cache refresh
- Ending time of the last cache refresh
- Number of URLs specified in the configuration file that were refreshed
- Number of pages refreshed during the previous night’s cache access log
- Number of URLs the cache agent refreshed
- Number of URLs remaining in the queue after reaching the maximum allowed number of URLs

The cache agent must have run at least once to display any information.

To view the Cache Statistics page, click the **Server Activity Monitor** folder and select **Cache Statistics**, as shown in Figure 63. Click **Refresh** to update the information.

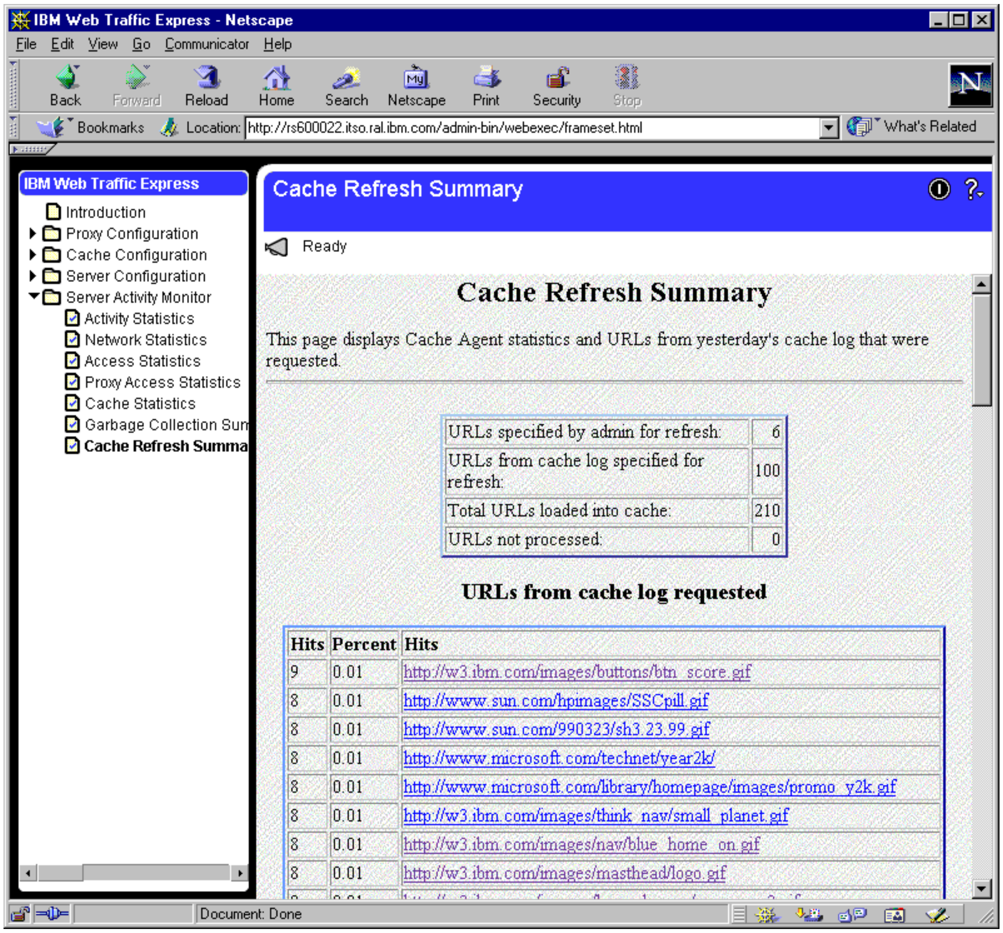

Cache Refresh Summary

This page displays Cache Agent statistics and URLs from yesterday's cache log that were requested.

URLs specified by admin for refresh:	6
URLs from cache log specified for refresh:	100
Total URLs loaded into cache:	210
URLs not processed:	0

URLs from cache log requested

Hits	Percent	Hits
9	0.01	http://w3.ibm.com/images/buttons/btn_score.gif
8	0.01	http://www.sun.com/hpimages/SSCpill.gif
8	0.01	http://www.sun.com/990323/sh3.23.99.gif
8	0.01	http://www.microsoft.com/technet/year2k/
8	0.01	http://www.microsoft.com/library/homepage/images/promo_y2k.gif
8	0.01	http://w3.ibm.com/images/think_nav/small_planet.gif
8	0.01	http://w3.ibm.com/images/nav/blue_home_on.gif
8	0.01	http://w3.ibm.com/images/masthead/logo.gif

Figure 63. IBM WTE GUI - Cache Refresh Summary

5.2 Configuration for Performance Improvements

This section describes the WTE configuration settings you can modify to tune your server.

5.2.1 Understanding Performance Settings

Each time your server receives a request from a client, it uses a thread to perform the requested action. If no threads are available, the server holds the request until more threads become available. The `MaxActiveThreads` directive in the WTE configuration file specifies the maximum number of active threads. Clients will experience slower response times or even failures if no active threads are available. The failure would be shown as an error such as:

```
Error 400 - Proxy Error:
Unable to connect to remote host or host not responding
```

If the number of active connections is very low, you can lower the amount of virtual memory used by lowering the `MaxActiveThreads` setting. To change the active threads started by WTE, click the **Server Configuration** folder, then the **System Management** folder, and select the **Performance** form. Click the **Submit** button to update the information after making any changes. The Performance form is shown in the following figure:

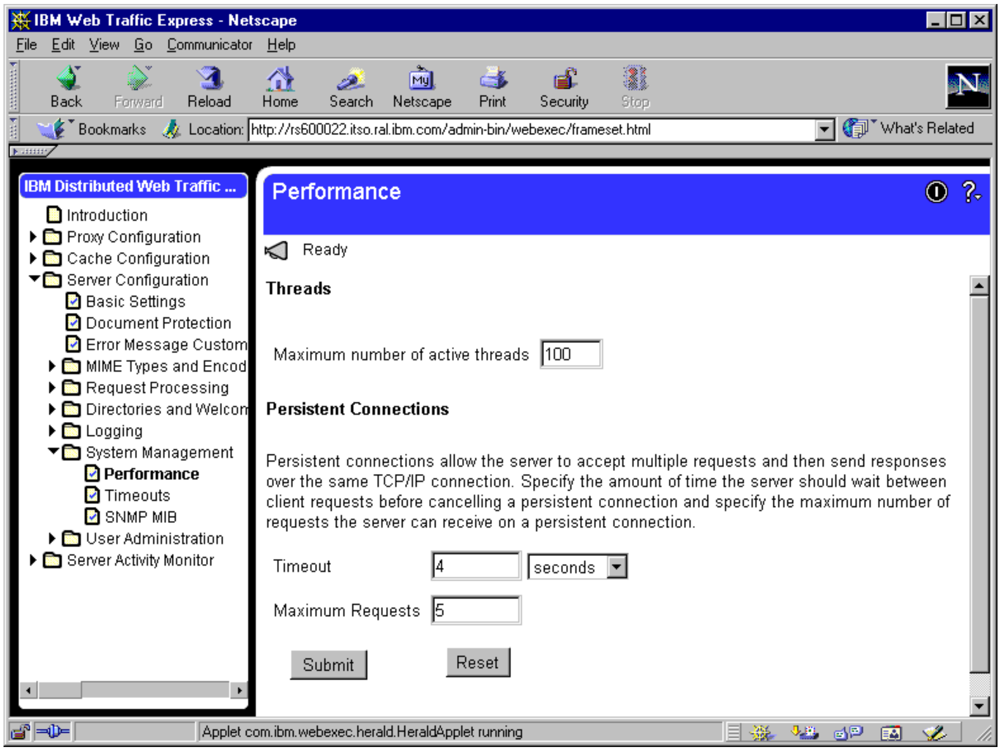

Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html

Performance

Ready

Threads

Maximum number of active threads: 100

Persistent Connections

Persistent connections allow the server to accept multiple requests and then send responses over the same TCP/IP connection. Specify the amount of time the server should wait between client requests before cancelling a persistent connection and specify the maximum number of requests the server can receive on a persistent connection.

Timeout: 4 seconds

Maximum Requests: 5

Applet com.ibm.webexec.herald.HeraldApplet running

Figure 64. IBM WTE GUI - Performance

The following factors may impact the server response time:

- Network speed
- Traffic on the local area network (LAN)
- Number of clients requesting from your server
- Number of threads set on your server
- Other applications running on your server
- System resources

Change the values of the Persistent Connections Timeout and Maximum Requests fields to specify the characteristics of a persistent connection. A persistent connection allows the server to accept multiple requests and to send responses over the same TCP/IP connection. By increasing the number of maximum requests, overall throughput will be increased as well because the server will not have to establish a separate TCP/IP connection for each request and response. Also, the TCP/IP connection is used more efficiently, due to the characteristics of TCP/IP. However, keep in mind that persistent connections require network bandwidth as well as a dedicated server thread.

Persistent timeout specifies the amount of time the server will wait between client requests before cancelling a persistent connection.

To view the active threads started by WTE, click the **Server Activity Statistics** folder and select **Activity Statistics**. Click **Refresh** to update the information. The following figure shows the Activity Statistics of the WTE:

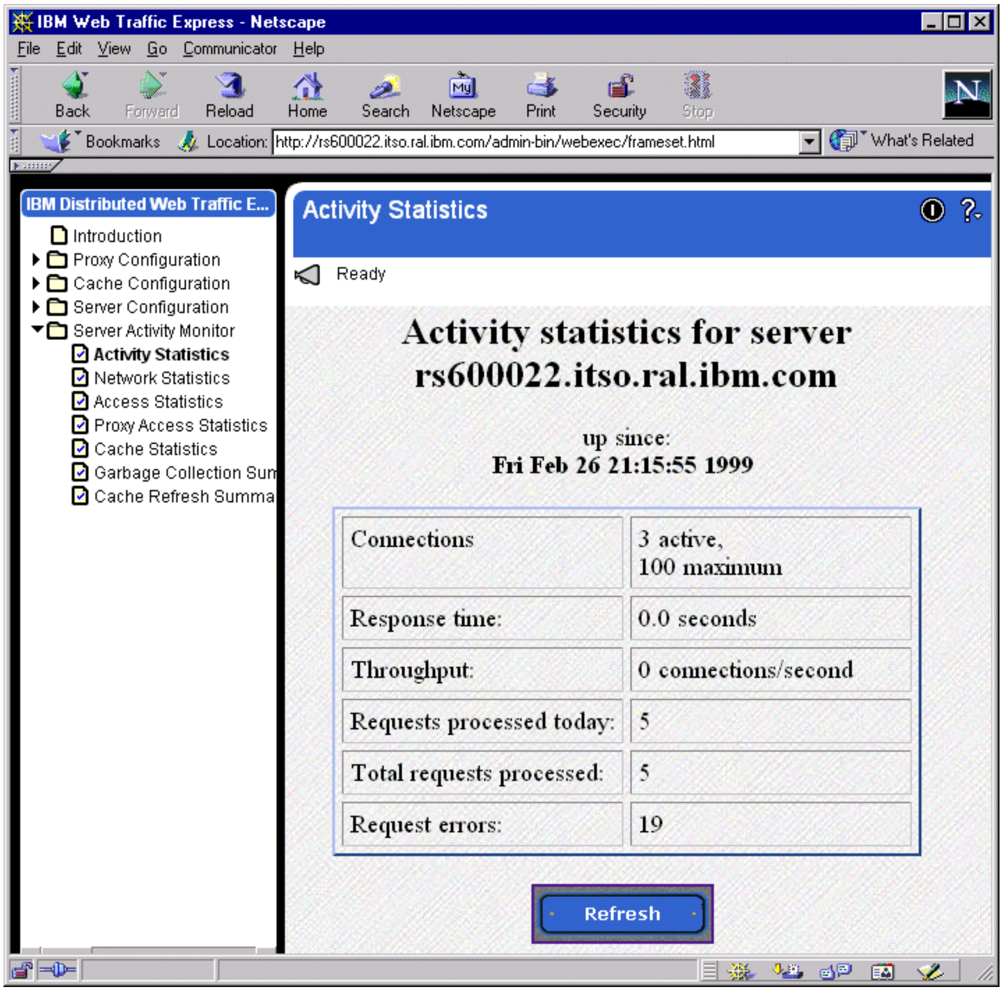

Activity statistics for server rs600022.itso.ral.ibm.com

up since:
Fri Feb 26 21:15:55 1999

Connections	3 active, 100 maximum
Response time:	0.0 seconds
Throughput:	0 connections/second
Requests processed today:	5
Total requests processed:	5
Request errors:	19

Figure 65. IBM WTE GUI - Activity Statistics

5.2.2 Timeouts – Closing Connections Automatically

Timeouts are used to control the server processing time. The default values are appropriate for most requests and need not be changed unless necessary.

To change the timeout values of the WTE, click the **Server Configuration folde**r, followed by the **System Management** folder and select the **Timeouts** form. Click the **Submit** button to update the information after making any changes. The Timeouts form is shown in the following figure:

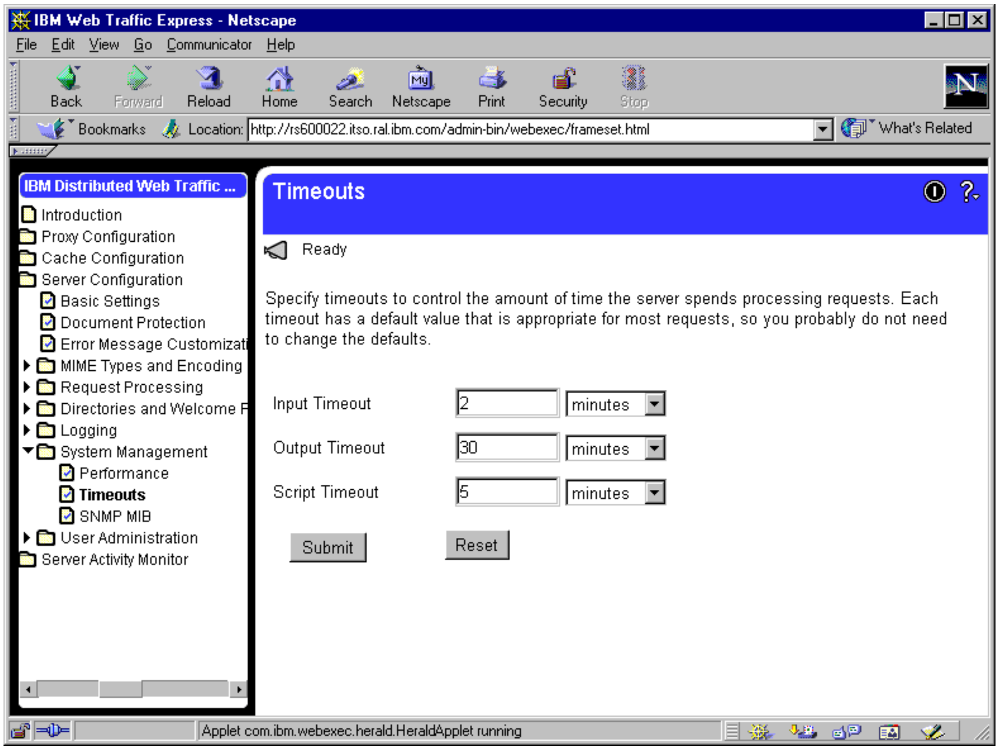

Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html

Timeouts

Ready

Specify timeouts to control the amount of time the server spends processing requests. Each timeout has a default value that is appropriate for most requests, so you probably do not need to change the defaults.

Input Timeout	2	minutes
Output Timeout	30	minutes
Script Timeout	5	minutes

Applet com.ibm.webexec.herald.HeraldApplet running

Figure 66. IBM WTE GUI - Timeouts

Notice that:

- The Input Timeout field is used to set the time allowed for a client to send a request after making a connection to the server. A client first connects to the server and then sends a request. If the client does not send a request within the amount of time specified by this directive, the server drops the connection. Take note that, if you are using persistent connections, the persistent connection timeout specifies the time to wait for the client to send another request (see Figure 64 on page 107).
- The Output Timeout is used to set the maximum time allowed for your server to send output to a client. The time limit applies to requests for local files and requests for which the server is acting as a proxy. However, it does not apply to requests that start a local CGI-BIN program. If the server does not send the complete response within the amount of time specified by this directive, the server drops the connection.
- The Script Timeout is used to set the time allowed for a CGI-BIN program started by the server to finish. When the time runs out, the server ends the program.

5.2.3 Configuring for Proxy Performance

Click the **Proxy Configuration** folder in the navigation frame, then select **Proxy Performance**, as shown in Figure 67.

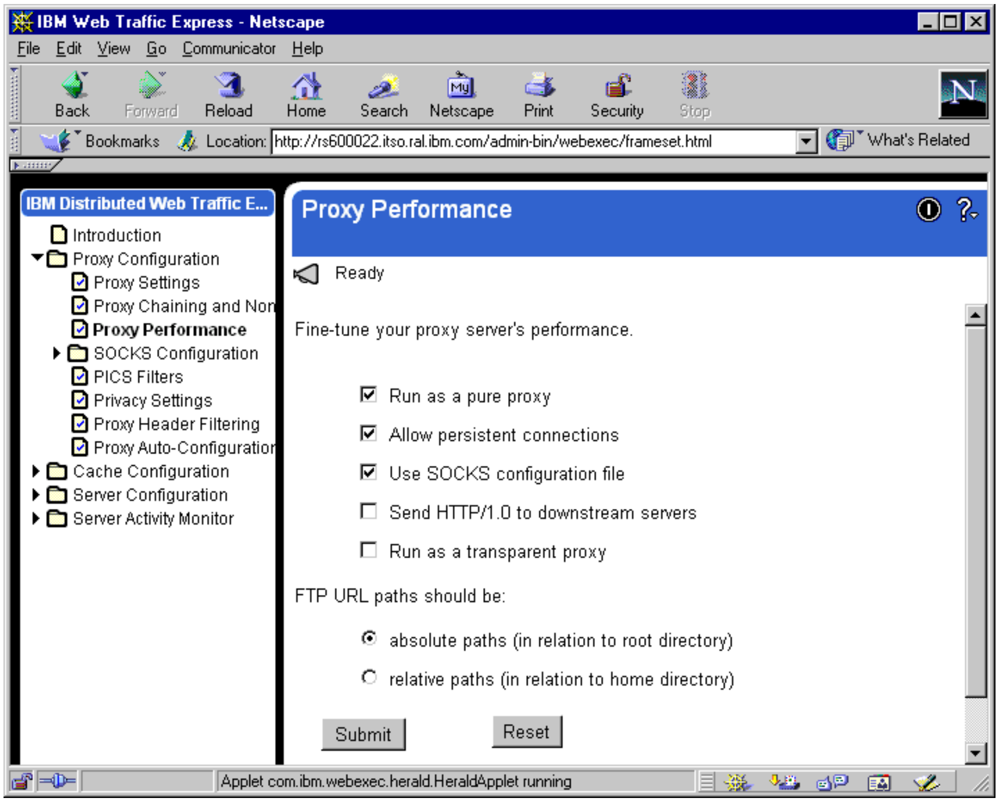
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Distributed Web Traffic E...
Introduction
Proxy Configuration
Proxy Settings
Proxy Chaining and Non
Proxy Performance
SOCKS Configuration
PICS Filters
Privacy Settings
Proxy Header Filtering
Proxy Auto-Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
Proxy Performance
Ready
Submit
Reset
Applet com.ibm.webexec.herald.HeraldApplet running

Fine-tune your proxy server's performance.

☑ Run as a pure proxy

☑ Allow persistent connections

☑ Use SOCKS configuration file

□ Send HTTP/1.0 to downstream servers

□ Run as a transparent proxy

FTP URL paths should be:

● absolute paths (in relation to root directory)

○ relative paths (in relation to home directory)

Figure 67. IBM WTE GUI - Proxy Performance

Fill in the Proxy Performance form by clicking the following fields:

- **Run as a pure proxy**

  Mark this box if you want to increase proxy performance by strictly running a proxy server. The alternative would be to run WTE as a content server too. We recommend that you run WTE as a pure proxy server.

- **Allow persistent connections**

  Use this setting to allow clients to force the server to keep an open connection with them. This decreases client's log time associated with requesting documents from the proxy. However, persistent connections require network bandwidth as well as a dedicated server thread to maintain. Do not allow persistent connections if your setup limits the number of available threads.

- **Use a SOCKS configuration file**

  Check this box if you want the proxy to look at the SOCKS configuration file to decide whether or not to connect through the SOCKS server.

- **Send HTTP/1.0 to downstream servers**

  Check this if you have downstream servers which do not correctly handle requests from HTTP/1.1 clients.

- **Run as a transparent proxy**

  Check this box if you want your WTE server to run as a transparent proxy. For more details on this, see Chapter 7, “Using WTE as a Transparent Proxy Server” on page 131.

To Identify how FTP URL paths should be resolved, check either:

- **absolute paths**

  This option resolves to fully-specified paths with respect to the root directory.

- **relative paths**

  This option resolves with respect to the home directory.

After you have completed your selections, click **Submit** to make the changes to the configuration file.

5.3 Using the Server Activity Monitor to Tune WTE

In 5.1, "WTE Server Activity Monitor" on page 99 we showed how to use the Server Activity Monitor to collect data about a WTE server's activities. In this section we show you how you can use the data collected with the Server Activity Monitor as an input to change the configuration of your WTE server, and improve its performance.

5.3.1 Activity Statistics

As we said in 5.1.1, "Activity Statistics" on page 99, the Activity Statistics page provides basic information on thread and connection activity, request statistics, and server response time (see Figure 65 on page 108).

By changing the following directives in the proxy configuration file, ibmproxy.conf, you can customize the server for your needs:

- `MaxActiveThreads`

  This directive specifies how many threads are in the server pool. You can increase or decrease the number of threads available depending on how much memory your server has. This directive will have an effect on many of the statistics on this page.

- `ProxyPersistence`

  This directive specifies whether the proxy will allow persistent connections from a client. This directive may have an effect on network throughput, depending on usage.

5.3.2 Network Statistics

As we said in 5.1.2, "Network Statistics" on page 100, the Network Statistics page provides information about the network the proxy is running on, such as data rate and bytes sent and received (see Figure 58 on page 101).

If your proxy server accesses the Internet through a SOCKS server, consider using the flexible-client SOCKS feature of WTE to improve performance. Refer to 2.6, "Flexible-Client SOCKS" on page 41 for more details on the flexible-client SOCKS feature of WTE.

By changing the `FlexibleSocks` directives in the WTE proxy configuration file, ibmproxy.conf, you can customize the server for your needs. This directive specifies whether the proxy will use the standard SOCKS configuration file when retrieving requests. Using flexible-client SOCKS reduces request latency and interruptions to the SOCKS server.

In the SOCKS configuration file, socks.conf, specify the hosts to connect to directly, the hosts to pass to the SOCKS server, and the hosts to which connection is denied.

Refer to Chapter 9, “WTE Flexible-Client SOCKS Support” on page 155 for more information on additional directives you can modify to further enhance performance when a flexible-client SOCKS configuration is in place.

5.3.3 Access Statistics

As we explained in 5.1.3, “Access Statistics” on page 101, the Access Statistics page provides information found in the access log files, including users accessing your proxy, IP addresses, user IDs (if pages are protected), date and time of access, and the HTTP method performed. Refer to 5.4, “Customizing Logs” on page 115 for details on how WTE logging activities can be customized.

5.3.4 Proxy Access Statistics

As we said in 5.1.4, “Proxy Access Statistics” on page 102, the Proxy Access Statistics page provides information on the proxy activity, such as which URLs were requested and if they were retrieved from the cache. Following the URLs are the return codes given to the client, and file size in bytes. By changing the following directives in the proxy configuration file, ibmproxy.conf, you can increase the number of cache hits:

- `AutoCacheRefresh`

  This directive turns the automatic mode of the cache agent on or off. The automatic mode will refresh a number of URLs from the previous night's access log. If clients are requesting many of the same documents every day, this will increase cache hits.

- `CacheMinHold`

  This directive allows you to override the expiry information in the header of URLs. Some sites routinely force documents to immediately expire when they actually have a longer lifetime. You can specify the URL mask for the URL to override and the time to keep the file. However, notice that when you override the expiry information, *clients could receive stale data*.

- `PureProxy`

  This directive determines whether the server is just a proxy, or both a proxy and a content server. As we already said several times, we recommend that you use WTE as a pure proxy server only, and not as a content server. This will increase your WTE server's performance. As a matter of fact, the `PureProxy` directive is by default set to `On`.

- `CacheLocalDomain`

  WTE supports persistent connections, where a separate thread is used to keep each client connection open to the proxy. In most cases, loading documents that reside on your intranet is faster than loading documents that reside on the Internet. Because of this, you might choose not to cache documents from your internal Web servers. To instruct the proxy to ignore sites on your domain when refreshing the cache, set `CacheLocalDomain` to `Off`.

5.3.5 Cache Statistics

As we explained in 5.1.5, “Cache Statistics” on page 103, the Cache Statistics page provides the following information on the cache and index:

- Whether the cache is currently operational, still being reindexed from a server start, or if garbage collection is running

  For example, if the cache is currently operational, you will see the following confirmation message on the Cache Statistics page (see Figure 61 on page 104):

  ```
  Cache operational
  ```

- A table showing the separate subcaches, their current size, the number of files in the subcache, and the size of the index for the subcache
- Whether any of the subcaches are full

  For example, a possible confirmation message you might see on this page is the following:

  ```
  None of the subcaches are full
  ```

In this case, tuning can be achieved by changing the following directives in the proxy configuration file, ibmproxy.conf:

- `Caching`

  This directive is used to activate or deactivate caching by setting its value to `On` or `Off` respectively.

- `CacheSize`

  This directive sets the maximum cache size in drive space.

- `CacheFiles`

  This directive sets the maximum cache size in number of files.

5.3.6 Garbage Collection Summary

As we said in 5.1.6, “Garbage Collection Summary” on page 104, the Garbage Collection Summary page provides the following information:

- Starting time of the last garbage collection
- Ending time of the last garbage collection
- Number of files, directories, and bytes in the cache after garbage collection last ran
- The size of the cache as a percentage of the maximum cache size
- Number of files, directories, and bytes removed during garbage collection
- Memory used during garbage collection

A sample Garbage Collection Summary report is shown in the following figure:

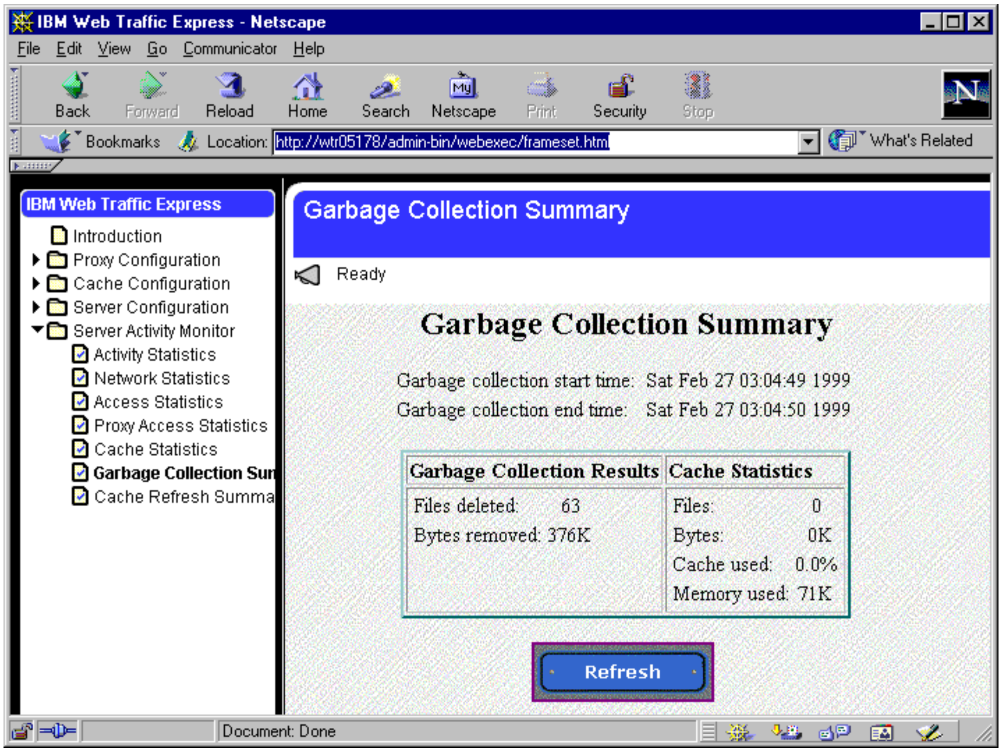

Location: http://wtr05178/admin-bin/webexec/frameset.html

Garbage Collection Summary

Garbage collection start time: Sat Feb 27 03:04:49 1999
Garbage collection end time: Sat Feb 27 03:04:50 1999

Garbage Collection Results	Cache Statistics
Files deleted: 63 Bytes removed: 376K	Files: 0 Bytes: 0K Cache used: 0.0% Memory used: 71K

Figure 68. IBM WTE GUI - Garbage Collection Summary

Garbage collection must have been run for this page to display any information. By changing the following directives in the proxy configuration file, ibmproxy.conf, you can customize the garbage collection to meet your requirements:

- `GCDailyGC`

  This directive sets the time garbage collection begins. This is set in local time using a 24-hour clock. For example,16:30 is 4:30 p.m.

- `GCMemUsage`

  This directive specifies the amount of memory (in KB) the garbage collection process uses. Smaller values are less efficient, but larger values require more resources. Therefore, average values are recommended.

- `CacheMinHold`

  This directive, which we discussed in 5.1.4, “Proxy Access Statistics” on page 102, allows you to override the expiry information in the header of URLs. Some sites routinely force documents to immediately expire when they actually have a longer lifetime. You can specify the URL mask for the URL to override and the time to keep the file. Notice that, when you override the expiry information, *clients could receive stale data*.

- `GCMaxInUse`

  This directive determines the maximum cache utilization after garbage collection has completed.

5.3.7 Cache Refresh Summary

As we said in 5.1.7, “Cache Refresh Summary” on page 105, the Cache Refresh Summary page provides information on the cache agent's last run. The cache agent must have run at least once to display any information. By changing the following directives in the proxy configuration file, ibmproxy.conf, you can customize the server to better meet your requirements:

- `CacheLocalDomain`

  This directive, discussed in 5.3.4, “Proxy Access Statistics” on page 112, instructs the proxy whether to cache pages on the local domain or not. If most of the traffic on your intranet is to local sites, you should set `CacheLocalDomain` to `On`.

- `LoadURL`

  This directive specifies a URL the cache agent will refresh on the next refresh run. If many clients are requesting the same page, but for some reason are not included in the access log, you can manually set the URL to be refreshed. Multiple `LoadURL` directives are allowed.

- `LoadTopCached`

  This directive specifies the number of most popular URLs to load using automatic cache refreshing.

- `MaxRunTime`

  This directive specifies the maximum time the cache agent is allowed to run. A value of `0` would allow the cache agent any amount of time to complete. Cache refreshing is memory-intensive. Therefore, if you are refreshing a very large cache, you may wish to limit the time the cache agent runs so normal operating hours are not interrupted.

5.4 Customizing Logs

The IBM WTE can be customized to provide access and error logs. The logs are specified by the following directives:

- `AccessLog`

  This log is used for logging local document requests.

- `ErrorLog`

  This log is used for logging any errors.

- `ProxyAccessLog`

  This log is used for logging proxy requests.

- `CacheAccessLog`

  This log is used for logging hits on the proxy cache, and is only valid if the WTE server is running as a proxy

Prerequisites

To enable logging of requests to the proxy cache, the following must be defined (see 4.3.6, "Enabling Basic Caching" on page 93):

- `Caching` must be set to `On` (default is `Off`).
- `CacheRoot` must be defined (by default, no `CacheRoot` is defined).
- `CacheAccessLog` must be defined (by default, no `CacheAccessLog` is defined).

It is recommended to perform log file space maintenance so as to ensure that the logs do not fill up the space allocated. This can be performed by using the Log Archiving forms.

When changing log file names and/or paths, ensure that the server has permission to write to the new log file.

The directory where the logs are installed by default change by platform, as shown in the following table

Table 5. Default Directories where Logs Are Installed

Platform	Log directory
AIX	/usr/lpp/internet/server_root/logs
Solaris	
Windows NT	C:\WWW\Logs

The default log file names are shown in the following table:

Table 6. Default Log File Names

Log	File name
AccessLog	httpd-log
ErrorLog	httpd-errors
ProxyAccessLog	httpd-proxy

As already noted in "Prerequisites" on page 116, `CacheAccessLog` does not have a default value. Its value is defined only when `Caching` is set to `On`. You have to specify the fully qualified path of the `CacheAccessLog` file name.

The server starts a new log file each day at midnight if it is running. Otherwise, the server starts a new log file the first time you start it on a given day. This feature cannot be changed.

When creating the file, the server uses the file name you specify and appends a date suffix. The date suffix is in the format *Mmmddyyyy*, where:

- *Mmm* is the first three letters of the month.
- *dd* is the day of the month.
- *yyyy* is the year.

An example of a log file name started on February 28, 1999 would be httpd-log.Feb281999.

5.4.1 WTE Logging

The WTE server logs activity in the proxy access, cache access, access and error logs. These logs are stored on the disk and at midnight every day, the server closes the current log files and creates new log files for the coming day. The path and name of these log files can be customized using the Log Files form. Click the **Server Configuration** folder, then click the **Logging** subfolder, and select **Log Files**, as shown below:

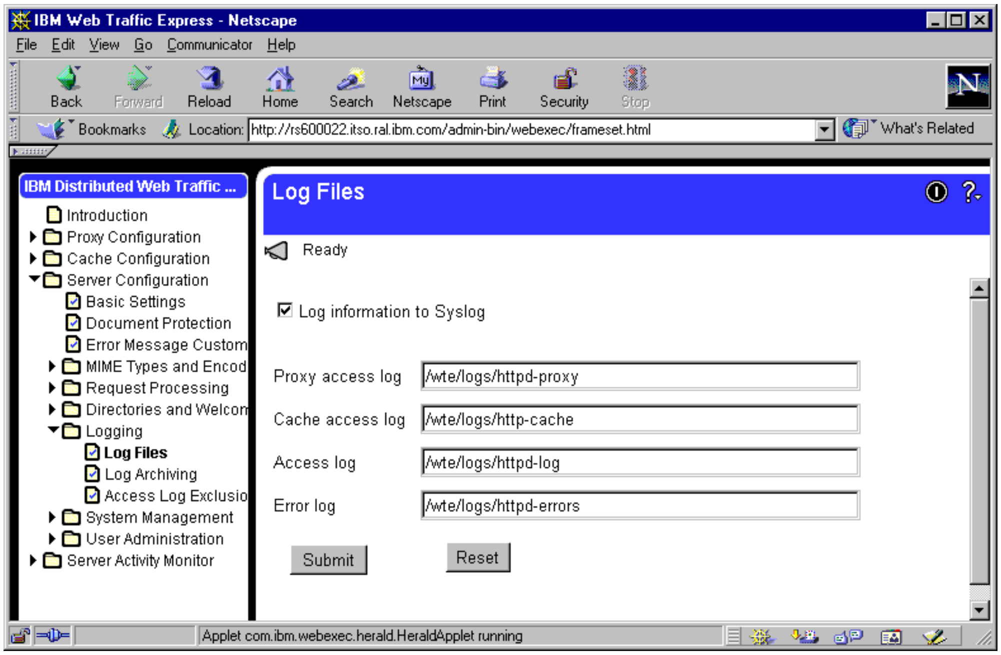

☑ Log information to Syslog

Proxy access log	/wte/logs/httpd-proxy
Cache access log	/wte/logs/http-cache
Access log	/wte/logs/httpd-log
Error log	/wte/logs/httpd-errors

Figure 69. IBM WTE GUI - Log Files

We click the **Log information to Syslog** box as we also want send the log information to the underlying operating system. On AIX and Solaris, the syslog file is /etc/syslog.conf. We configured it through the following steps:

1. Add the following lines to the /etc/syslog.conf file

```
user.err /wte/logs/wte.log # For error information
user.info /wte/logs/wte.log # For access information
```

2. Restart the syslogd daemon with the following command:

```
refresh -s syslogd
```

Notice that we changed the path of the log files to the file system /wte, which we created on our platform (see “The /usr and /wte File Systems” on page 53 and “The /opt/WSPP Directory and /wte File System” on page 60). The default path is /usr/lpp/internet/server_root. Press **Submit** to make the modification.

You have to restart the WTE server to activate the modification (see 4.3.3, “Directives not Changed on Restart” on page 89). Refer to 3.5, “Starting and Stopping WTE” on page 70.

5.4.2 Log Archiving

Log archiving is recommended as a means to maintain WTE server log files. The Log Archiving form supports two log archiving methods:

- Compress
- Purge

You can also opt for no log archiving.

On our AIX platform, we set the Log Archiving method to **Compress**. In particular, we configure our WTE server to compress logs older than 30 days and delete them only after 90 days. We chose to use the `/usr/bin/compress` command to compress our log files. This configuration can be issued in the Log Archiving form, accessible by clicking the **Server Configuration** folder, then on the **Logging** subfolder, and finally on the **Log Archiving** menu item, as shown below:

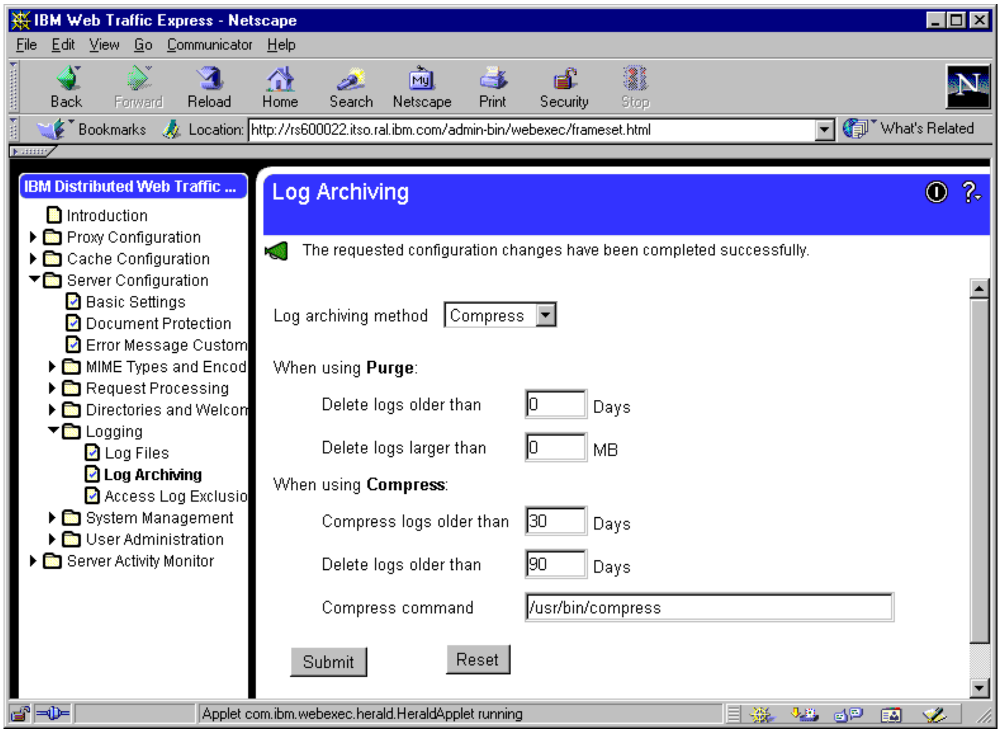

Log Archiving

The requested configuration changes have been completed successfully.

Log archiving method: Compress

When using **Purge**:

Delete logs older than 0 Days

Delete logs larger than 0 MB

When using **Compress**:

Compress logs older than 30 Days

Delete logs older than 90 Days

Compress command /usr/bin/compress

Figure 70. IBM WTE GUI - Log Archiving

5.4.3 Access Log Exclusions

Access log exclusions allow you to control what is logged. You can exclude data based on specific directories and/or files, user-agents, host names or IP addresses, methods, MIME types, and return codes. Click the **Server Configuration** folder, then click the **Logging** subfolder and select the **Access Log Exclusions** form. As this is a long form, we will show you the form in sections.

Let us suppose that we do not want to log the requests to the /Docs and /admin-bin directories, which are the documentation and forms directories respectively, as well as the /tunetips.html file. Then, we will put the entries as shown in the form section below:

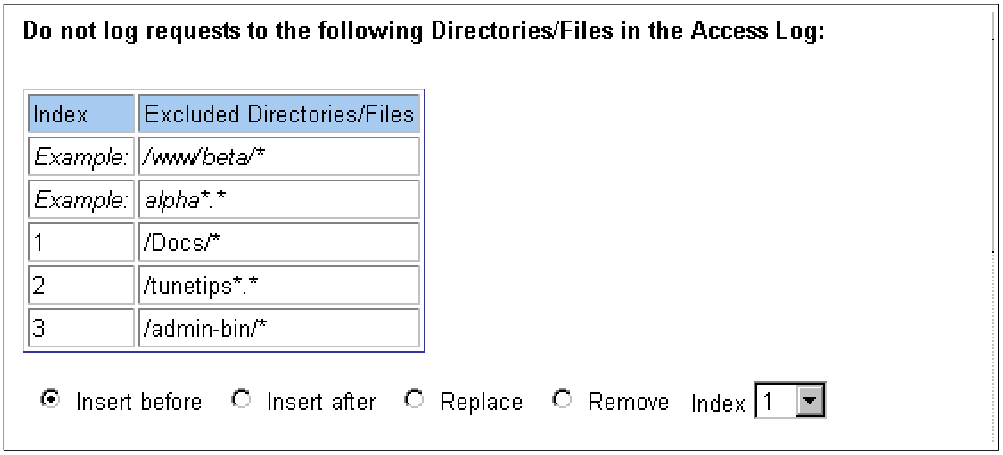

Do not log requests to the following Directories/Files in the Access Log:

Index	Excluded Directories/Files
Example:	*/www/beta/**
Example:	*alpha*.**
1	/Docs/*
2	/tunetips*.*
3	/admin-bin/*

● Insert before ○ Insert after ○ Replace ○ Remove Index 1

Figure 71. (Part 1 of 4). Access Log Exclusions

We do not put any entry in the user-agents section, as shown in the figure below:

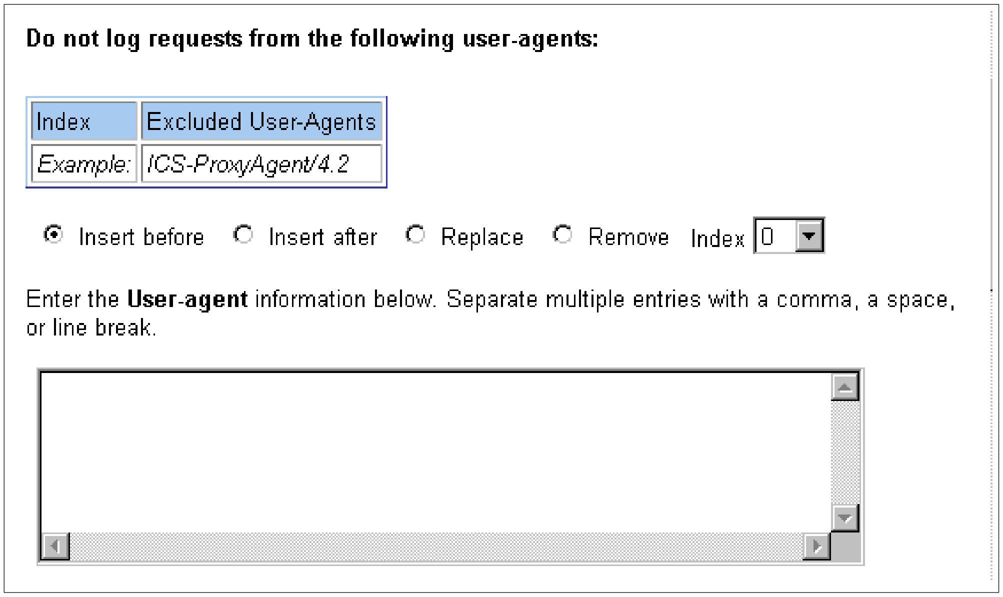

Do not log requests from the following user-agents:

Index	Excluded User-Agents
Example:	*ICS-ProxyAgent/4.2*

● Insert before ○ Insert after ○ Replace ○ Remove Index 0

Enter the **User-agent** information below. Separate multiple entries with a comma, a space, or line break.

Figure 72. (Part 2 of 4). Access Log Exclusions

The following figure shows that we have already configured our WTE server to not log requests coming from the machine with host name wtr05128.itso.ral.ibm.com. If you want to add other machine host names or IP addresses, you can use the text box provided. Notice that the wildcard character (*) is accepted.

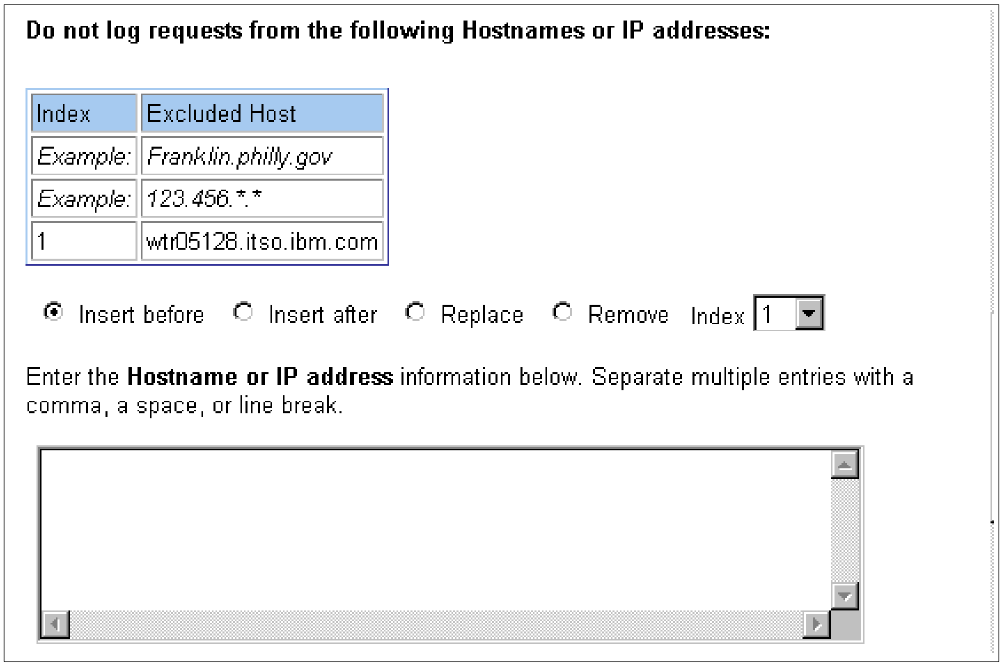

Do not log requests from the following Hostnames or IP addresses:

Index	Excluded Host
Example:	*Franklin.philly.gov*
Example:	*123.456.*.**
1	wtr05128.itso.ibm.com

● Insert before ○ Insert after ○ Replace ○ Remove Index 1

Enter the **Hostname or IP address** information below. Separate multiple entries with a comma, a space, or line break.

Figure 73. (Part 3 of 4). Access Log Exclusions

Since we do not want to log requests for files of the image/gif type, we only mark the **image/gif** box, as shown in the last section of the Access Log Exclusions form, shown below:

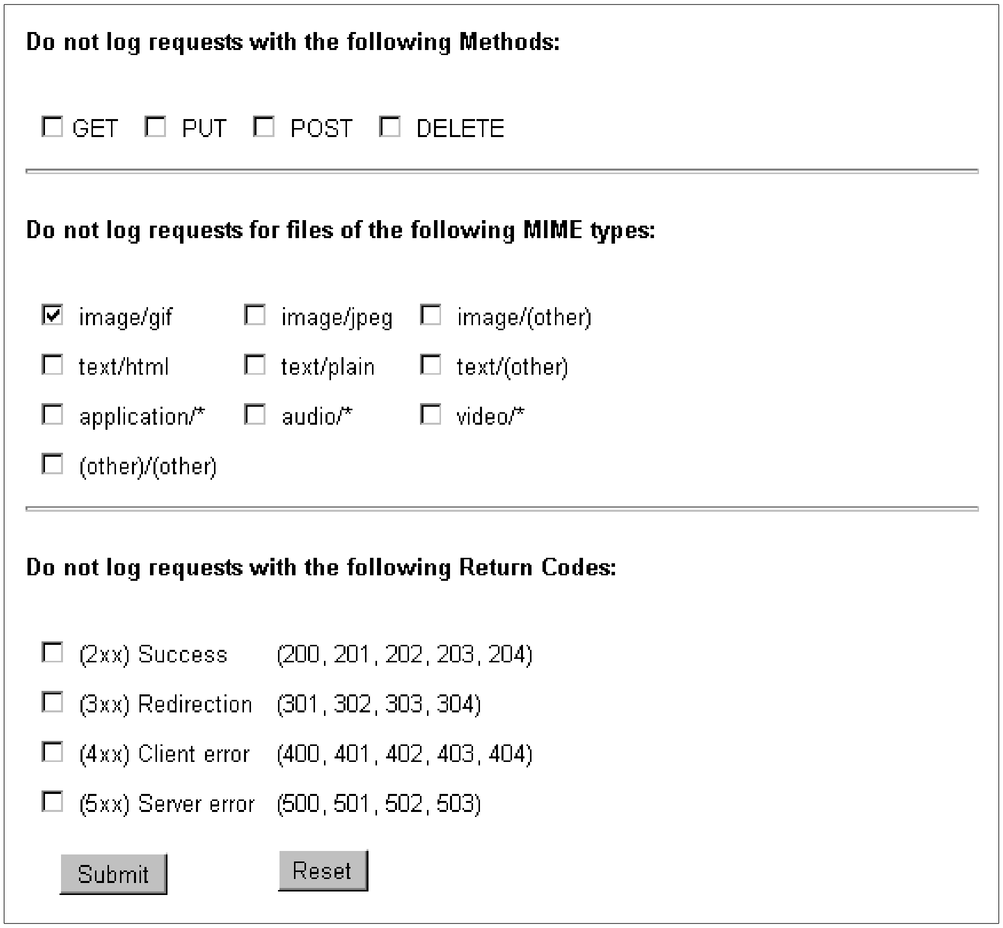

Do not log requests with the following Methods:

□ GET □ PUT □ POST □ DELETE

Do not log requests for files of the following MIME types:

☑ image/gif	□ image/jpeg	□ image/(other)
□ text/html	□ text/plain	□ text/(other)
□ application/*	□ audio/*	□ video/*
□ (other)/(other)		

Do not log requests with the following Return Codes:

□ (2xx) Success	(200, 201, 202, 203, 204)
□ (3xx) Redirection	(301, 302, 303, 304)
□ (4xx) Client error	(400, 401, 402, 403, 404)
□ (5xx) Server error	(500, 501, 502, 503)

Figure 74. (Part 4 of 4). Access Log Exclusions

Chapter 6. Handling Header Information with WTE

The proxy function of IBM Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2, allows several options to selectively hide the HTTP header information about the browser or the request that passes through it. Headers contain information about the `User-Agent` (consisting of browser and operating system data), `Client-IP` (consisting of the IP address of the requestor), and a `Referer` (providing the destination server with the URL of the referring link to this page). This header information can be blocked for privacy reasons, but there are advantages and disadvantages. While blocking header information increases client anonymity, it also limits the client from receiving customized materials.

When requesting documents, Web clients send headers that provide additional information about the browser or the request. Header generation occurs automatically when a request is sent. An example header is shown in the following screen:

```
User-Agent: Mozilla 4.5/WinNT
Client-IP: 9.24.106.76
Referer: http://www.ics.raleigh.ibm.com/WebTrafficExpress/main.html
```

Descriptions of the fields in this header are:

- `User-Agent`

  The value of this field provides browser and operating system information.

- `Client-IP`

  The value of this field provides the IP address of the client requesting the URL.

- `Referer`

  The value of this field provides the destination server with the URL of the link referring to this page.

The following is a log captured by the Windows NT Network Monitor, showing the header information of a typical HTTP request, revealing the header information described above:

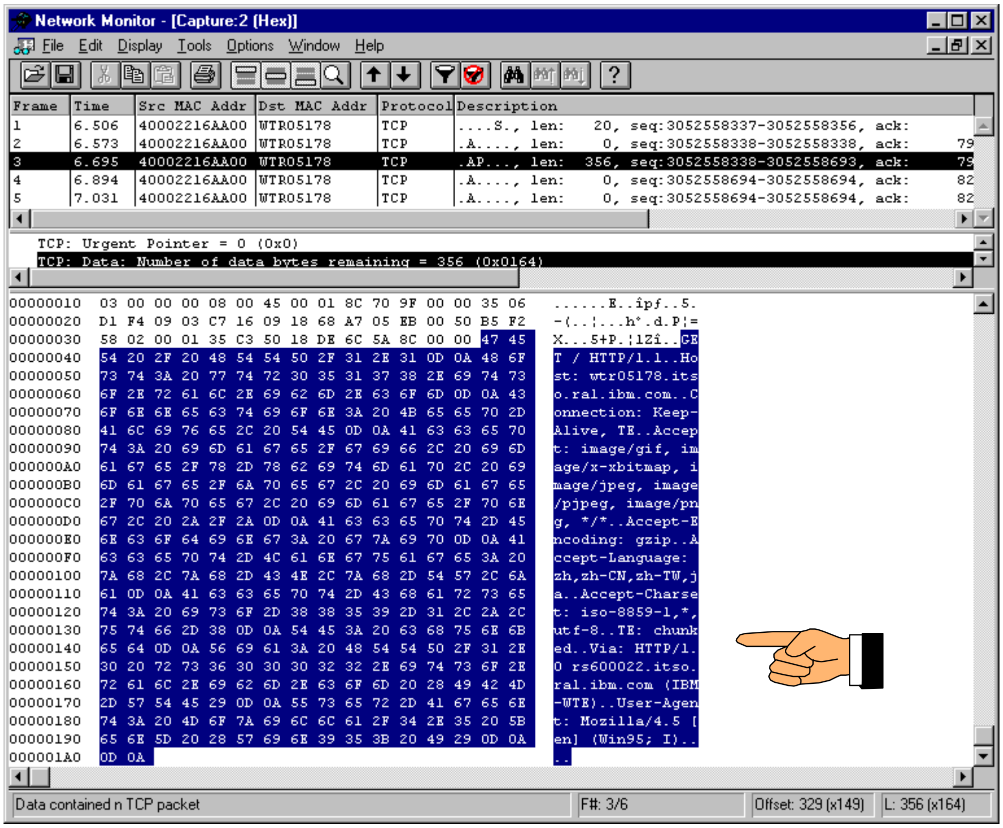
Network Monitor - [Capture:2 (Hex)]
File Edit Display Tools Options Window Help
?

```
Frame  Time   Src MAC Addr Dst MAC Addr   Protocol Description
1      6.506  40002216AA00 WTR05178       TCP      ....S., len:   20, seq:3052558337-3052558356, ack:
2      6.573  40002216AA00 WTR05178       TCP      .A...., len:    0, seq:3052558338-3052558338, ack:     79
3      6.695  40002216AA00 WTR05178       TCP      .AP..., len:  356, seq:3052558338-3052558693, ack:     79
4      6.894  40002216AA00 WTR05178       TCP      .A...., len:    0, seq:3052558694-3052558694, ack:     82
5      7.031  40002216AA00 WTR05178       TCP      .A...., len:    0, seq:3052558694-3052558694, ack:     82
```

```
  TCP: Urgent Pointer = 0 (0x0)
  TCP: Data: Number of data bytes remaining = 356 (0x0164)
```

```
00000010  03 00 00 00 08 00 45 00 01 8C 70 9F 00 00 35 06   ......E..îpf..5.
00000020  D1 F4 09 03 C7 16 09 18 68 A7 05 EB 00 50 B5 F2   -(..¦...h°.d.P¦=
00000030  58 02 00 01 35 C3 50 18 DE 6C 5A 8C 00 00 47 45   X...5+P.¦lZî..GE
00000040  54 20 2F 20 48 54 54 50 2F 31 2E 31 0D 0A 48 6F   T / HTTP/1.1..Ho
00000050  73 74 3A 20 77 74 72 30 35 31 37 38 2E 69 74 73   st: wtr05178.its
00000060  6F 2E 72 61 6C 2E 69 62 6D 2E 63 6F 6D 0D 0A 43   o.ral.ibm.com..C
00000070  6F 6E 6E 65 63 74 69 6F 6E 3A 20 4B 65 65 70 2D   onnection: Keep-
00000080  41 6C 69 76 65 2C 20 54 45 0D 0A 41 63 63 65 70   Alive, TE..Accep
00000090  74 3A 20 69 6D 61 67 65 2F 67 69 66 2C 20 69 6D   t: image/gif, im
000000A0  61 67 65 2F 78 2D 78 62 69 74 6D 61 70 2C 20 69   age/x-xbitmap, i
000000B0  6D 61 67 65 2F 6A 70 65 67 2C 20 69 6D 61 67 65   mage/jpeg, image
000000C0  2F 70 6A 70 65 67 2C 20 69 6D 61 67 65 2F 70 6E   /pjpeg, image/pn
000000D0  67 2C 20 2A 2F 2A 0D 0A 41 63 63 65 70 74 2D 45   g, */*..Accept-E
000000E0  6E 63 6F 64 69 6E 67 3A 20 67 7A 69 70 0D 0A 41   ncoding: gzip..A
000000F0  63 63 65 70 74 2D 4C 61 6E 67 75 61 67 65 3A 20   ccept-Language: 
00000100  7A 68 2C 7A 68 2D 43 4E 2C 7A 68 2D 54 57 2C 6A   zh,zh-CN,zh-TW,j
00000110  61 0D 0A 41 63 63 65 70 74 2D 43 68 61 72 73 65   a..Accept-Charse
00000120  74 3A 20 69 73 6F 2D 38 38 35 39 2D 31 2C 2A 2C   t: iso-8859-1,*,
00000130  75 74 66 2D 38 0D 0A 54 45 3A 20 63 68 75 6E 6B   utf-8..TE: chunk
00000140  65 64 0D 0A 56 69 61 3A 20 48 54 54 50 2F 31 2E   ed..Via: HTTP/1.
00000150  30 20 72 73 36 30 30 30 32 32 2E 69 74 73 6F 2E   0 rs600022.itso.
00000160  72 61 6C 2E 69 62 6D 2E 63 6F 6D 20 28 49 42 4D   ral.ibm.com (IBM
00000170  2D 57 54 45 29 0D 0A 55 73 65 72 2D 41 67 65 6E   -WTE)..User-Agen
00000180  74 3A 20 4D 6F 7A 69 6C 6C 61 2F 34 2E 35 20 5B   t: Mozilla/4.5 [
00000190  65 6E 5D 20 28 57 69 6E 39 35 3B 20 49 29 0D 0A   en] (Win95; I)..
000001A0  0D 0A                                             ..
```

```
Data contained n TCP packet        F#: 3/6        Offset: 329 (x149)        L: 356 (x164)
```

Figure 75. NT Network Monitor -Header Information 1

6.1 Configuring Header Options

Header options can be configured using the Configuration and Administration forms or by manually editing the WTE proxy configuration file, ibmproxy.conf. Three machines will be used in our environment as shown in the following diagram:

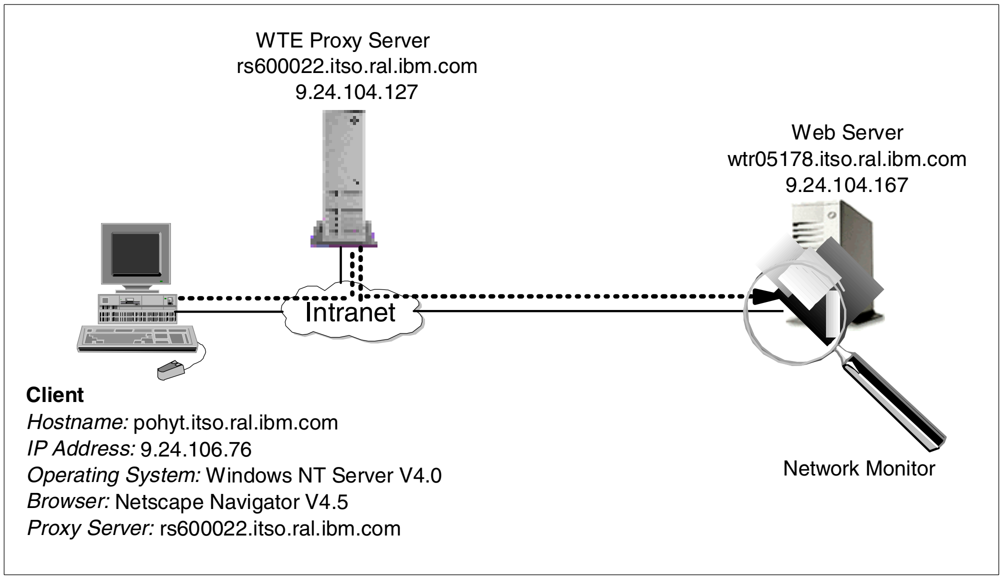
WTE Proxy Server
rs600022.itso.ral.ibm.com
9.24.104.127
Web Server
wtr05178.itso.ral.ibm.com
9.24.104.167
Intranet
Client
Hostname: pohyt.itso.ral.ibm.com
IP Address: 9.24.106.76
Operating System: Windows NT Server V4.0
Browser: Netscape Navigator V4.5
Proxy Server: rs600022.itso.ral.ibm.com
Network Monitor

Figure 76. Environment for Header Information Testing

We will describe header configuration using the WTE Configuration and Administration forms.

On the navigation menu:

1. Click the **Proxy Configuration** folder.
2. Click **Privacy Settings**.

As we want to hide the IP address of the client, its browser and its operating system header information, we fill in the Privacy Settings form as follows:

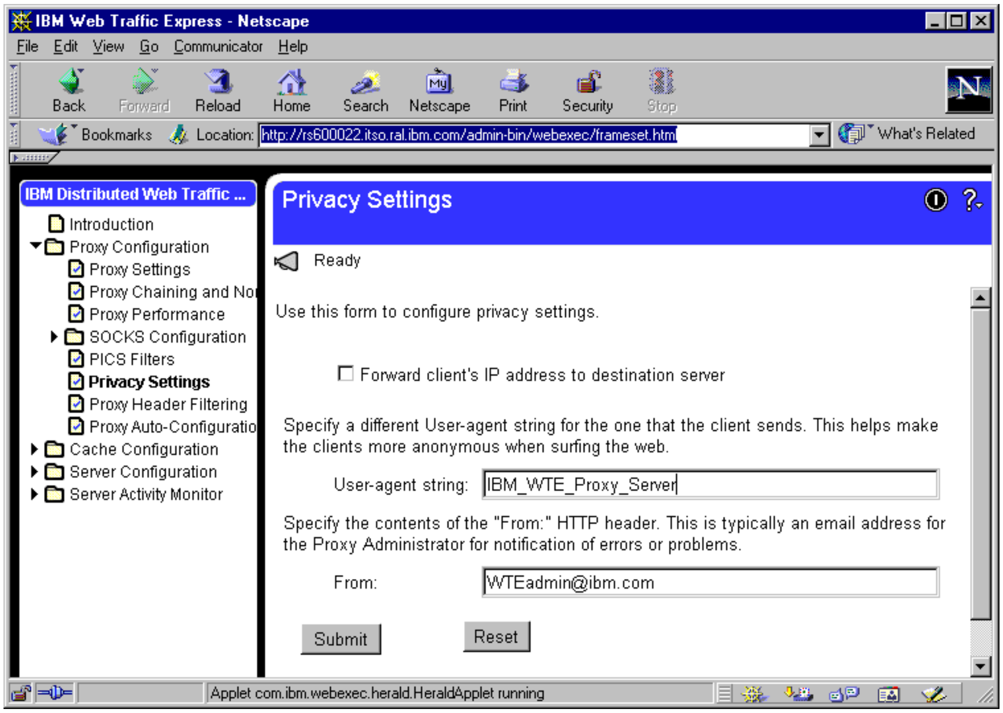
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Distributed Web Traffic ...
Introduction
Proxy Configuration
Proxy Settings
Proxy Chaining and No
Proxy Performance
SOCKS Configuration
PICS Filters
Privacy Settings
Proxy Header Filtering
Proxy Auto-Configuratio
Cache Configuration
Server Configuration
Server Activity Monitor
Privacy Settings
Ready
Applet com.ibm.webexec.herald.HeraldApplet running

Use this form to configure privacy settings.

☐ Forward client's IP address to destination server

Specify a different User-agent string for the one that the client sends. This helps make the clients more anonymous when surfing the web.

User-agent string: IBM_WTE_Proxy_Server

Specify the contents of the "From:" HTTP header. This is typically an email address for the Proxy Administrator for notification of errors or problems.

From: WTEadmin@ibm.com

Submit Reset

Figure 77. Proxy Configuration - Privacy Settings

On the Privacy Settings form:

1. For security reasons, we do not select the **Forward client's IP address to destination server** box. This directive specifies whether the WTE proxy server should forward the requesting client's IP address to the destination server or not. If the box is checked, a header field will be generated:

   `Client-IP:` *`IP_address_of_clientX`*

2. To improve client anonymity, we entered `IBM_WTE_Proxy_Server` in the User-agent string field. The value entered will overwrite the value sent by the client, which should have contained browser and operating system information.
3. For notification of any errors, we provided the e-mail address of the WTE administrator, `WTEadmin@ibm.com`, so that users can contact the WTE administrator if they encounter any errors or problems. The value entered should become the value of the `From:` header field.
4. Click the **Submit** button to make the changes to the WTE configuration file, ibmproxy.conf.
5. Click the **Restart** icon to refresh the WTE server.

To trace what header information is being sent, we will start the Windows NT Network Monitor on the Web server machine, as shown in the following figure:

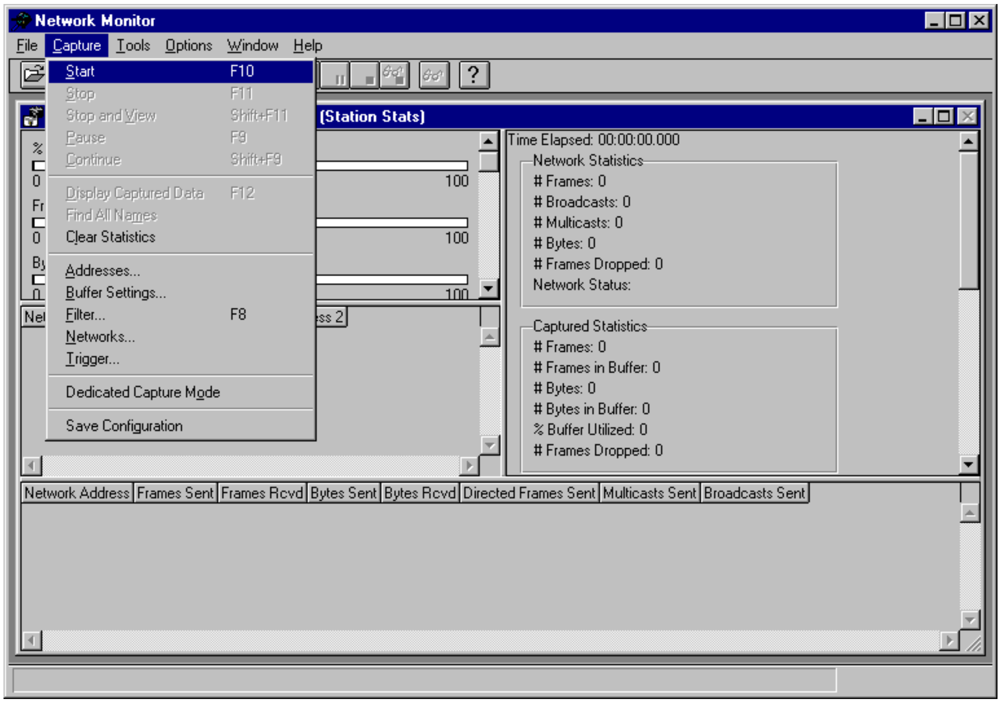

Time Elapsed: 00:00:00.000

Network Statistics
- # Frames: 0
- # Broadcasts: 0
- # Multicasts: 0
- # Bytes: 0
- # Frames Dropped: 0
- Network Status:

Captured Statistics
- # Frames: 0
- # Frames in Buffer: 0
- # Bytes: 0
- # Bytes in Buffer: 0
- % Buffer Utilized: 0
- # Frames Dropped: 0

Network Address	Frames Sent	Frames Rcvd	Bytes Sent	Bytes Rcvd	Directed Frames Sent	Multicasts Sent	Broadcasts Sent

Figure 78. Windows NT Network Monitor - Starting

From a client browser, we invoked the URL `http://wtr05178.ral.itso.ibm.com`. Once we got the requested page, we stopped the Windows NT Network Monitor as shown below:

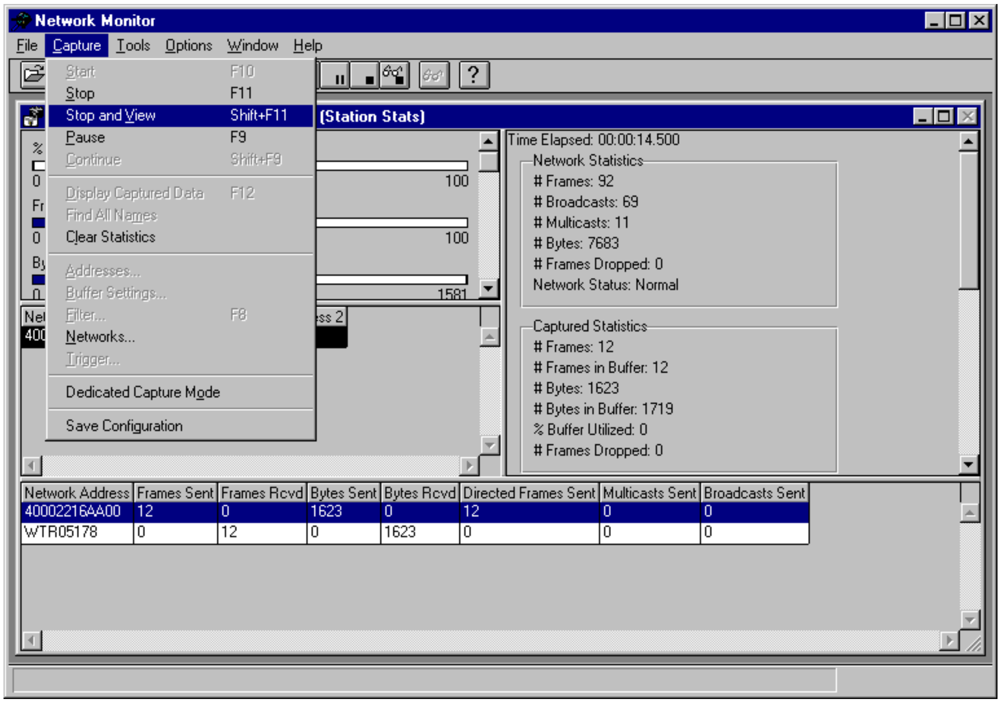

Time Elapsed: 00:00:14.500

Network Statistics
- # Frames: 92
- # Broadcasts: 69
- # Multicasts: 11
- # Bytes: 7683
- # Frames Dropped: 0
- Network Status: Normal

Captured Statistics
- # Frames: 12
- # Frames in Buffer: 12
- # Bytes: 1623
- # Bytes in Buffer: 1719
- % Buffer Utilized: 0
- # Frames Dropped: 0

Network Address	Frames Sent	Frames Rcvd	Bytes Sent	Bytes Rcvd	Directed Frames Sent	Multicasts Sent	Broadcasts Sent
40002216AA00	12	0	1623	0	12	0	0
WTR05178	0	12	0	1623	0	0	0

Figure 79. NT Network Monitor - Stopping

When the data captured is compared to Figure 75 on page 124, we can see that:

- The `From:` header field is `WTEadmin@ibm.com`.
- The `User-Agent:` field no longer shows the browser and operating system information of the client, but the string `IBM_WTE_Proxy_Server`.
- The IP address of the browser is successfully hidden, masked by the host name of the WTE proxy server itself.

This is shown in the following figure:

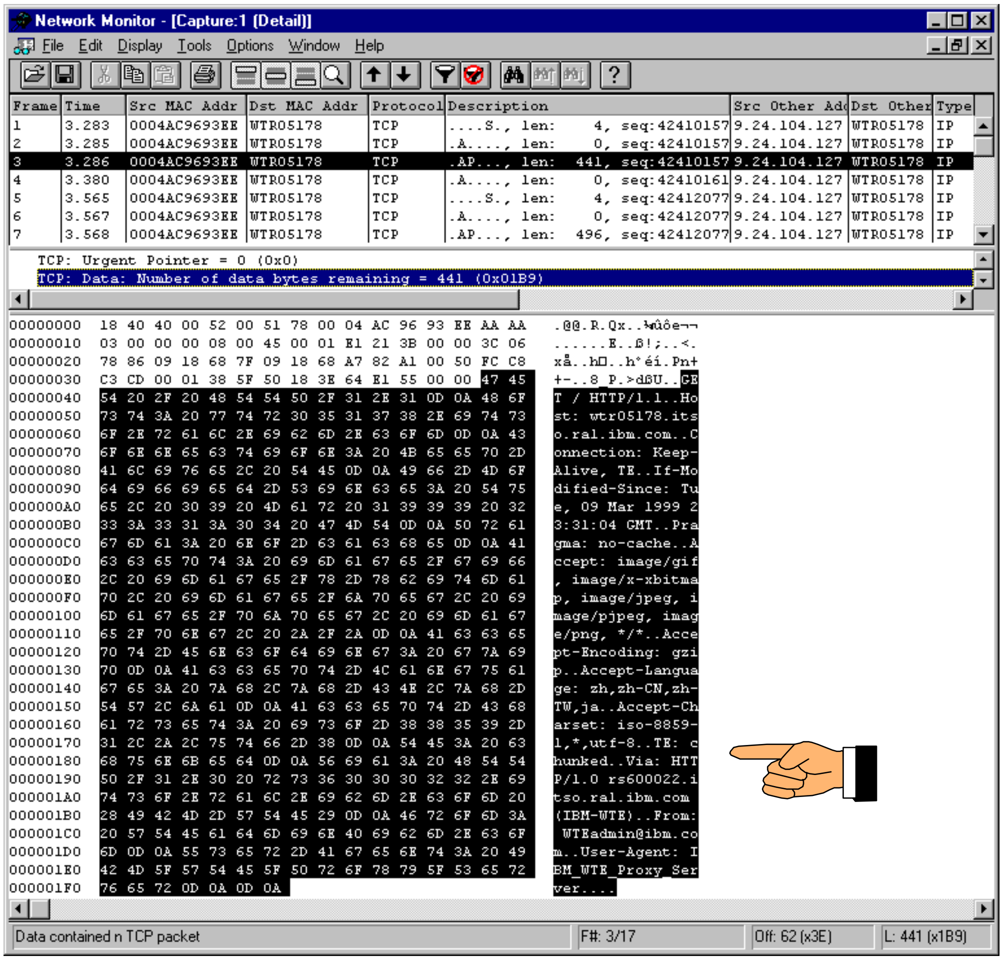
Network Monitor - [Capture:1 (Detail)]
File Edit Display Tools Options Window Help

Frame	Time	Src MAC Addr	Dst MAC Addr	Protocol	Description	Src Other Add	Dst Other	Type
1	3.283	0004AC9693EE	WTR05178	TCP	....S., len: 4, seq:42410157	9.24.104.127	WTR05178	IP
2	3.285	0004AC9693EE	WTR05178	TCP	.A...., len: 0, seq:42410157	9.24.104.127	WTR05178	IP
3	3.286	0004AC9693EE	WTR05178	TCP	.AP..., len: 441, seq:42410157	9.24.104.127	WTR05178	IP
4	3.380	0004AC9693EE	WTR05178	TCP	.A...., len: 0, seq:42410161	9.24.104.127	WTR05178	IP
5	3.565	0004AC9693EE	WTR05178	TCP	....S., len: 4, seq:42412077	9.24.104.127	WTR05178	IP
6	3.567	0004AC9693EE	WTR05178	TCP	.A...., len: 0, seq:42412077	9.24.104.127	WTR05178	IP
7	3.568	0004AC9693EE	WTR05178	TCP	.AP..., len: 496, seq:42412077	9.24.104.127	WTR05178	IP

```
    TCP: Urgent Pointer = 0 (0x0)
    TCP: Data: Number of data bytes remaining = 441 (0x01B9)
```

```
00000000  18 40 40 00 52 00 51 78 00 04 AC 96 93 EE AA AA   .@@.R.Qx..¼ûôe¬¬
00000010  03 00 00 00 08 00 45 00 01 E1 21 3B 00 00 3C 06   ......E..ß!;..<.
00000020  78 86 09 18 68 7F 09 18 68 A7 82 A1 00 50 FC C8   xå..h□..h°éí.Pn+
00000030  C3 CD 00 01 38 5F 50 18 3E 64 E1 55 00 00 47 45   +-..8_P.>dßU..GE
00000040  54 20 2F 20 48 54 54 50 2F 31 2E 31 0D 0A 48 6F   T / HTTP/1.1..Ho
00000050  73 74 3A 20 77 74 72 30 35 31 37 38 2E 69 74 73   st: wtr05178.its
00000060  6F 2E 72 61 6C 2E 69 62 6D 2E 63 6F 6D 0D 0A 43   o.ral.ibm.com..C
00000070  6F 6E 6E 65 63 74 69 6F 6E 3A 20 4B 65 65 70 2D   onnection: Keep-
00000080  41 6C 69 76 65 2C 20 54 45 0D 0A 49 66 2D 4D 6F   Alive, TE..If-Mo
00000090  64 69 66 69 65 64 2D 53 69 6E 63 65 3A 20 54 75   dified-Since: Tu
000000A0  65 2C 20 30 39 20 4D 61 72 20 31 39 39 39 20 32   e, 09 Mar 1999 2
000000B0  33 3A 33 31 3A 30 34 20 47 4D 54 0D 0A 50 72 61   3:31:04 GMT..Pra
000000C0  67 6D 61 3A 20 6E 6F 2D 63 61 63 68 65 0D 0A 41   gma: no-cache..A
000000D0  63 63 65 70 74 3A 20 69 6D 61 67 65 2F 67 69 66   ccept: image/gif
000000E0  2C 20 69 6D 61 67 65 2F 78 2D 78 62 69 74 6D 61   , image/x-xbitma
000000F0  70 2C 20 69 6D 61 67 65 2F 6A 70 65 67 2C 20 69   p, image/jpeg, i
00000100  6D 61 67 65 2F 70 6A 70 65 67 2C 20 69 6D 61 67   mage/pjpeg, imag
00000110  65 2F 70 6E 67 2C 20 2A 2F 2A 0D 0A 41 63 63 65   e/png, */*..Acce
00000120  70 74 2D 45 6E 63 6F 64 69 6E 67 3A 20 67 7A 69   pt-Encoding: gzi
00000130  70 0D 0A 41 63 63 65 70 74 2D 4C 61 6E 67 75 61   p..Accept-Langua
00000140  67 65 3A 20 7A 68 2C 7A 68 2D 43 4E 2C 7A 68 2D   ge: zh,zh-CN,zh-
00000150  54 57 2C 6A 61 0D 0A 41 63 63 65 70 74 2D 43 68   TW,ja..Accept-Ch
00000160  61 72 73 65 74 3A 20 69 73 6F 2D 38 38 35 39 2D   arset: iso-8859-
00000170  31 2C 2A 2C 75 74 66 2D 38 0D 0A 54 45 3A 20 63   1,*,utf-8..TE: c
00000180  68 75 6E 6B 65 64 0D 0A 56 69 61 3A 20 48 54 54   hunked..Via: HTT
00000190  50 2F 31 2E 30 20 72 73 36 30 30 30 32 32 2E 69   P/1.0 rs600022.i
000001A0  74 73 6F 2E 72 61 6C 2E 69 62 6D 2E 63 6F 6D 20   tso.ral.ibm.com 
000001B0  28 49 42 4D 2D 57 54 45 29 0D 0A 46 72 6F 6D 3A   (IBM-WTE)..From:
000001C0  20 57 54 45 61 64 6D 69 6E 40 69 62 6D 2E 63 6F    WTEadmin@ibm.co
000001D0  6D 0D 0A 55 73 65 72 2D 41 67 65 6E 74 3A 20 49   m..User-Agent: I
000001E0  42 4D 5F 57 54 45 5F 50 72 6F 78 79 5F 53 65 72   BM_WTE_Proxy_Ser
000001F0  76 65 72 0D 0A 0D 0A                              ver....
```

Data contained n TCP packet | F#: 3/17 | Off: 62 (x3E) | L: 441 (x1B9)

Figure 80. NT Network Monitor -Header Information 2

You can also selectively block client header information using the Proxy Header Filtering form.

On the navigation menu:

1. Click **Proxy Configuration** folder.
2. Click **Proxy Header Filtering**.

We will demonstrate an example of blocking the User-Agent information, by filling in the Proxy Header Filtering form as follows:

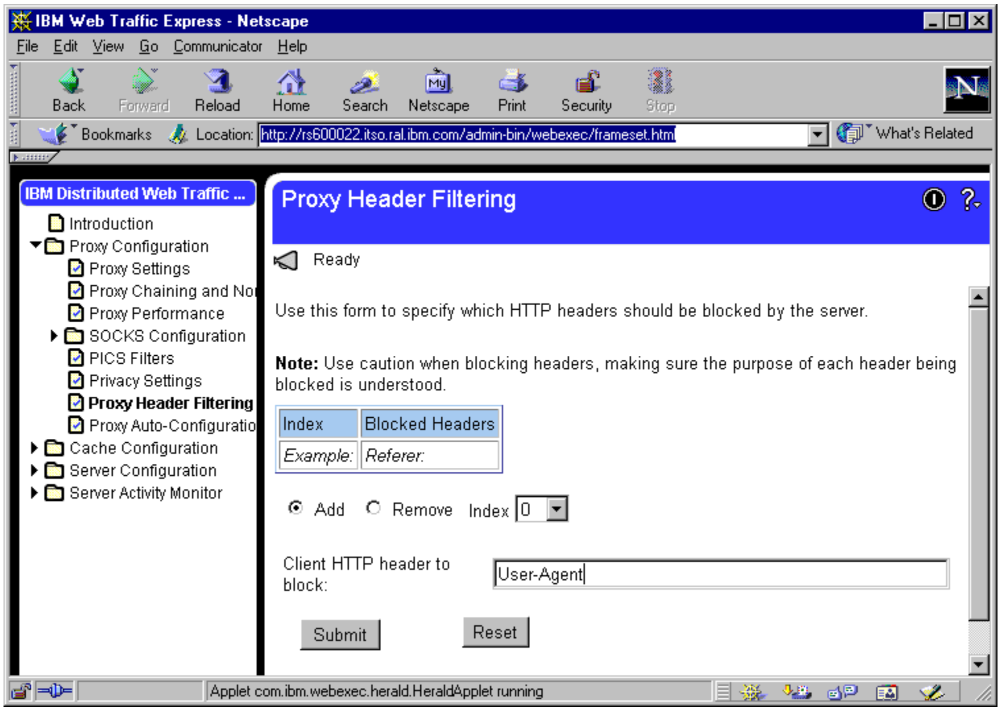

Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html

Proxy Header Filtering

Ready

Use this form to specify which HTTP headers should be blocked by the server.

Note: Use caution when blocking headers, making sure the purpose of each header being blocked is understood.

Index	Blocked Headers
Example:	*Referer:*

● Add ○ Remove Index 0

Client HTTP header to block: User-Agent

Applet com.ibm.webexec.herald.HeraldApplet running

Figure 81. Proxy Configuration - Proxy Header Filtering

After filling in the form, it is necessary to perform the following operations:

1. Click the **Submit** button to make the changes to the WTE ibmproxy.conf configuration file.
2. Click the **Restart** icon to refresh the httpd daemon.

As above, we will set up the Windows NT Network Monitor and request the same URL as before. The following is the data captured from the Windows NT Network Monitor. This time, the User-Agent field is totally blocked; in other words, it is not displayed at all:

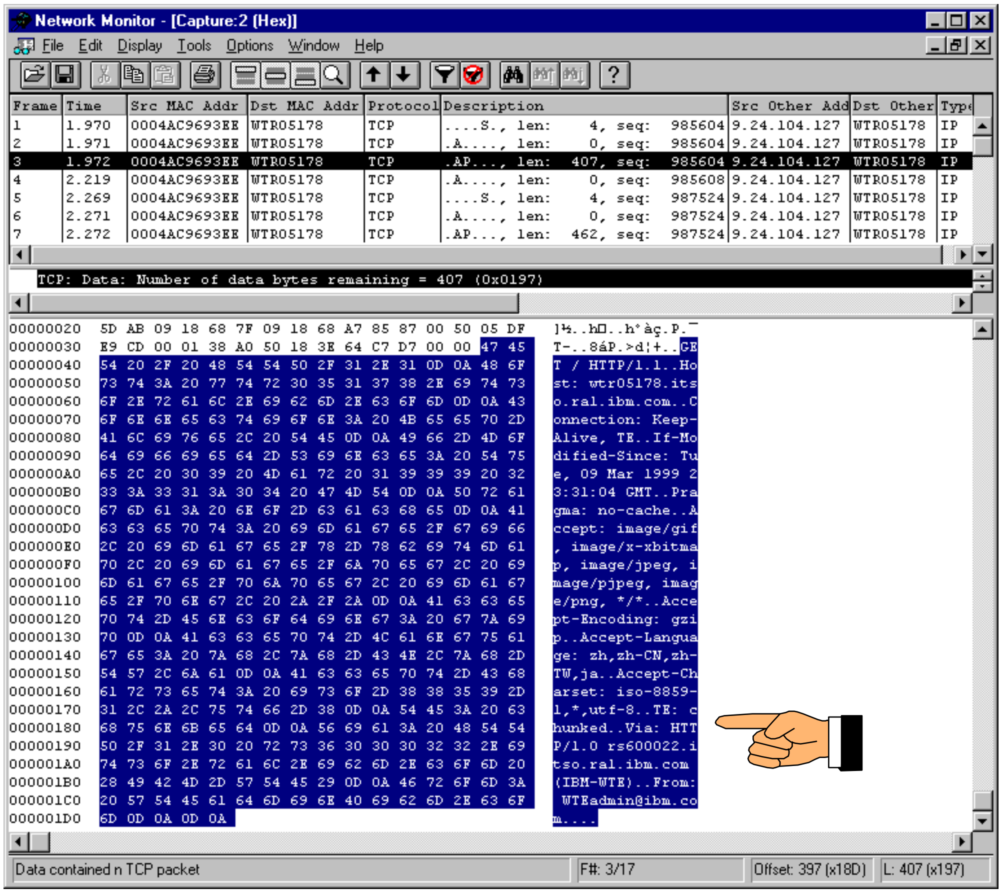

Network Monitor - [Capture:2 (Hex)]
File Edit Display Tools Options Window Help

```
Frame Time  Src MAC Addr Dst MAC Addr Protocol Description                       Src Other Add Dst Other Type
1     1.970 0004AC9693EE WTR05178     TCP      ....S., len:    4, seq:   985604 9.24.104.127  WTR05178  IP
2     1.971 0004AC9693EE WTR05178     TCP      .A...., len:    0, seq:   985604 9.24.104.127  WTR05178  IP
3     1.972 0004AC9693EE WTR05178     TCP      .AP..., len:  407, seq:   985604 9.24.104.127  WTR05178  IP
4     2.219 0004AC9693EE WTR05178     TCP      .A...., len:    0, seq:   985608 9.24.104.127  WTR05178  IP
5     2.269 0004AC9693EE WTR05178     TCP      ....S., len:    4, seq:   987524 9.24.104.127  WTR05178  IP
6     2.271 0004AC9693EE WTR05178     TCP      .A...., len:    0, seq:   987524 9.24.104.127  WTR05178  IP
7     2.272 0004AC9693EE WTR05178     TCP      .AP..., len:  462, seq:   987524 9.24.104.127  WTR05178  IP

TCP: Data: Number of data bytes remaining = 407 (0x0197)

00000020  5D AB 09 18 68 7F 09 18 68 A7 85 87 00 50 05 DF   ]¼..h□..h°àç.P.¯
00000030  E9 CD 00 01 38 A0 50 18 3E 64 C7 D7 00 00 47 45   T-..8áP.>d¦+..GE
00000040  54 20 2F 20 48 54 54 50 2F 31 2E 31 0D 0A 48 6F   T / HTTP/1.1..Ho
00000050  73 74 3A 20 77 74 72 30 35 31 37 38 2E 69 74 73   st: wtr05178.its
00000060  6F 2E 72 61 6C 2E 69 62 6D 2E 63 6F 6D 0D 0A 43   o.ral.ibm.com..C
00000070  6F 6E 6E 65 63 74 69 6F 6E 3A 20 4B 65 65 70 2D   onnection: Keep-
00000080  41 6C 69 76 65 2C 20 54 45 0D 0A 49 66 2D 4D 6F   Alive, TE..If-Mo
00000090  64 69 66 69 65 64 2D 53 69 6E 63 65 3A 20 54 75   dified-Since: Tu
000000A0  65 2C 20 30 39 20 4D 61 72 20 31 39 39 39 20 32   e, 09 Mar 1999 2
000000B0  33 3A 33 31 3A 30 34 20 47 4D 54 0D 0A 50 72 61   3:31:04 GMT..Pra
000000C0  67 6D 61 3A 20 6E 6F 2D 63 61 63 68 65 0D 0A 41   gma: no-cache..A
000000D0  63 63 65 70 74 3A 20 69 6D 61 67 65 2F 67 69 66   ccept: image/gif
000000E0  2C 20 69 6D 61 67 65 2F 78 2D 78 62 69 74 6D 61   , image/x-xbitma
000000F0  70 2C 20 69 6D 61 67 65 2F 6A 70 65 67 2C 20 69   p, image/jpeg, i
00000100  6D 61 67 65 2F 70 6A 70 65 67 2C 20 69 6D 61 67   mage/pjpeg, imag
00000110  65 2F 70 6E 67 2C 20 2A 2F 2A 0D 0A 41 63 63 65   e/png, */*..Acce
00000120  70 74 2D 45 6E 63 6F 64 69 6E 67 3A 20 67 7A 69   pt-Encoding: gzi
00000130  70 0D 0A 41 63 63 65 70 74 2D 4C 61 6E 67 75 61   p..Accept-Langua
00000140  67 65 3A 20 7A 68 2C 7A 68 2D 43 4E 2C 7A 68 2D   ge: zh,zh-CN,zh-
00000150  54 57 2C 6A 61 0D 0A 41 63 63 65 70 74 2D 43 68   TW,ja..Accept-Ch
00000160  61 72 73 65 74 3A 20 69 73 6F 2D 38 38 35 39 2D   arset: iso-8859-
00000170  31 2C 2A 2C 75 74 66 2D 38 0D 0A 54 45 3A 20 63   1,*,utf-8..TE: c
00000180  68 75 6E 6B 65 64 0D 0A 56 69 61 3A 20 48 54 54   hunked..Via: HTT
00000190  50 2F 31 2E 30 20 72 73 36 30 30 30 32 32 2E 69   P/1.0 rs600022.i
000001A0  74 73 6F 2E 72 61 6C 2E 69 62 6D 2E 63 6F 6D 20   tso.ral.ibm.com
000001B0  28 49 42 4D 2D 57 54 45 29 0D 0A 46 72 6F 6D 3A   (IBM-WTE)..From:
000001C0  20 57 54 45 61 64 6D 69 6E 40 69 62 6D 2E 63 6F    WTEadmin@ibm.co
000001D0  6D 0D 0A 0D 0A                                    m....

Data contained n TCP packet    F#: 3/17    Offset: 397 (x18D)    L: 407 (x197)
```

Figure 82. NT Network Monitor -Header Information 3

This completes our demonstration of the capabilities offered by WTE V2 of handling header information.

Chapter 7. Using WTE as a Transparent Proxy Server

Transparent proxying means that the client software is totally unaware of the existence of the intermediate proxy server. Normally, if a client browser uses a proxy server, then the browser must be configured to specify the address and port of the proxy server. With transparent proxy, clients do not have to configure their browsers.

The benefits of this proxy server technology are the following:

- Transparent proxying minimizes user errors, since clients do not configure their browsers.
- Transparent proxy can be deployed where appropriate without reconfiguring the clients.
- Clients are *forced* to use the proxy server, which automatically benefits from PICS filtering implementation and caching.

In this chapter, we study the transparent proxy support offered by IBM Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2. Transparent proxy is currently available only on the AIX platform.

7.1 How Transparent Proxy Works

In a transparent proxy scenario, a router or switch in the network is set up to redirect port 80 traffic to the WTE proxy server, instead of to its next network hop. The WTE server will accept all of this traffic (regardless of its actual destination) and process each request, by reconstructing the full URL being requested. WTE obtains the requested URL from the `HOST:` header sent by the client; after this, the request can be processed normally. If a `HOST:` header does not exist for the request, then WTE will use the destination address of the socket to fulfill the request. The WTE proxy server can also recognize local requests (for status pages, for example) and does not proxy those.

Other requests (for example, SSL and FTP requests) are not tunnelled through a transparent proxy. The router that redirects traffic to the proxy server only redirects traffic going to port 80. It does not redirect the traffic going to other ports, and requests directed to ports other than 80 are sent to the destination server without passing through the WTE transparent proxy. Therefore, SSL requests, which by default are directed to port 443, are sent directly to the secure server the user is communicating with. Therefore, the transparent proxy never sees the traffic directed to a different port. On the other hand, proxying SSL requests with a transparent proxy would not be useful at all as SSL pages cannot even be cached. Furthermore, SSL tunneling also adds latency to the proxy server.

SecureWay Network Dispatcher in a Transparent Proxy Scenario

As you can see, in order to implement a transparent proxy environment with WTE, it is necessary that a router be properly configured to route specific requests to the WTE transparent proxy server, and all other requests directly to the Internet. This role could be played also by a machine running IBM SecureWay Network Dispatcher (ND), the Load Balancing component of IBM WebSphere Performance Pack Version 2. We will discuss this in “If the Router Role Is Played by ND” on page 135.

Another way to use ND in a transparent proxy scenario is to implement WTE proxy server load balancing. This possibility is further discussed in “Load Balancing Multiple WTE Transparent Proxy Servers” on page 133.

Note that in WebSphere Performance Pack Version 2, transparent proxy is supported on the AIX platform only and works for HTTP requests only.

Authentication using user ID and password will work with transparent proxy. However, proxy authentication (see Chapter 11, “WTE Proxy Server Protection Scenario” on page 175) and authentication using the client’s IP address will not. In fact, in a transparent proxy scenario, browsers are not aware of the transparent proxy, and are not able to handle user ID and password requests from the proxy server.

7.2 WTE Transparent Proxy Scenario

In this section, we describe our experience with WTE transparent proxy.

The following figure offers a graphical representation of our working environment and shows how transparent proxying works with the HTTP protocol (port 80) and with traffic directed to other ports, such as SSL (port 443):

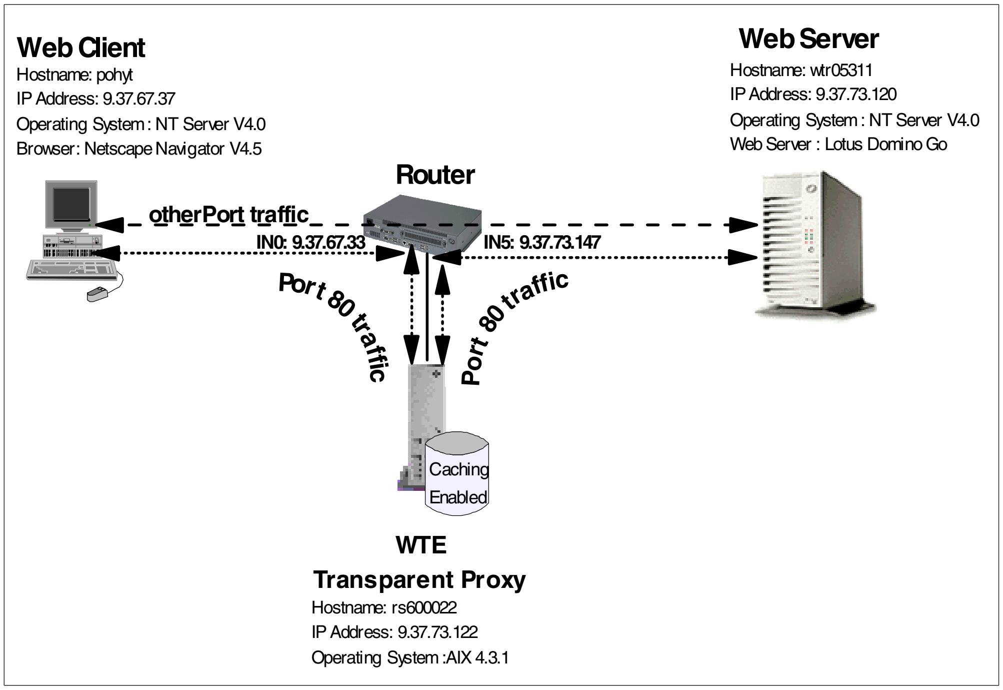

Figure 83. Transparent Proxy Server Scenario Representation

Table 7 on page 133 is a summary of the components we used in this scenario:

Table 7. Transparent Proxy Server Scenario Components

Role	Operating System	Host Name	IP Address	Default Gateway
Web Client	Windows NT Server V4.0	pohyt	9.37.67.37	9.37.67.33
WTE Transparent Proxy Server	AIX Version 4.3.1	rs600022	9.37.73.122	9.37.73.147
Web Server	Windows NT Server V4.0	wtr05311	9.37.73.120	9.37.73.147

When a Web browser is not configured to use a proxy server, it sends partial URLs. For example, if we want to get `http://www.ibm.com/`, the browser would connect to `www.ibm.com` and request the Web server home directory /. With a transparent proxy server in the network, the browser still tries to connect to www.ibm.com. But the proxy server intercepts the request, generates the full URL for this request, and satisfies it for the client.

Load Balancing Multiple WTE Transparent Proxy Servers

The scenario represented in Figure 83 on page 133 is very interesting because it demonstrates how WTE transparent proxying works. However, with WebSphere Performance Pack Version 2, it is possible to implement more complex and more reliable transparent proxy environments, where multiple WTE transparent proxy servers are load balanced by an ND machine and share the same cache using IBM AFS Enterprise File System (AFS).

7.2.1 How Client Machines Must Be Set Up for Transparent Proxy

On the client machine, whose host name is pohyt and IP address 9.37.67.37, we set the Web browser to use direct connection to the Internet. As we are using

Netscape Navigator V4.5, the direct proxy configuration is set in the Preferences menu as shown:

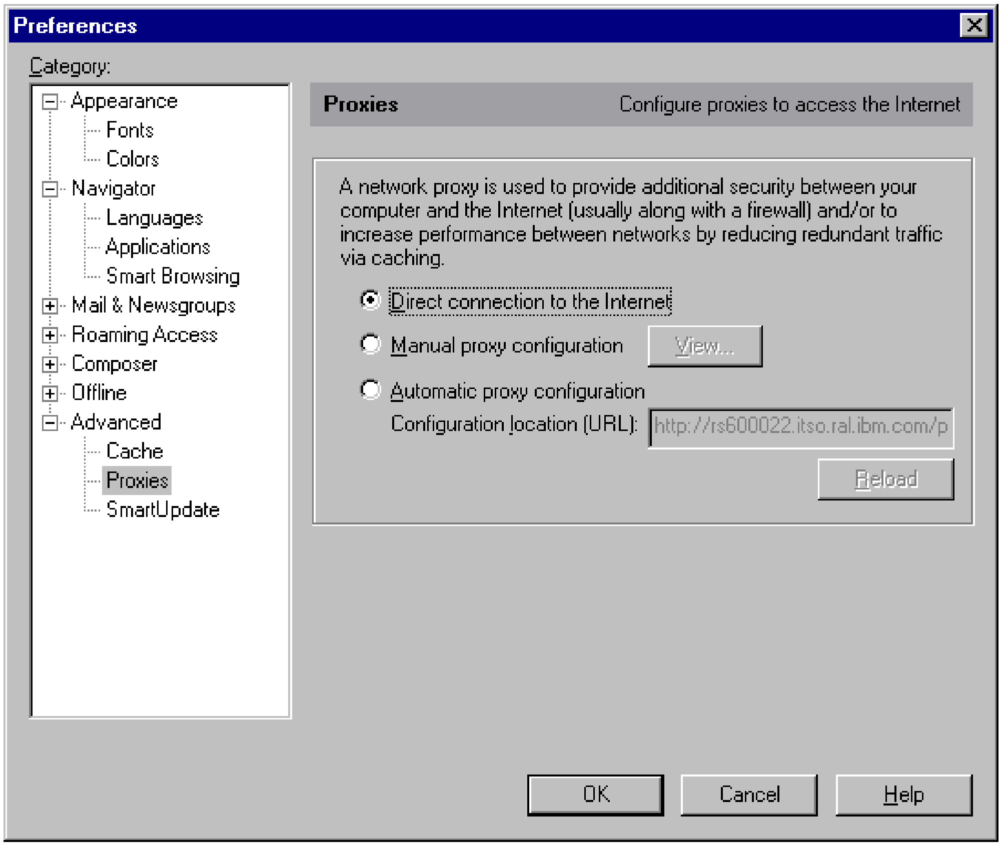

A network proxy is used to provide additional security between your computer and the Internet (usually along with a firewall) and/or to increase performance between networks by reducing redundant traffic via caching.

● Direct connection to the Internet

○ Manual proxy configuration

○ Automatic proxy configuration

Configuration location (URL): http://rs600022.itso.ral.ibm.com/p

Figure 84. Direct Connection to the Internet Configuration

If you are using Netscape Navigator, as shown above, on the Preferences menu click the **Direct connection to the Internet** button. Then click **OK** to activate the changes.

The reason for this configuration is that the client does not need to be aware that a transparent proxy server will proxy all the port 80 HTTP requests. This will be transparent to the user. For this reason, the browser on the client machine can be configured for a direct connection to the Internet. The only configuration requirement is that on the client machine, the default gateway be set to a router that redirects HTTP traffic (port 80) to the transparent proxy, and all the traffic for other ports directly to the Internet (or to another router). The router configuration is described in the next section.

7.2.2 Setting Up the IBM 2210 Router for Transparent Proxy

In this section, we describe a brief procedure to configure the IBM 2210 Router.

First of all, it is necessary to establish a connection to the IBM 2210 Router via an ASCII terminal or from the serial port of a personal computer with some terminal emulation software, such as the Windows NT HyperTerminal. Then, it is necessary to perform the router basic configuration. Finally, you will configure the router for the transparent proxy environment.

If the Router Role Is Played by ND

As we said in “SecureWay Network Dispatcher in a Transparent Proxy Scenario” on page 132, an ND machine with (at least) two interfaces can play the role of the router in a transparent proxy environment.

To implement this solution, you will need to use the wildcard port/cluster feature of ND V2.1.

For an ND machine to do policy-based routing, you should define a wildcard cluster address (0.0.0.0), add to this port 80 and add to port 80 the WTE transparent proxy server. Then you need to add another wildcard cluster (0.0.0.0), add to this the wildcard port (0), and add to this another router as a server.

This way, traffic to any IP address on port 80 will be forwarded to the WTE transparent proxy server, and traffic on any IP address on any other port will be forwarded to another router. Notice that the clients have to be configured to use the ND box as their default gateway.

The solution we have just described represents a simple and not very expensive way to implement transparent proxy using only WebSphere Performance Pack Version 2.

7.2.2.1 Basic Configuration

The following session screen shows the sequence of configuration steps that we have performed for the router initial configuration. All the user inputs are highlighted in boldfaced text.

```
Data_Center1 *
Data_Center1 *t 6
Gateway user configuration
Data_Center1 Config>qconfig

Router Quick Configuration for the following:
o   Interfaces
o   Multilink PPP (w/o DIALs)
o   Dial Circuits (w/o DIALs)
o   Dial-in Access to LANs (DIALs)
o   Bridging
        Spanning Tree Bridge (STB)
        Source Routing Bridge (SRB)
        Source Routing/Transparent Bridge (SR/TB)
        Source Routing Transparent Bridge (SRT)
o   Protocols
        IP (including OSPF, RIP and SNMP)
        IPX
o   Booting
o   Service Port

Event Logging will be enabled for all configured subsystems
with logging level 'Standard'

Note:  Please be warned that any existing configuration for a particular item
will be removed if that item is configured through Quick Configuration

**********************************************************
Interface Configuration
**********************************************************

Type 'Yes' to Configure Interfaces
Type 'No' to skip Interface Configuration
Type 'Quit' to exit Quick Config

Configure Interfaces? (Yes, No, Quit): [Yes]
Type 'r' any time at this level to restart Interface Configuration

Intf 0 is Token Ring
Speed in Mb/sec (4, 16): [16]
Connector (STP, UTP): [UTP]

Intf 1 is WAN Frame Relay
Encapsulation for WAN interface 1  (PPP, Frame Relay, V34): [Frame Relay]
Cable type (RS-232 DTE, RS-232 DCE, V.35 DTE, V.35 DCE, V.36 DTE, X.21 DTE,
    X.21 DCE): [RS-232 DTE]

Intf 2 is WAN Frame Relay
Encapsulation for WAN interface 2  (PPP, Frame Relay, V34): [Frame Relay]
Cable type (RS-232 DTE, RS-232 DCE, V.35 DTE, V.35 DCE, V.36 DTE, X.21 DTE,
    X.21 DCE): [RS-232 DTE]

Intf 3 is WAN Frame Relay
Encapsulation for WAN interface 3  (PPP, Frame Relay, V34): [Frame Relay]
Cable type (RS-232 DTE, RS-232 DCE, V.35 DTE, V.35 DCE, V.36 DTE, X.21 DTE,
    X.21 DCE): [V.35 DTE]

Intf 4 is WAN Frame Relay
Encapsulation for WAN interface 4  (PPP, Frame Relay, V34): [Frame Relay]
Cable type (RS-232 DTE, RS-232 DCE, V.35 DTE, V.35 DCE, V.36 DTE, X.21 DTE,
    X.21 DCE): [RS-232 DCE]
Internal clock speed (decimal) (2400 - 64000): [64000]
```

```
Intf 5 is Token Ring
Speed in Mb/sec (4, 16): [16]
Connector (STP, UTP): [UTP]

Intf 6 is NULL Device

This is all configured device information:

Intf 0 is Token Ring, Speed 16 Mb/sec, Connector UTP
Intf 1 is WAN Frame Relay, RS-232 DTE cable
Intf 2 is WAN Frame Relay, RS-232 DTE cable
Intf 3 is WAN Frame Relay, V.35 DTE cable
Intf 4 is WAN Frame Relay, RS-232 DCE cable,
 internal clock speed 64000 bits/second
Intf 5 is Token Ring, Speed 16 Mb/sec, Connector UTP
Intf 6 is NULL Device

Save this configuration? (Yes, No): [Yes]

Device  configuration saved

**********************************************************
Bridging Configuration
**********************************************************

Type 'Yes' to Configure Bridging
Type 'No' to skip Bridging Configuration
Type 'Quit' to exit Quick Config

Configure Bridging? (Yes, No, Quit): [Yes] no

**********************************************************
Protocol Configuration
**********************************************************

Type 'Yes' to Configure Protocols
Type 'No' to skip Protocol Configuration
Type 'Quit' to exit Quick Config

Configure Protocols? (Yes, No, Quit): [Yes]
Type 'r' any time at this level to restart Protocol Configuration

Configure IP? (Yes, No): [Yes]
Type 'r' any time at this level to restart IP Configuration
IP Configuration is already present
Configure IP anyway? (Yes, No): [No] yes

Configuring Per-Interface IP Information

Configuring Interface 0 (Token Ring)
Configure IP on this interface? (Yes, No): [Yes]
IP Address: [9.37.67.33] 9.37.67.33
Address Mask: [255.255.255.0] 255.255.255.0
```

```
Configuring Interface 1 (WAN Frame Relay)
Configure IP on this interface? (Yes, No): [Yes] no

Configuring Interface 2 (WAN Frame Relay)
Configure IP on this interface? (Yes, No): [Yes] no

Configuring Interface 3 (WAN Frame Relay)
Configure IP on this interface? (Yes, No): [Yes] no

Configuring Interface 4 (WAN Frame Relay)
Configure IP on this interface? (Yes, No): [Yes] no

Configuring Interface 5 (Token Ring)
Configure IP on this interface? (Yes, No): [Yes]
IP Address: [9.37.73.147] 9.37.73.147
Address Mask: [255.255.255.0] 255.255.255.0

Configuring Interface 6 (NULL Device)
Configure IP on this interface? (Yes, No): [Yes] no

Per-Interface IP Configuration complete

Configuring IP Routing Information
Enable Dynamic Routing? (Yes, No): [Yes] no

Only Static Routing Enabled

Routing Configuration Complete

Configuring SNMP Information

SNMP will be configured with the following parameters:

     Community: public
     Access:    read_trap

If you plan to use the graphical configuration tool
to download a configuration, it requires the definition
of a community name with read_write_trap access.

Define community with read_write_trap access ? (Yes, No): [Yes] no

SNMP Configuration Complete

This is the information you have entered:

     Interface #     IP Address         Address Mask
          0          9.37.67.33         255.255.255.0
          5          9.37.73.147        255.255.255.0

Only STATIC Routing present.
```

```
SNMP has been configured with the following parameters:

     Community: public
     Access:    read_trap

If you plan to use the graphical configuration tool to
download a configuration, you will need to use the SNMP configuration
environment to define a community name with read_write_trap access.


Save this configuration? (Yes, No): [Yes]

IP configuration saved

Configure IPX? (Yes, No): [Yes] no

***********************************************************
Booting Configuration
***********************************************************

Type ’Yes’ to Configure Booting
Type ’No’ to skip Booting Configuration
Type ’Quit’ to exit Quick Config

Configure Booting? (Yes, No, Quit): [Yes]
Type ’r’ any time at this level to restart Booting Configuration

Previous Boot information

Booting Method:       IBD Boot
IBD Load Name:        MRS33.img

Create an IBD boot record using this information? (Yes, No): [Yes]

Boot configuration saved

***********************************************************
Service Port Configuration
***********************************************************

Type 'Yes' to Configure Service Ports
Type 'No' to skip Service Ports Configuration
Type 'Quit' to exit Quick Config

Configure service port? (Yes, No, Quit): [Yes] no

Quick Config Done
Restart the router for this configuration to take effect.

Restart the router? (Yes, No): [Yes]


RESTARTING THE ROUTER........
```

```
Copyright Notices:

Licensed Materials - Property of IBM
Multiprotocol Routing Services
(C) Copyright IBM Corp. 1996, 1999
All Rights Reserved. US Gov. Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.


MOS Operator Console

NOTE: The cc folks have decided that Command Completion will be DISABLED
BY DEFAULT for existing customers, and enabled by default only for new configs.
If you want Command Completion, you must explicitly enable it. (Defect 34288)
This warning message is only temporary, but the 'tip' below will appear
to the customer with an existing config.  Kind of like online Release Notes.

Command Completion is currently DISABLED.  To enable this option,
enter 'CONFIGURATION', then 'ENABLE COMMAND-COMPLETION', then Control-P.
For help using the Command Line Interface, press ESCAPE, then '?'.

Data_Center1 *
```

7.2.2.2 Configuration for a Transparent Proxy Environment

To implement transparent proxy, we must configure the router to direct all port 80 traffic to the WTE transparent proxy server. We will list the steps we performed to configure the IBM 2210 Router for transparent proxy. We will configure the internal traffic (9.37.67.*x*) on the IN0 network interface and the external traffic (9.37.73.*x*) on the IN5 network interface of the router.

The following session screen shows the sequence of configuration steps that we performed to configure the IBM 2210 Router for transparent proxy. Once again, user inputs are highlighted in boldfaced text.

```
Data_Center1*t 6
Gateway user configuration
Data_Center1 Config>p ip
Internet protocol user configuration
Data_Center1 IP config>li add                     # List the Router Addresses
IP addresses for each interface:
   intf     0   9.37.67.33       255.255.255.0    Local wire broadcast, fill 1
   intf     1                                     IP disabled on this interface
   intf     2                                     IP disabled on this interface
   intf     3                                     IP disabled on this interface
   intf     4                                     IP disabled on this interface
   intf     5   9.37.73.147      255.255.255.0    Local wire broadcast, fill 1
   intf     6                                     IP disabled on this interface
Internal IP address: 10.10.100.1
Data_Center1 IP config>set acc on
Data_Center1 IP config>add pac                    # Add Filter
Packet-filter name []? in0
Filter incoming or outgoing traffic? [IN]?
Which interface is this filter for [0]?
Data_Center1 IP config>update pac in0             # Define Filter
Data_Center1 Packet-filter 'in0' Config>add acc   # Add access for Port 80
Access Control type [E]? i
Internet source [0.0.0.0]?
Source mask [0.0.0.0]?
Internet destination [0.0.0.0]?
Destination mask [0.0.0.0]?
Starting protocol number ([0] for all protocols) [0]? 6
Ending protocol number [6]?
Starting DESTINATION port number ([0] for all ports) [0]? 80
Ending DESTINATION port number [80]?
Starting SOURCE port number ([0] for all ports) [0]?
TOS/Precedence filter mask (00-FF - [0] for none) [0]?
TOS/Precedence modification mask (00-FF - [0] for none) [0]?
Use policy-based routing? [No]: Y
Next hop gateway address []? 9.37.73.122
Use default route if next hop gateway unreachable? [Yes]: n
Enable logging? [No]:
Data_Center1 Packet-filter 'in0' Config>
Data_Center1 Packet-filter 'in0' Config>lis acc   # List Filter Rules
Access Control is: enabled
Access Control facility: USER
List of access control records:

1   Type=I    Source=0.0.0.0          Dest  =0.0.0.0          Prot=  6
              SMask =0.0.0.0          DMask =0.0.0.0
              SPorts=    0-65535      DPorts=   80-80
                                      Log=N
              PbrGw=9.37.73.122       UseDefRte=N
Data_Center1 Packet-filter 'in0' Config>
Data_Center1 Packet-filter 'in0' Config>add acc   #Add access for all other Ports
Access Control type [E]? i
Internet source [0.0.0.0]?
Source mask [0.0.0.0]?
Internet destination [0.0.0.0]?
Destination mask [0.0.0.0]?
Starting protocol number ([0] for all protocols) [0]?
Starting DESTINATION port number ([0] for all ports) [0]?
Starting SOURCE port number ([0] for all ports) [0]?
```

```
Filter on ICMP Type ([-1] for all types) [-1]?
TOS/Precedence filter mask (00-FF - [0] for none) [0]?
TOS/Precedence modification mask (00-FF - [0] for none) [0]?
Use policy-based routing? [No]:
Enable logging? [No]:
Data_Center1 Packet-filter 'in0' Config>
Data_Center1 Packet-filter 'in0' Config>li acc
Access Control is: enabled
Access Control facility: USER
List of access control records:

1   Type=I    Source=0.0.0.0          Dest  =0.0.0.0          Prot=  6
              SMask =0.0.0.0          DMask =0.0.0.0
              SPorts=    0-65535      DPorts=   80-80
                                      Log=N
              PbrGw=9.37.73.122       UseDefRte=N

2   Type=I    Source=0.0.0.0          Dest  =0.0.0.0          Prot=  0-255
              SMask =0.0.0.0          DMask =0.0.0.0
              SPorts=    0-65535      DPorts=    0-65535
                      T/C= **/**      Log=N
Data_Center1 Packet-filter 'in0' Config>exit
Data_Center1 IP config>
Data_Center1 IP config>li pac
Name               Dir  Intf  State  Src-Addr-Ver
in0                In   0     On     Off

Access Control is: enabled

Data_Center1 IP Config> CTRL ^P                    # Enter Control P
Data_Center1 *restart
Are you sure you want to restart the gateway? (Yes or [No]): y

Copyright Notices:
Licensed Materials - Property of IBM Multiprotocol Routing Services(C) Copyright I
1996, 1999
All Rights Reserved. US Gov. Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.


MOS Operator Console
NOTE: The cc folks have decided that Command Completion will be DISABLED
BY DEFAULT for existing customers, and enabled by default only for new configs.
If you want Command Completion, you must explicitly enable it. (Defect 34288)
This warning message is only temporary, but the 'tip' below will appear
to the customer with an existing config.  Kind of like online Release Notes.

ommand Completion is currently DISABLED.  To enable this option,
enter 'CONFIGURATION', then 'ENABLE COMMAND-COMPLETION', then Control-P.
For help using the Command Line Interface, press ESCAPE, then '?'.

Data_Center1 *
```

7.2.3 How to Set Up WTE Transparent Proxy

Transparent proxy is as powerful as it is simple to configure. It can be configured using the Configuration and Administration forms or by manually editing the ibmproxy.conf file.

We will configure the WTE as a transparent proxy using the Configuration and Administration forms. On the navigation menu:

1. Click **Proxy Configuration** folder.
2. Click **Proxy Performance**.

You will be presented with the Proxy Performance settings form as shown in the following figure:

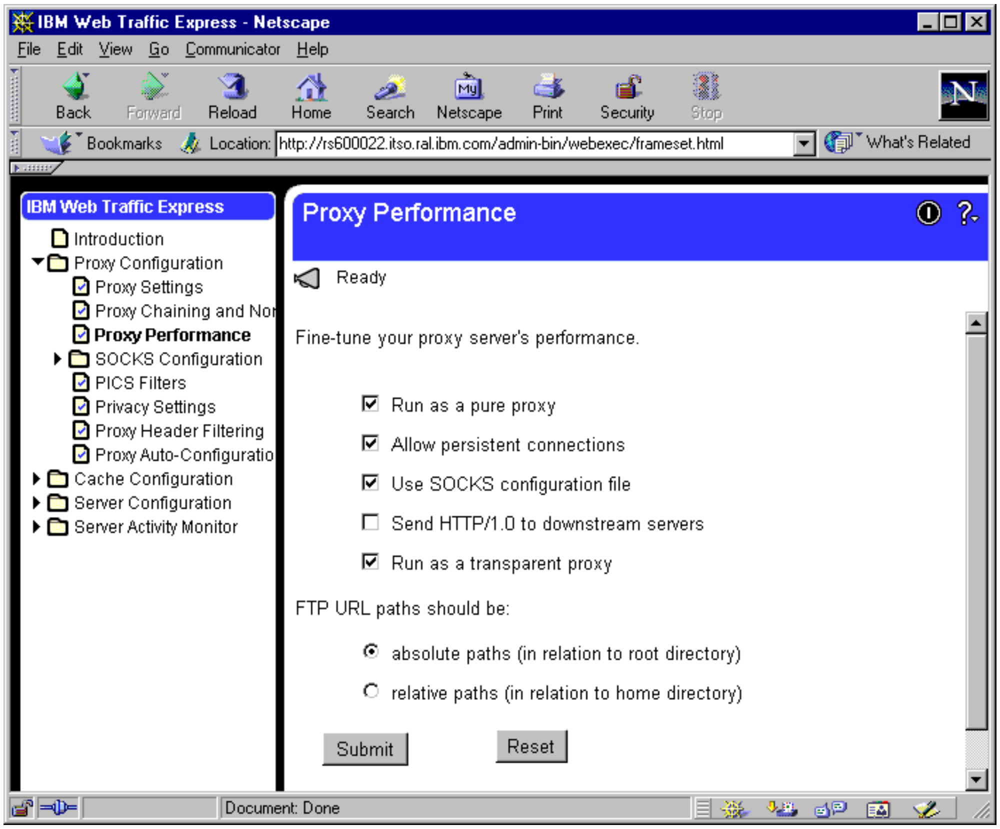

Fine-tune your proxy server's performance.

☑ Run as a pure proxy

☑ Allow persistent connections

☑ Use SOCKS configuration file

□ Send HTTP/1.0 to downstream servers

☑ Run as a transparent proxy

FTP URL paths should be:

● absolute paths (in relation to root directory)

○ relative paths (in relation to home directory)

Figure 85. Proxy Performance - Run as a Transparent Proxy

On the Performance Settings form:

1. Click the **Run as a transparent proxy** check box.
2. Click **Submit** to update the ibmproxy.conf file.
3. Click the **Restart** icon to restart the WTE server.

In order for the configuration above to be successful, verify that in the WTE ibmproxy.conf configuration file, `BindSpecific` is set to `Off` and `Port` is set to `80`.

7.2.4 Transparent Proxy Experience

We implement the transparent proxy scenario with the environment set up as shown on Figure 83 on page 133. As we wanted to monitor the proxy access on the transparent proxy server, we use the AIX `tail` command and the WTE Proxy Access Statistics form as follows:

1. On the WTE proxy server, rs600022, we monitor the proxy access using the following command:

```
tail -f /wte/logs/httpd-proxy.Mar261999
```

2. On the Web server, wtr05311, we set up the Windows NT Network Monitor to monitor the incoming network traffic.
3. On the Web client, pohyt, we entered the HTTP request http://9.37.73.120/.
4. On the Web client, pohyt, we tested the FTP protocol by entering the FTP request ftp://9.37.73.120/.

We successfully received the Web page (through the HTTP protocol) and the file request (through the FTP protocol) that we requested in Step 3 on page 144 and Step 4 on page 144. On the WTE proxy server, we confirmed that the HTTP requests were coming from the Web client, pohyt, whose IP address was 9.37.67.37. The FTP requests that we made were *never* registered by the proxy server. The following is the WTE proxy server log:

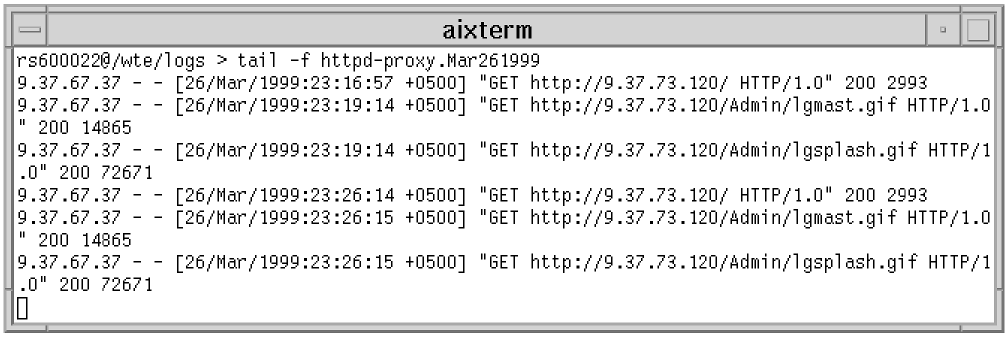

```
rs600022@/wte/logs > tail -f httpd-proxy.Mar261999
9.37.67.37 - - [26/Mar/1999:23:16:57 +0500] "GET http://9.37.73.120/ HTTP/1.0" 200 2993
9.37.67.37 - - [26/Mar/1999:23:19:14 +0500] "GET http://9.37.73.120/Admin/lgmast.gif HTTP/1.0
" 200 14865
9.37.67.37 - - [26/Mar/1999:23:19:14 +0500] "GET http://9.37.73.120/Admin/lgsplash.gif HTTP/1
.0" 200 72671
9.37.67.37 - - [26/Mar/1999:23:26:14 +0500] "GET http://9.37.73.120/ HTTP/1.0" 200 2993
9.37.67.37 - - [26/Mar/1999:23:26:15 +0500] "GET http://9.37.73.120/Admin/lgmast.gif HTTP/1.0
" 200 14865
9.37.67.37 - - [26/Mar/1999:23:26:15 +0500] "GET http://9.37.73.120/Admin/lgsplash.gif HTTP/1
.0" 200 72671
```

Figure 86. WTE Proxy Server Log

On the Web server, the following is the network traffic captured by the Windows NT Network Monitor while the HTTP requests were being received and served. It shows that the HTTP requests were coming from the WTE server:

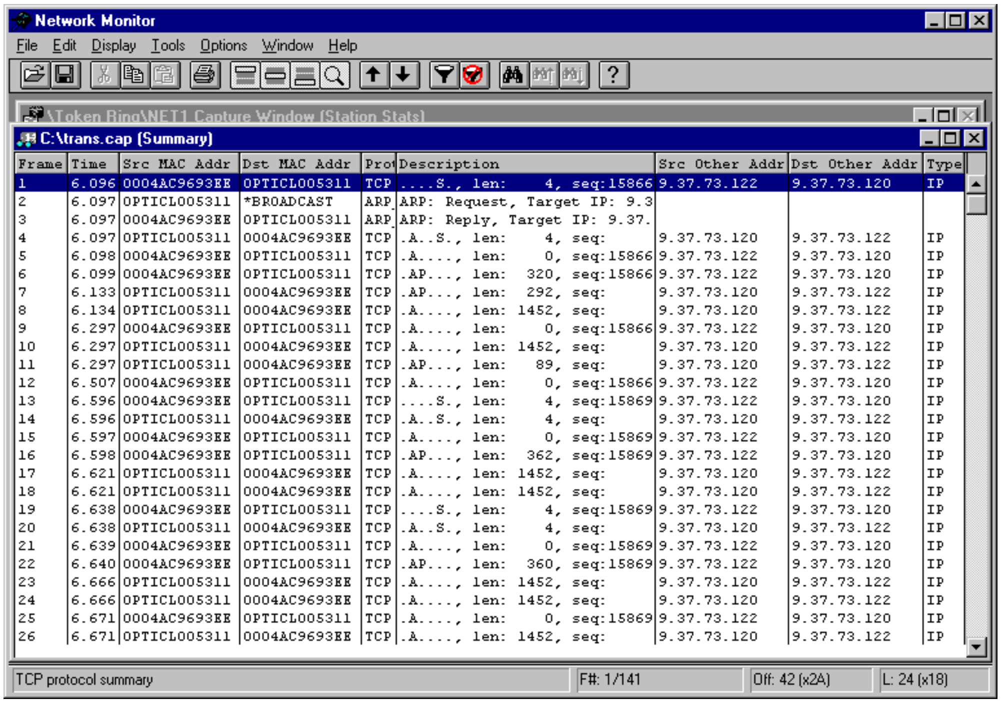

Frame	Time	Src MAC Addr	Dst MAC Addr	Prot	Description	Src Other Addr	Dst Other Addr	Type
1	6.096	0004AC9693EE	OPTICL005311	TCP	....S., len: 4, seq:15866	9.37.73.122	9.37.73.120	IP
2	6.097	OPTICL005311	*BROADCAST	ARP	ARP: Request, Target IP: 9.3			
3	6.097	0004AC9693EE	OPTICL005311	ARP	ARP: Reply, Target IP: 9.37.			
4	6.097	OPTICL005311	0004AC9693EE	TCP	.A..S., len: 4, seq:	9.37.73.120	9.37.73.122	IP
5	6.098	0004AC9693EE	OPTICL005311	TCP	.A...., len: 0, seq:15866	9.37.73.122	9.37.73.120	IP
6	6.099	0004AC9693EE	OPTICL005311	TCP	.AP..., len: 320, seq:15866	9.37.73.122	9.37.73.120	IP
7	6.133	OPTICL005311	0004AC9693EE	TCP	.AP..., len: 292, seq:	9.37.73.120	9.37.73.122	IP
8	6.134	OPTICL005311	0004AC9693EE	TCP	.A...., len: 1452, seq:	9.37.73.120	9.37.73.122	IP
9	6.297	0004AC9693EE	OPTICL005311	TCP	.A...., len: 0, seq:15866	9.37.73.122	9.37.73.120	IP
10	6.297	OPTICL005311	0004AC9693EE	TCP	.A...., len: 1452, seq:	9.37.73.120	9.37.73.122	IP
11	6.297	OPTICL005311	0004AC9693EE	TCP	.AP..., len: 89, seq:	9.37.73.120	9.37.73.122	IP
12	6.507	0004AC9693EE	OPTICL005311	TCP	.A...., len: 0, seq:15866	9.37.73.122	9.37.73.120	IP
13	6.596	0004AC9693EE	OPTICL005311	TCP	....S., len: 4, seq:15869	9.37.73.122	9.37.73.120	IP
14	6.596	OPTICL005311	0004AC9693EE	TCP	.A..S., len: 4, seq:	9.37.73.120	9.37.73.122	IP
15	6.597	0004AC9693EE	OPTICL005311	TCP	.A...., len: 0, seq:15869	9.37.73.122	9.37.73.120	IP
16	6.598	0004AC9693EE	OPTICL005311	TCP	.AP..., len: 362, seq:15869	9.37.73.122	9.37.73.120	IP
17	6.621	OPTICL005311	0004AC9693EE	TCP	.A...., len: 1452, seq:	9.37.73.120	9.37.73.122	IP
18	6.621	OPTICL005311	0004AC9693EE	TCP	.A...., len: 1452, seq:	9.37.73.120	9.37.73.122	IP
19	6.638	0004AC9693EE	OPTICL005311	TCP	....S., len: 4, seq:15869	9.37.73.122	9.37.73.120	IP
20	6.638	OPTICL005311	0004AC9693EE	TCP	.A..S., len: 4, seq:	9.37.73.120	9.37.73.122	IP
21	6.639	0004AC9693EE	OPTICL005311	TCP	.A...., len: 0, seq:15869	9.37.73.122	9.37.73.120	IP
22	6.640	0004AC9693EE	OPTICL005311	TCP	.AP..., len: 360, seq:15869	9.37.73.122	9.37.73.120	IP
23	6.666	OPTICL005311	0004AC9693EE	TCP	.A...., len: 1452, seq:	9.37.73.120	9.37.73.122	IP
24	6.666	OPTICL005311	0004AC9693EE	TCP	.A...., len: 1452, seq:	9.37.73.120	9.37.73.122	IP
25	6.671	0004AC9693EE	OPTICL005311	TCP	.A...., len: 0, seq:15869	9.37.73.122	9.37.73.120	IP
26	6.671	OPTICL005311	0004AC9693EE	TCP	.A...., len: 1452, seq:	9.37.73.120	9.37.73.122	IP

Figure 87. NT Network Monitor - http Requests

On the Web server, the network traffic captured while the FTP requests were being issued shows that the requests did not go via the WTE server. They passed through the router directly.

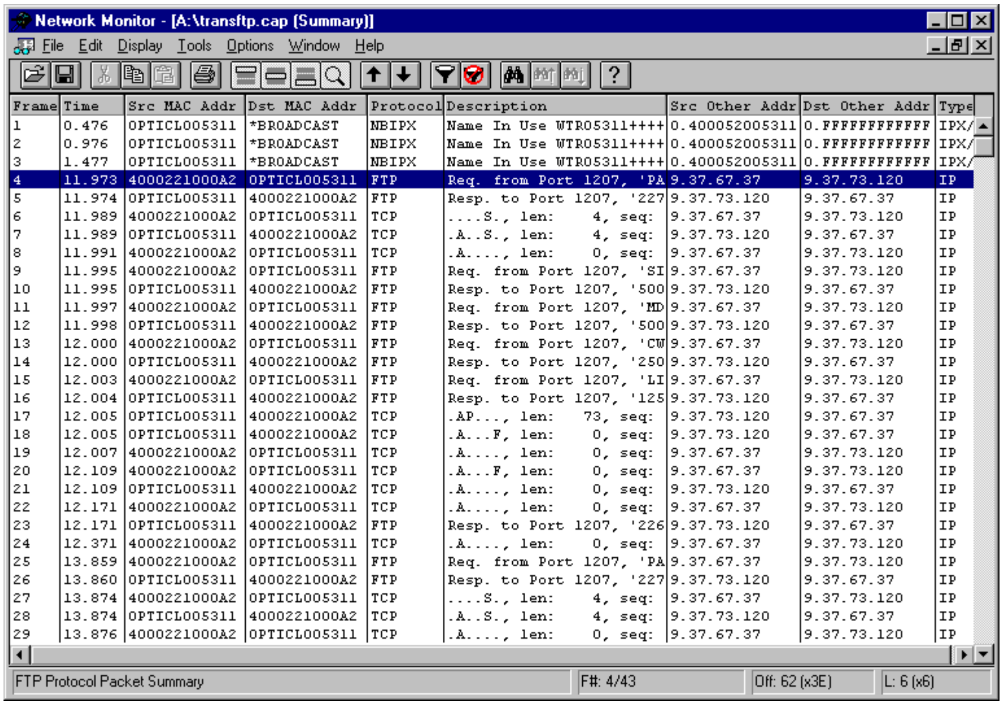
Network Monitor - [A:\transftp.cap (Summary)]
File Edit Display Tools Options Window Help

Frame	Time	Src MAC Addr	Dst MAC Addr	Protocol	Description	Src Other Addr	Dst Other Addr	Type
1	0.476	OPTICL005311	*BROADCAST	NBIPX	Name In Use WTR05311++++	0.400052005311	0.FFFFFFFFFFFF	IPX/
2	0.976	OPTICL005311	*BROADCAST	NBIPX	Name In Use WTR05311++++	0.400052005311	0.FFFFFFFFFFFF	IPX/
3	1.477	OPTICL005311	*BROADCAST	NBIPX	Name In Use WTR05311++++	0.400052005311	0.FFFFFFFFFFFF	IPX/
4	11.973	4000221000A2	OPTICL005311	FTP	Req. from Port 1207, 'PA	9.37.67.37	9.37.73.120	IP
5	11.974	OPTICL005311	4000221000A2	FTP	Resp. to Port 1207, '227	9.37.73.120	9.37.67.37	IP
6	11.989	4000221000A2	OPTICL005311	TCP	....S., len: 4, seq:	9.37.67.37	9.37.73.120	IP
7	11.989	OPTICL005311	4000221000A2	TCP	.A..S., len: 4, seq:	9.37.73.120	9.37.67.37	IP
8	11.991	4000221000A2	OPTICL005311	TCP	.A...., len: 0, seq:	9.37.67.37	9.37.73.120	IP
9	11.995	4000221000A2	OPTICL005311	FTP	Req. from Port 1207, 'SI	9.37.67.37	9.37.73.120	IP
10	11.995	OPTICL005311	4000221000A2	FTP	Resp. to Port 1207, '500	9.37.73.120	9.37.67.37	IP
11	11.997	4000221000A2	OPTICL005311	FTP	Req. from Port 1207, 'MD	9.37.67.37	9.37.73.120	IP
12	11.998	OPTICL005311	4000221000A2	FTP	Resp. to Port 1207, '500	9.37.73.120	9.37.67.37	IP
13	12.000	4000221000A2	OPTICL005311	FTP	Req. from Port 1207, 'CW	9.37.67.37	9.37.73.120	IP
14	12.000	OPTICL005311	4000221000A2	FTP	Resp. to Port 1207, '250	9.37.73.120	9.37.67.37	IP
15	12.003	4000221000A2	OPTICL005311	FTP	Req. from Port 1207, 'LI	9.37.67.37	9.37.73.120	IP
16	12.004	OPTICL005311	4000221000A2	FTP	Resp. to Port 1207, '125	9.37.73.120	9.37.67.37	IP
17	12.005	OPTICL005311	4000221000A2	TCP	.AP..., len: 73, seq:	9.37.73.120	9.37.67.37	IP
18	12.005	OPTICL005311	4000221000A2	TCP	.A...F, len: 0, seq:	9.37.73.120	9.37.67.37	IP
19	12.007	4000221000A2	OPTICL005311	TCP	.A...., len: 0, seq:	9.37.67.37	9.37.73.120	IP
20	12.109	4000221000A2	OPTICL005311	TCP	.A...F, len: 0, seq:	9.37.67.37	9.37.73.120	IP
21	12.109	OPTICL005311	4000221000A2	TCP	.A...., len: 0, seq:	9.37.73.120	9.37.67.37	IP
22	12.171	4000221000A2	OPTICL005311	TCP	.A...., len: 0, seq:	9.37.67.37	9.37.73.120	IP
23	12.171	OPTICL005311	4000221000A2	FTP	Resp. to Port 1207, '226	9.37.73.120	9.37.67.37	IP
24	12.371	4000221000A2	OPTICL005311	TCP	.A...., len: 0, seq:	9.37.67.37	9.37.73.120	IP
25	13.859	4000221000A2	OPTICL005311	FTP	Req. from Port 1207, 'PA	9.37.67.37	9.37.73.120	IP
26	13.860	OPTICL005311	4000221000A2	FTP	Resp. to Port 1207, '227	9.37.73.120	9.37.67.37	IP
27	13.874	4000221000A2	OPTICL005311	TCP	....S., len: 4, seq:	9.37.67.37	9.37.73.120	IP
28	13.874	OPTICL005311	4000221000A2	TCP	.A..S., len: 4, seq:	9.37.73.120	9.37.67.37	IP
29	13.876	4000221000A2	OPTICL005311	TCP	.A...., len: 0, seq:	9.37.67.37	9.37.73.120	IP

FTP Protocol Packet Summary | F#: 4/43 | Off: 62 (x3E) | L: 6 (x6)

Figure 88. NT Network Monitor - ftp Requests

Hence, we received the confirmation that we had successfully implemented a WTE transparent proxy server. This completes our example on the WTE transparent proxy scenario.

Chapter 8. WTE Proxy Autoconfiguration

IBM Web Traffic Express (WTE) Version 2, the caching and filtering component of IBM WebSphere Performance Pack Version 2, now supports *automatic proxy configuration*, a technology that matches a feature implemented in Netscape Navigator 2.0 or later and Microsoft Internet Explorer Version 4.0 or later. Automatic proxy configuration implies that the client browser does not need to be configured to point to a specific proxy or SOCKS server, as long as it points to a *proxy automatic configuration* (PAC) file.

The advantage of this feature is that it allows the server administrator to update only this file if there is any change to the proxy or SOCKS server configuration. Client browsers need not be re-configured. This feature can also be used to reroute requests when servers are down, to balance workload or to send specific URLs to specific proxies. Note that *new PAC files are only reloaded when a browser is restarted*.

8.1 How Proxy Autoconfiguration Works

Automatic proxy configuration is a browser function that enables more dynamic server selection. A PAC file is a JavaScript file that consists of functions that get called by the client browser before each URL retrieval. The function will return values indicating whether a proxy server, a SOCKS server, or a direct connection will be used to service the request. This file can also redirect the request if the initial connection to be used is down.

When a client's browser is set to autoproxy, it will call the JavaScript PAC file each time a URL is requested by the user. Below is an example of a PAC file:

```
//Created by WTE remote config-DO NOT EDIT
//PRI PROXY rs60022.itso.ral.ibm.com 80 END
//SEC PROXY wtr05178.itso.ral.ibm.com 80 END
//DIR yes END
function FindProxyForURL(url, host)
  {
  return "PROXY rs60022.itso.ral.ibm.com:80; " +
         "PROXY wtr05178.itso.ral.ibm.com:80; " +
         "DIRECT";
  }
```

Notice that a PAC file has a .pac extension.

This script will instruct the browser to use:

- The machine rs600022.itso.ral.ibm.com listening on port 80 as the primary proxy
- The machine wtr05178.itso.ral.ibm.com listening on port 80 as the secondary proxy
- A direct connection to the Internet if both proxies fail

In the scenario represented in Figure 89, you can see the path of the traffic when the primary proxy is functional:

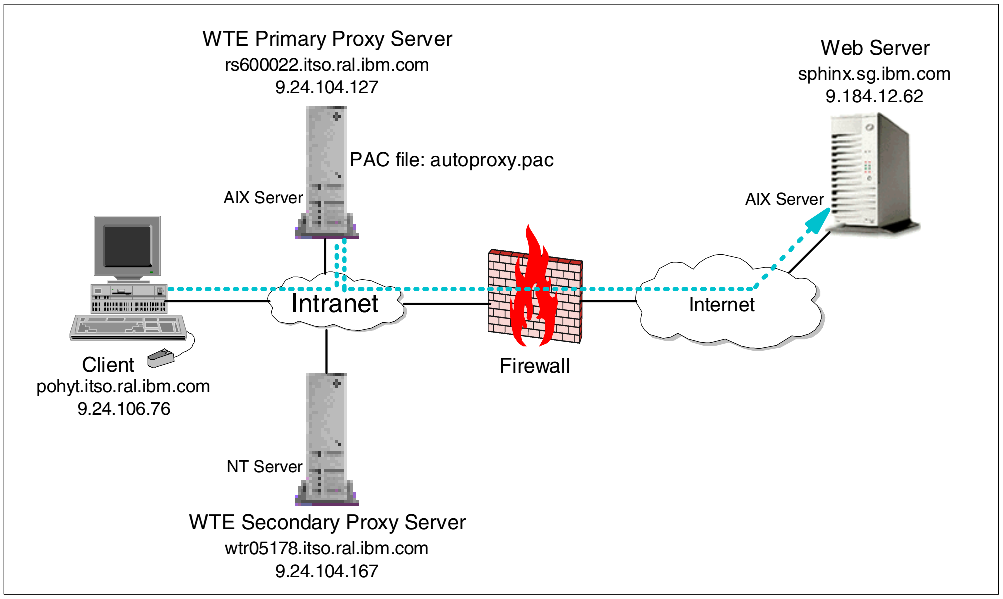

Figure 89. Automatic Proxy - Using the Primary Proxy Server

Should the primary proxy fail, the client browser will automatically point to the secondary proxy server indicated in the PAC file and re-direct the traffic to it as shown in Figure 90 on page 149. The client browser will record the fact that the primary proxy server is down and will continue to use the secondary proxy. If both the primary and secondary proxies are unavailable, the browser will attempt to connect using a direct connection to the Internet.

Using IBM SecureWay Network DIspatcher

In such a scenario, should the primary proxy server fail, newly joined clients will not be able to retrieve the PAC file. Only the browsers that connected before to the primary proxy will have the PAC file, and will be able to connect to the secondary proxy server.

To avoid this limitation, we recommend using a machine running IBM SecureWay Network Dispatcher (ND) to load balance between two or more proxy servers. The ND software is able to understand when one or more servers in the cluster are down, and will temporarily exclude them from the cluster. This way, ND can ensure proxy server load balancing and high availability.

In order to configure a stable environment, we recommend the use of two ND machines running in high availability mode.

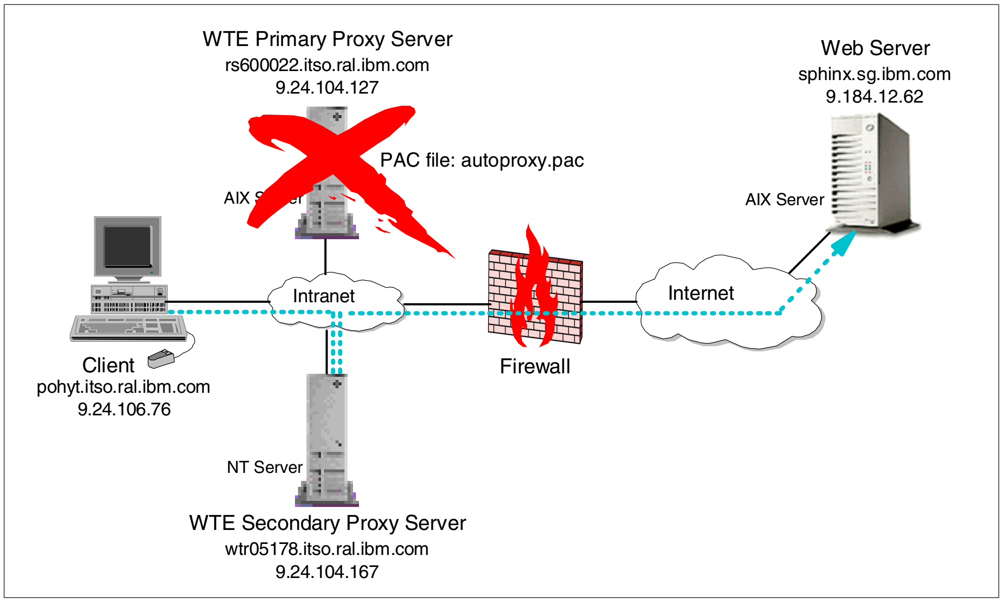

Figure 90. Automatic Proxy - Using the Secondary Proxy Server

8.2 How to Set Up Automatic Proxy

Proxy autoconfiguration can be configured using the Configuration and Administration forms. We recommend verifying that both the primary and secondary proxies are functioning correctly:

1. Set your browser to point to the primary proxy server.
2. Point to a Web server that is outside the firewall.
3. Repeat Step 1 and Step 2 but point to the secondary proxy server.

After verifying that both the proxies are functioning, on the navigation menu:

1. Click **Proxy Configuration** folder.
2. Click **Proxy Auto-Configuration**.
3. Click the **Create a new file** radio button.
4. Enter the PAC file name in the appropriate field, without forgetting that the file name must have a .pac extension.

We entered `autoproxy.pac` as our proxy autoconfiguration file name. Figure 91 shows the Proxy Auto-Configuration form:

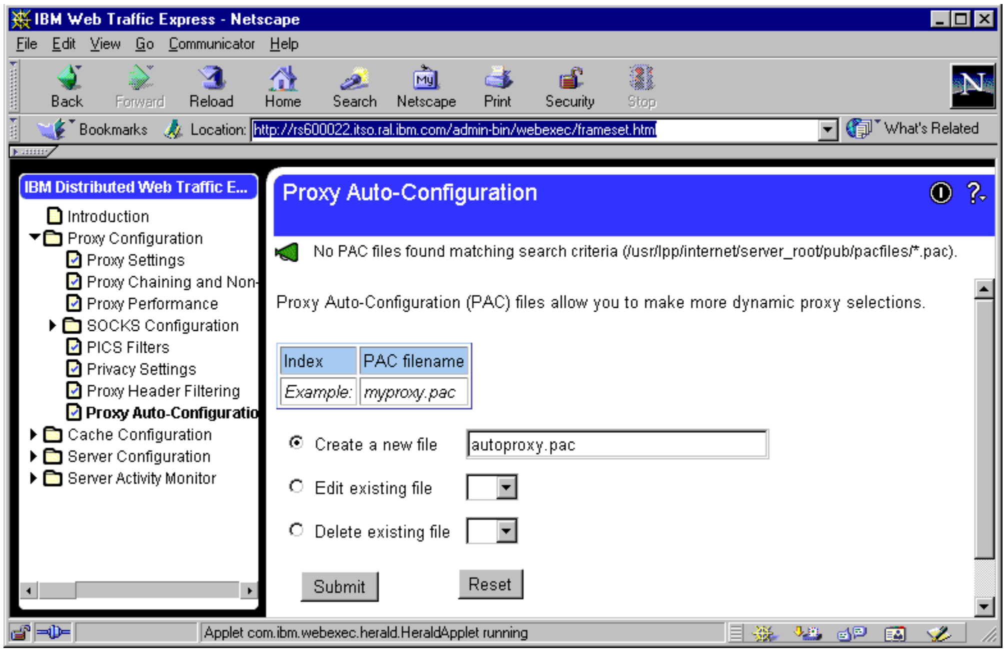
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Distributed Web Traffic E...
Introduction
Proxy Configuration
Proxy Settings
Proxy Chaining and Non-
Proxy Performance
SOCKS Configuration
PICS Filters
Privacy Settings
Proxy Header Filtering
Proxy Auto-Configuratio
Cache Configuration
Server Configuration
Server Activity Monitor
Submit
Reset
Applet com.ibm.webexec.herald.HeraldApplet running

Proxy Auto-Configuration

No PAC files found matching search criteria (/usr/lpp/internet/server_root/pub/pacfiles/*.pac).

Proxy Auto-Configuration (PAC) files allow you to make more dynamic proxy selections.

Index	PAC filename
Example:	*myproxy.pac*

● Create a new file autoproxy.pac

○ Edit existing file

○ Delete existing file

Figure 91. Proxy Configuration - Proxy Auto-Configuration Form 1

Click the **Submit** button to create the PAC file. After clicking the **Submit** button, the next form requires you to enter host name and port number for the primary and secondary proxy or SOCKS servers. Mark the **Route DIRECT** check box if you want the browser to attempt a direct connection if both Proxies/SOCKS servers are down. We entered the primary and secondary server host names and port numbers as shown in Figure 92 on page 150:

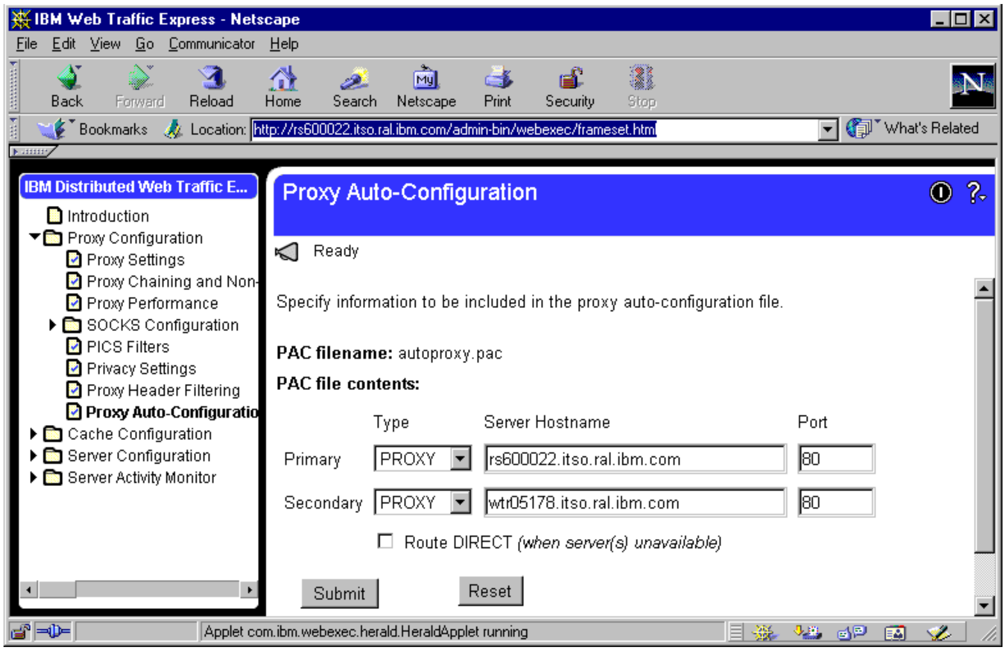
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Distributed Web Traffic E...
Introduction
Proxy Configuration
Proxy Settings
Proxy Chaining and Non-
Proxy Performance
SOCKS Configuration
PICS Filters
Privacy Settings
Proxy Header Filtering
Proxy Auto-Configuratio
Cache Configuration
Server Configuration
Server Activity Monitor
Submit
Reset
Applet com.ibm.webexec.herald.HeraldApplet running

Proxy Auto-Configuration

Ready

Specify information to be included in the proxy auto-configuration file.

PAC filename: autoproxy.pac

PAC file contents:

	Type	Server Hostname	Port
Primary	PROXY	rs600022.itso.ral.ibm.com	80
Secondary	PROXY	wtr05178.itso.ral.ibm.com	80

□ Route DIRECT *(when server(s) unavailable)*

Figure 92. Proxy Configuration - Proxy Auto-Configuration Form 2

Click the **Submit** button to update the autoproxy.pac file. After processing the form, WTE displays the PAC file name in the Proxy Auto-Configuration form, as shown in Figure 93:

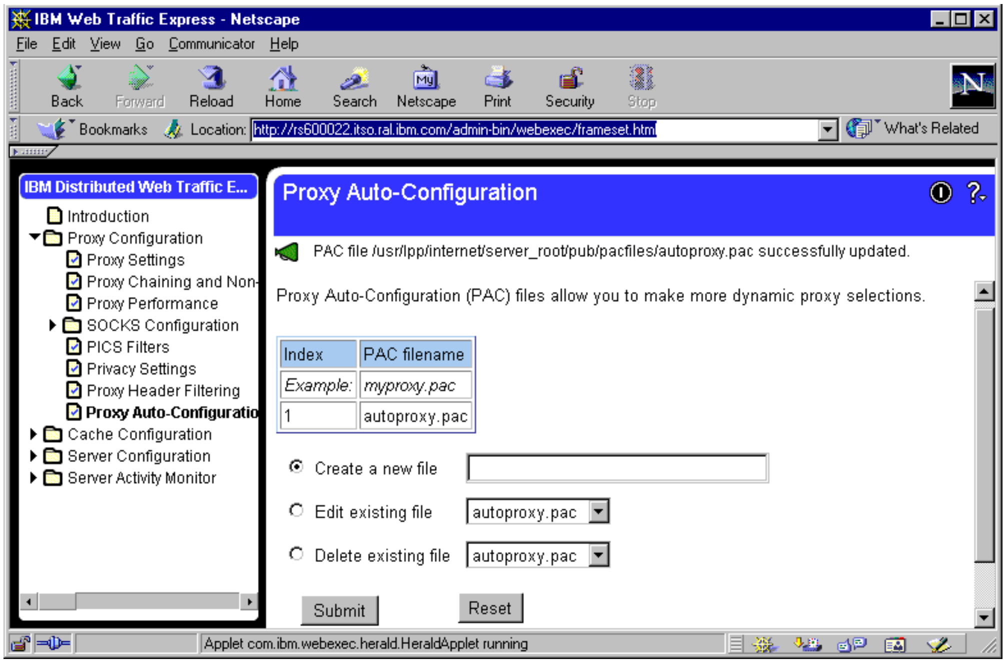
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Distributed Web Traffic E...
Introduction
Proxy Configuration
Proxy Settings
Proxy Chaining and Non-
Proxy Performance
SOCKS Configuration
PICS Filters
Privacy Settings
Proxy Header Filtering
Proxy Auto-Configuratio
Cache Configuration
Server Configuration
Server Activity Monitor
Proxy Auto-Configuration
Submit
Reset
Applet com.ibm.webexec.herald.HeraldApplet running

PAC file /usr/lpp/internet/server_root/pub/pacfiles/autoproxy.pac successfully updated.

Proxy Auto-Configuration (PAC) files allow you to make more dynamic proxy selections.

Index	PAC filename
Example:	*myproxy.pac*
1	autoproxy.pac

● Create a new file

○ Edit existing file autoproxy.pac

○ Delete existing file autoproxy.pac

Figure 93. Proxy Configuration - Proxy Auto-configuration 3

8.3 Automatic Proxy Configuration Experience

To demonstrate how the proxy autoconfiguration works, we use the test environment shown in Figure 89 on page 148 and Figure 90 on page 149. For the purpose of tracing the flow of the requests, we first disable caching on both of the proxies. On the client machine, pohyt, we set the Web browser to use automatic proxy configuration. As we are using Netscape Navigator V4.5, the automatic proxy configuration is set in the Preferences menu as shown:

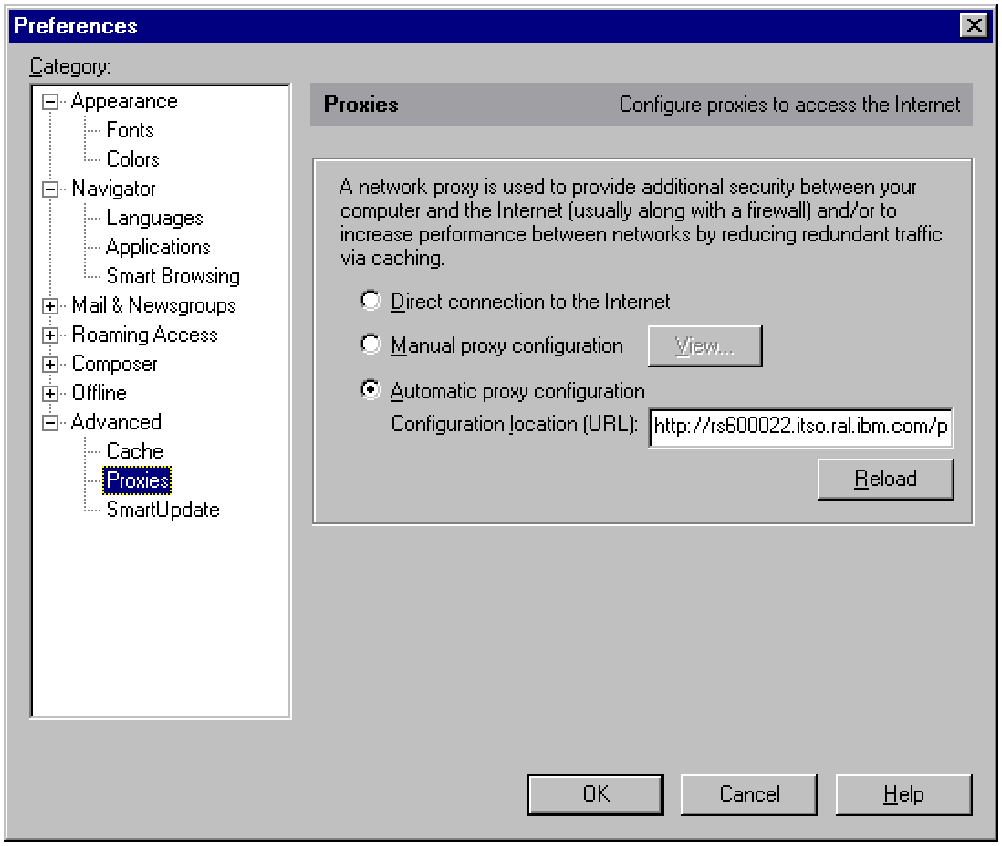

A network proxy is used to provide additional security between your computer and the Internet (usually along with a firewall) and/or to increase performance between networks by reducing redundant traffic via caching.

○ Direct connection to the Internet

○ Manual proxy configuration [View...]

● Automatic proxy configuration

Configuration location (URL): http://rs600022.itso.ral.ibm.com/p

Figure 94. Proxy Autoconfiguration - Netscape Browser Preferences Setting

We click **Automatic proxy configuration**, and enter the automatic proxy configuration URL as `http://rs6000.itso.ral.ibm.com/pacfiles/autoproxy`. Then, we click **Reload** to activate the changes.

Since we wanted to monitor the proxy access on the primary proxy server, we used the AIX `tail` command and the WTE Proxy Access Statistics form as follows:

1. On the primary proxy server, rs600022, we entered the following command:

```
tail -f /wte/logs/httpd-proxy.Mar111999
```

2. On the secondary proxy server, wtr05178, we used the WTE Proxy Access Statistics form, which can be started from the WTE Configuration and Administration Forms.

3. On the Web server machine sphinx, located in Singapore, we entered the following command:

```
tail -f /usr/lpp/internet/server_root/logs/httpd-log.Mar121999
```

8.3.1 Access the Proxies

From the client machine, pohyt, on which the proxy configuration was set to automatic proxy, we requested the URL `http://sphinx.sg.ibm.com/sphinx.html`, and the URL `http://sphinx.sg.ibm.com/sphinx11.html`. What happened to our requests?

1. The request from the Web browser was taken by the primary proxy server, rs600022, which services the URL via its designated SOCKS server to the Web server sphinx. On the primary proxy server, rs600022, we saw the two URL requests coming from IP address 9.24.106.76, which was the client machine, pohyt, as shown:

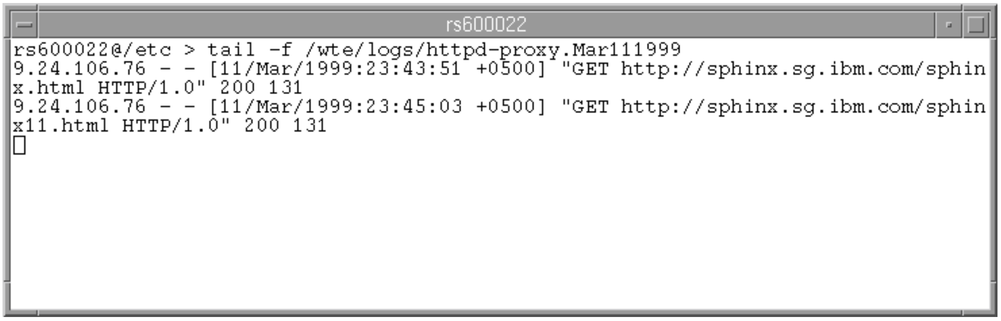

```
rs600022@/etc > tail -f /wte/logs/httpd-proxy.Mar111999
9.24.106.76 - - [11/Mar/1999:23:43:51 +0500] "GET http://sphinx.sg.ibm.com/sphin
x.html HTTP/1.0" 200 131
9.24.106.76 - - [11/Mar/1999:23:45:03 +0500] "GET http://sphinx.sg.ibm.com/sphin
x11.html HTTP/1.0" 200 131
```

Figure 95. Proxy Access Log File on the Primary Proxy Server, rs600022

2. On the Web server, sphinx, we also saw two URL requests coming from the IP address 9.3.199.22, which was the primary proxy server’s designated SOCKS server, as shown:

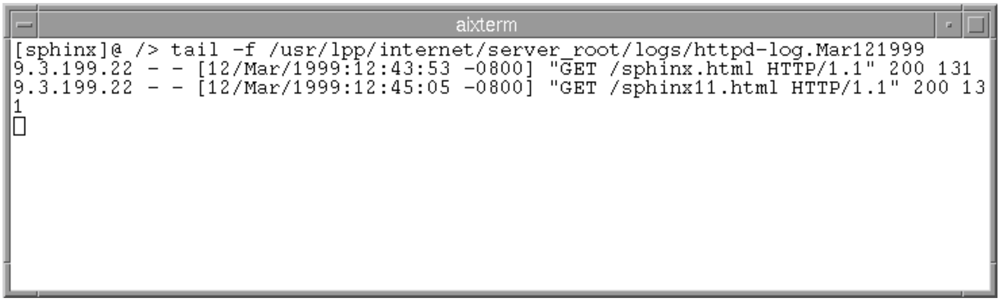

```
[sphinx]@ /> tail -f /usr/lpp/internet/server_root/logs/httpd-log.Mar121999
9.3.199.22 - - [12/Mar/1999:12:43:53 -0800] "GET /sphinx.html HTTP/1.1" 200 131
9.3.199.22 - - [12/Mar/1999:12:45:05 -0800] "GET /sphinx11.html HTTP/1.1" 200 13
1
```

Figure 96. Httpd Access Log File on Web Server, sphinx

No access was made to the secondary proxy server, wtr05178, at this time (see Figure 90 on page 149).

Next, we disable the primary proxy server, rs600022, by switching it off, thus simulating a system failure. From the client machine, pohyt, without shutting down the browser, we try again to request the URL `http://sphinx.sg.ibm.com/sphinx.html`, and the URL `http://sphinx.sg.ibm.com/sphinx11.html`.

The first request takes a while to process as the browser still tries to access the primary proxy server. After a few seconds, we can see that it proceeds to use the secondary proxy server. An explanation of the steps follows:

1. This time, the request from the Web browser was taken by the secondary proxy server, wtr05178, which services the URL via its designated SOCKS server to the Web server sphinx. On the secondary proxy server, wtr05178, we see two URL requests coming from IP address 9.24.106.76, which was the client machine, pohyt, as shown:

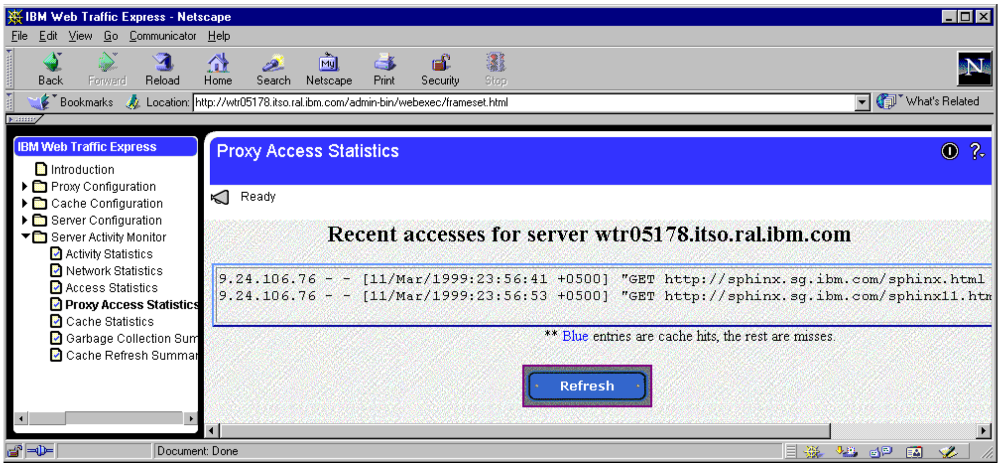

Recent accesses for server wtr05178.itso.ral.ibm.com

```
9.24.106.76 - - [11/Mar/1999:23:56:41 +0500] "GET http://sphinx.sg.ibm.com/sphinx.html
9.24.106.76 - - [11/Mar/1999:23:56:53 +0500] "GET http://sphinx.sg.ibm.com/sphinx11.htm
```

** Blue entries are cache hits, the rest are misses.

Figure 97. Proxy Access Statistics on Secondary Proxy Server, wtr05178

2. On the Web server sphinx, we also saw two more URL requests coming from the machine having IP address 9.37.3.157. This was our secondary proxy's designated SOCKS server as shown:

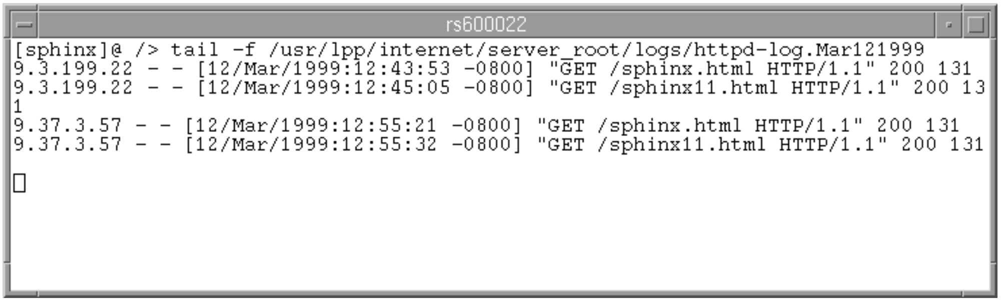

```
[sphinx]@ /> tail -f /usr/lpp/internet/server_root/logs/httpd-log.Mar121999
9.3.199.22 - - [12/Mar/1999:12:43:53 -0800] "GET /sphinx.html HTTP/1.1" 200 131
9.3.199.22 - - [12/Mar/1999:12:45:05 -0800] "GET /sphinx11.html HTTP/1.1" 200 13
1
9.37.3.57 - - [12/Mar/1999:12:55:21 -0800] "GET /sphinx.html HTTP/1.1" 200 131
9.37.3.57 - - [12/Mar/1999:12:55:32 -0800] "GET /sphinx11.html HTTP/1.1" 200 131
```

Figure 98. Httpd Access Log File on Web Server, sphinx

The example we have just described shows a clear demonstration of how WTE proxy automatic configuration works.

Chapter 9. WTE Flexible-Client SOCKS Support

Flexible-client SOCKS in IBM Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2, allows the caching and filtering proxy server to reside behind a firewall or SOCKS server without sharing the same physical machine.

In this chapter we examine the flexible-client SOCKS support offered by WTE and we demonstrate how to set up a flexible-client SOCKS environment.

9.1 WTE vs Traditional Proxy Servers

With flexible-client SOCKS support, the WTE server does not have to forward all the requests to the SOCKS server. On the contrary, requests for Web content located in the intranet can be routed directly to the destination content server. This way, workload on the SOCKS server is dramatically reduced, and the end user will experience a better response time.

This is not the case with a traditional proxy server, which has to forward all the requests to the firewall. It is the firewall, then, that forwards the requests to the content server, regardless of whether the content server is located in the intranet or in the Internet. For this reason, network congestion increases, and the workload on the firewall machine is higher. Figure 99 on page 155 shows the way a traditional proxy server works.

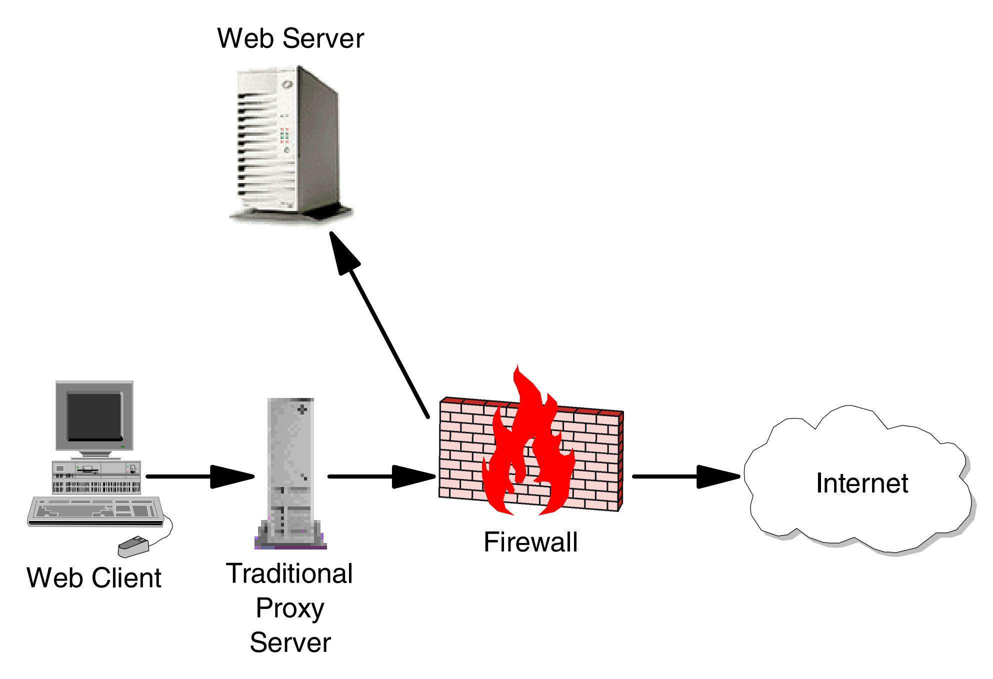

Figure 99. Traditional Proxy Server

The flexible-client SOCKS enhancement, implemented in WTE, improves security by allowing a firewall server to be isolated from the proxy server. It reduces the load on the firewall as the internal requests are handled by the proxy server. Moreover, the administrator can easily specify which requests the proxy server sends to the SOCKS server and which requests it redirects back to the local domain.

Note

WTE is a SOCKS4 client, and it does not include a SOCKS server. A SOCKS server is included in the IBM eNetwork Firewall.

Figure 100 shows an example of the flexible-client SOCKS implementation, which allows requests for the local intranet to be routed directly to the destination content server; only requests for Internet servers are routed to the SOCKS server.

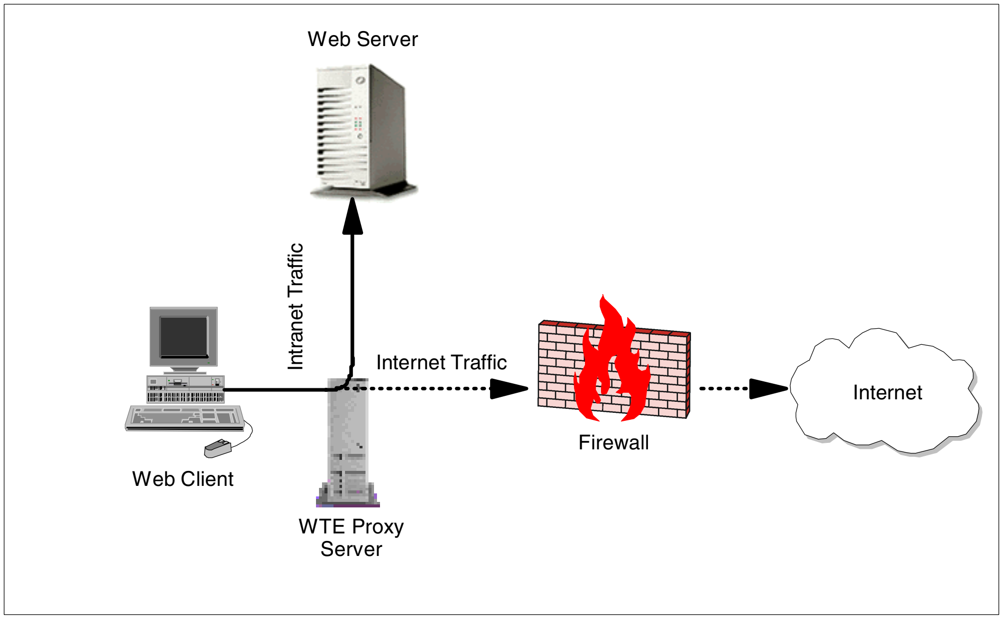

Figure 100. WTE with Flexible-Client SOCKS

The flexible-client SOCKS provided by WTE allows you to specify which IP addresses or domains should be contacted directly by the proxy server and which ones should be contacted through the SOCKS server.

If your system does not already have a SOCKS configuration file, a default SOCKS configuration file will be installed with WTE. This file is named socks.conf on AIX and Solaris, and socks.cnf on Windows NT.

9.2 Flexible-Client SOCKS Scenario

Flexible-client SOCKS can be configured using the Configuration and Administration forms or by editing the SOCKS configuration file. The following diagram shows a graphical representation of the flexible-client SOCKS scenario we implemented:

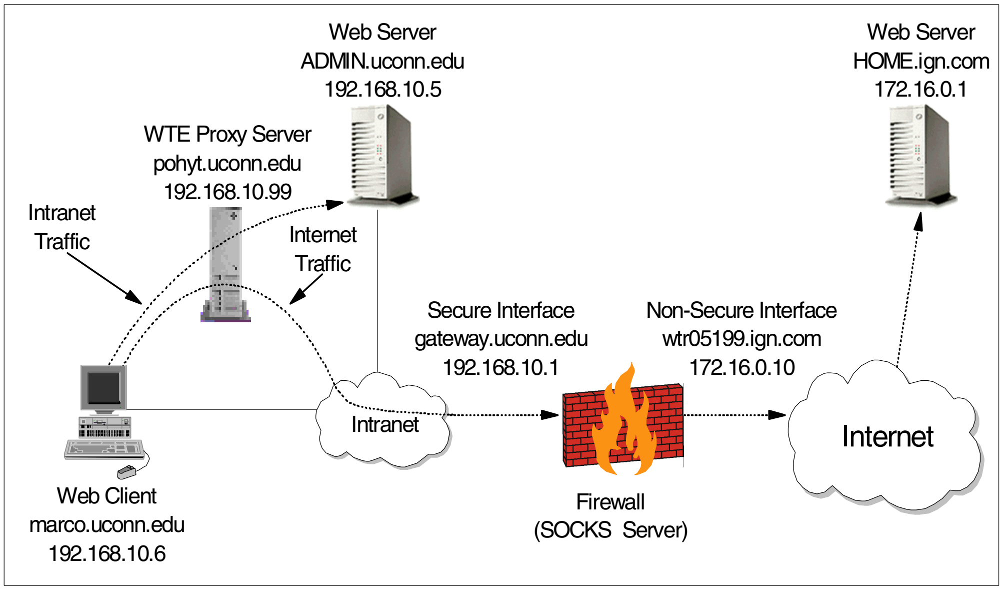

Figure 101. Schematic of the Flexible-Client SOCKS Implementation

Table 8 on page 157 is a summary of the components we used in this scenario:

Table 8. Flexible-Client SOCKS Scenario Components

Role	Operating System	Host Name	IP Address
Web client	Windows NT Server V4.0	marco.uconn.edu	192.168.10.6
WTE proxy server	Windows NT Server V4.0	pohyt.uconn.edu	192.168.10.99
Firewall (SOCKS server)	Windows NT Server V4.0	gateway.uconn.edu	192.168.10.1
		wtr05199.ign.com	172.16.0.1
Intranet Web Server	Windows NT Server V4.0	ADMIN.uconn.edu	192.168.10.5
Internet Web Server	Windows NT Server V4.0	HOME.ign.com	172.16.0.1

9.2.1 How to Set Up Flexible-Client SOCKS

The Configuration and Administration forms will be used in our implementation to set up the flexible-client SOCKS configuration in the WTE server. Follow the steps listed below:

1. Click the **Proxy Configuration** folder from the navigation menu.
2. Click the **SOCKS Configuration** folder.
3. Click **Direct Connections** for requests that should *not* pass through the SOCKS server. Click **SOCKS Connections** for requests that should pass through the SOCKS server. Take note that you can **Insert before**, **Insert after**, **Replace** or **Remove** the entries in the index by clicking the buttons and the Index numbers.

   In the Direct Connections form, fill in the IP Address, Subnet Mask and the Port Number fields for all the connections that are to be direct.

   For example, to have direct connections for all HTTP (port 80) traffic in the network 192.168.10, we entered IP Address `192.168.10.0`, and Subnet Mask

255.255.255.0, set Port to **Equal to** and entered Port Number 80 as shown in Figure 102.

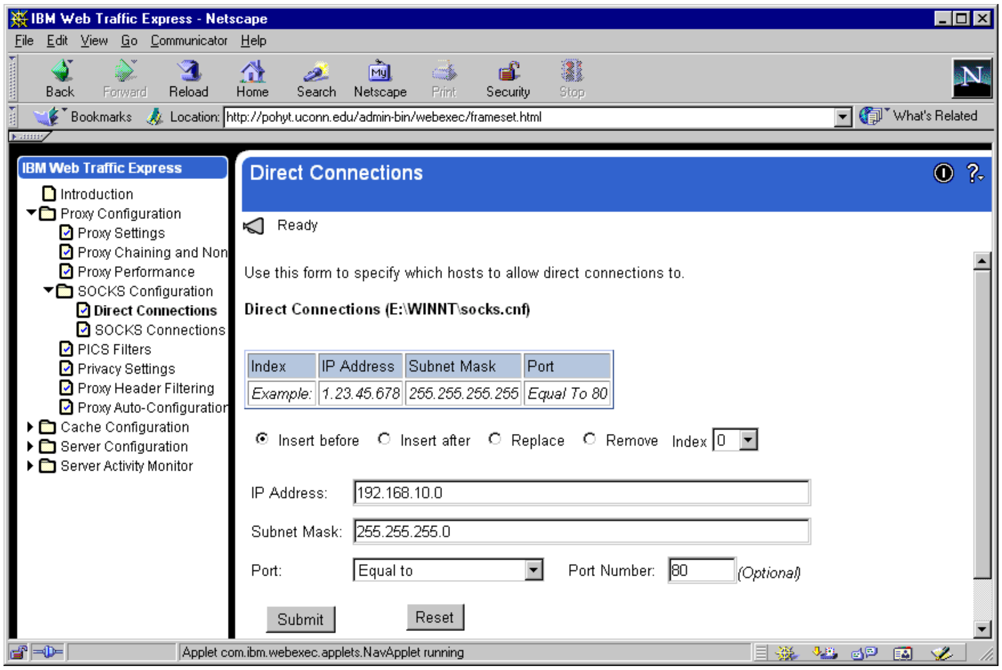

Direct Connections

Ready

Use this form to specify which hosts to allow direct connections to.

Direct Connections (E:\WINNT\socks.cnf)

Index	IP Address	Subnet Mask	Port
Example:	*1.23.45.678*	*255.255.255.255*	*Equal To 80*

● Insert before ○ Insert after ○ Replace ○ Remove Index 0

IP Address: 192.168.10.0

Subnet Mask: 255.255.255.0

Port: Equal to Port Number: 80 *(Optional)*

Figure 102. (Part 1 of 2). SOCKS Configuration - Direct Connections

Click **Submit** to update the SOCKS configuration file. The following figure shows a successful submission with the entries updated in the index table:

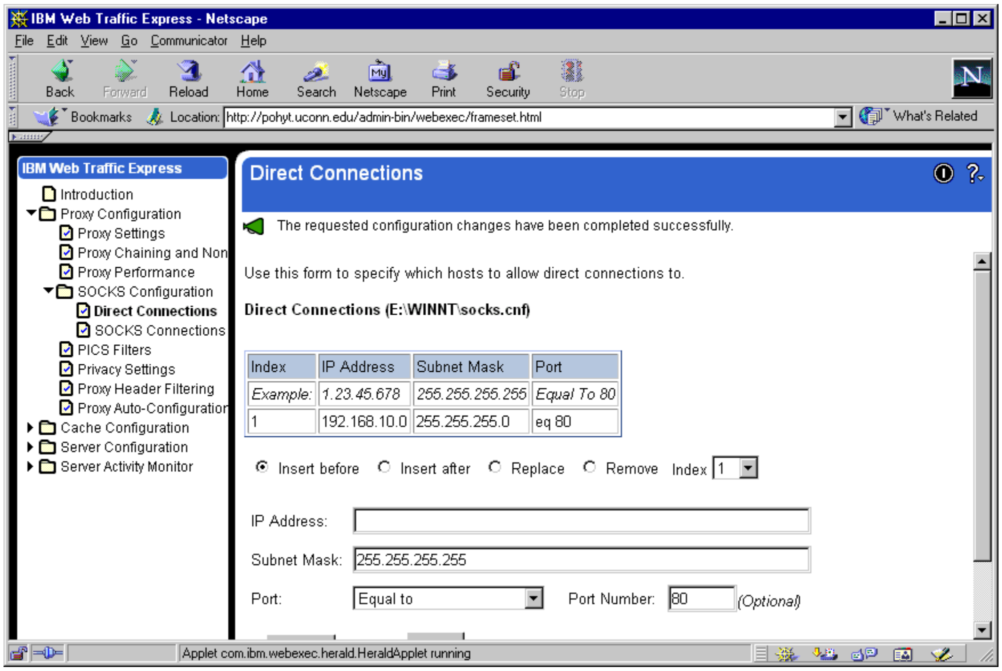

Direct Connections

The requested configuration changes have been completed successfully.

Use this form to specify which hosts to allow direct connections to.

Direct Connections (E:\WINNT\socks.cnf)

Index	IP Address	Subnet Mask	Port
Example:	*1.23.45.678*	*255.255.255.255*	*Equal To 80*
1	192.168.10.0	255.255.255.0	eq 80

● Insert before ○ Insert after ○ Replace ○ Remove Index 1

IP Address:

Subnet Mask: 255.255.255.255

Port: Equal to Port Number: 80 *(Optional)*

Figure 103. (Part 2 of 2). SOCKS Configuration - Direct Connections

In the SOCKS Connections form, fill in the IP Address, Subnet Mask and Port Number of the connections that are to use a SOCKS server.

For example, to have SOCKS connections for all HTTP (Port 80) traffic for all the other networks, we entered IP Address `0.0.0.0`, and Subnet Mask `0.0.0.0`, we set Port to **Equal to** and entered Port Number `80` as shown in Figure 104 on page 159.

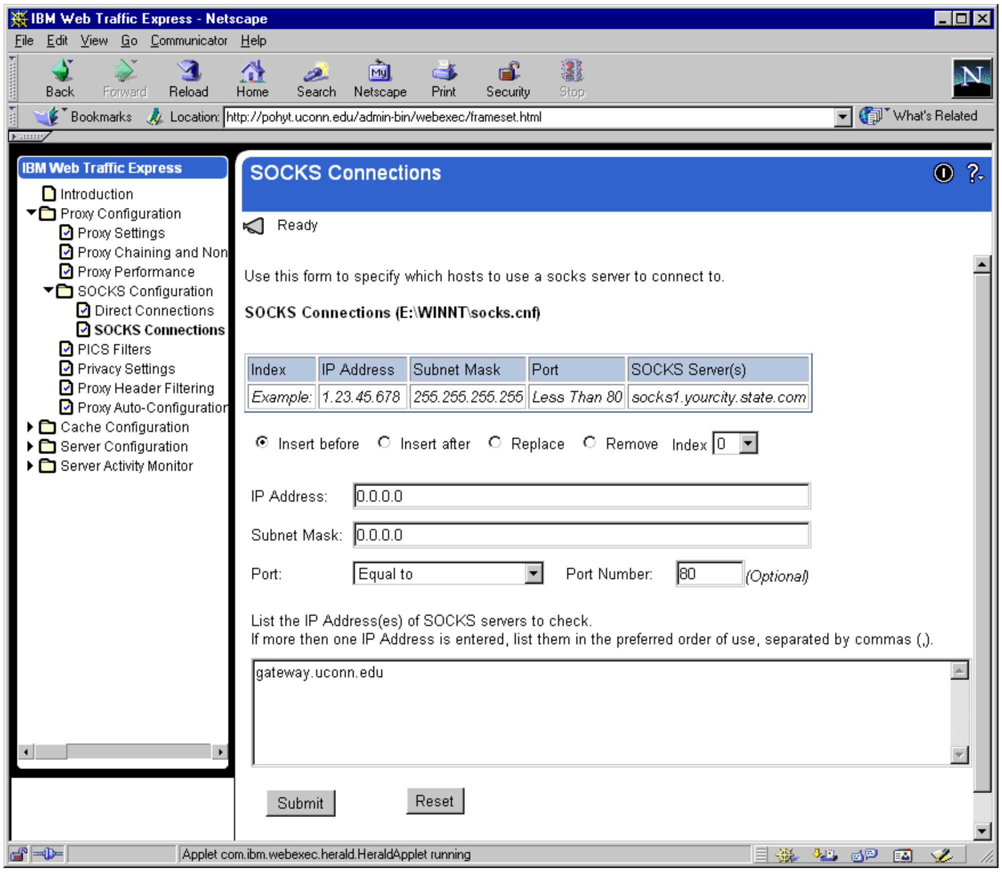

SOCKS Connections

Ready

Use this form to specify which hosts to use a socks server to connect to.

SOCKS Connections (E:\WINNT\socks.cnf)

Index	IP Address	Subnet Mask	Port	SOCKS Server(s)
Example:	*1.23.45.678*	*255.255.255.255*	*Less Than 80*	*socks1.yourcity.state.com*

● Insert before ○ Insert after ○ Replace ○ Remove Index 0

IP Address: 0.0.0.0

Subnet Mask: 0.0.0.0

Port: Equal to Port Number: 80 *(Optional)*

List the IP Address(es) of SOCKS servers to check.
If more then one IP Address is entered, list them in the preferred order of use, separated by commas (,).

gateway.uconn.edu

Submit Reset

Figure 104. (Part 1 of 2). SOCKS Configuration - SOCKS Connections

Click **Submit** to update the SOCKS configuration file. The following figure shows a successful submission with the entries updated in the index table:

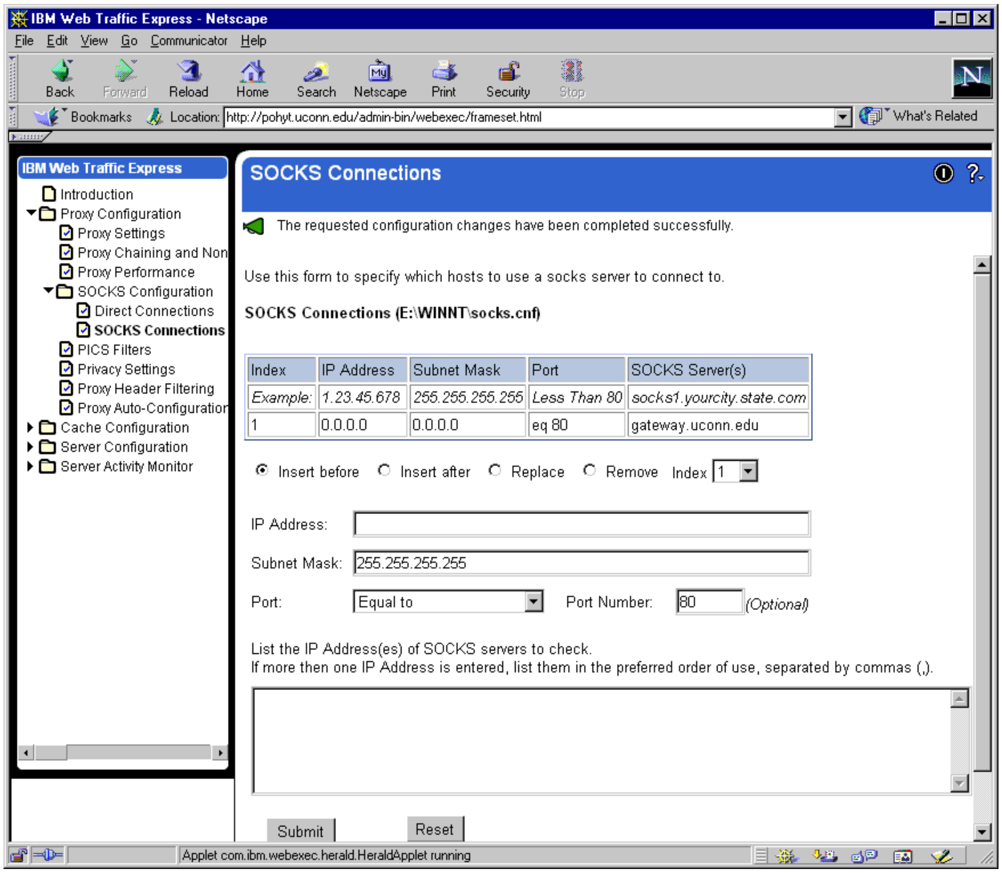

SOCKS Connections

The requested configuration changes have been completed successfully.

Use this form to specify which hosts to use a socks server to connect to.

SOCKS Connections (E:\WINNT\socks.cnf)

Index	IP Address	Subnet Mask	Port	SOCKS Server(s)
Example:	*1.23.45.678*	*255.255.255.255*	*Less Than 80*	*socks1.yourcity.state.com*
1	0.0.0.0	0.0.0.0	eq 80	gateway.uconn.edu

● Insert before ○ Insert after ○ Replace ○ Remove Index 1

IP Address:

Subnet Mask: 255.255.255.255

Port: Equal to Port Number: 80 *(Optional)*

List the IP Address(es) of SOCKS servers to check.
If more then one IP Address is entered, list them in the preferred order of use, separated by commas (,).

Submit Reset

Figure 105. SOCKS Configuration - SOCKS Connections 2

4. After successfully updating the SOCKS entries, we have to restart the WTE server to activate the changes. Click the **Restart** button to restart the WTE server.

9.2.2 Flexible-Client SOCKS Experience

To demonstrate how the flexible-client SOCKS configuration works, we set the Web client machine, marco, to use the WTE proxy server, pohyt. We also activated the Windows NT Network Monitor on the intranet Web server, ADMIN, so that we could verify that the traffic directed to this Web server was coming directly from the WTE proxy server instead of the SOCKS server.

On the Web client's browser, we entered the URL `http://ADMIN.uconn.edu` and it returned the ADMIN's home page. What we were interested in was the data captured by the Windows NT Network Monitor. The following figure shows a summary of the data captured:

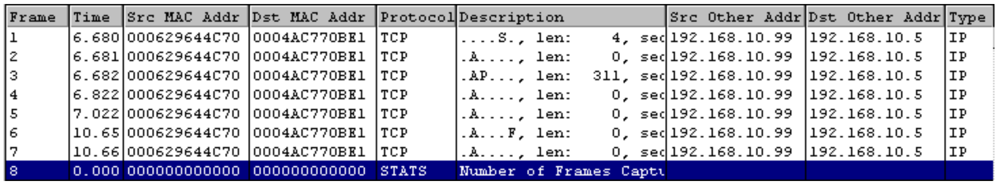

Frame	Time	Src MAC Addr	Dst MAC Addr	Protocol	Description	Src Other Addr	Dst Other Addr	Type
1	6.680	000629644C70	0004AC770BE1	TCP	....S., len: 4, sec	192.168.10.99	192.168.10.5	IP
2	6.681	000629644C70	0004AC770BE1	TCP	.A...., len: 0, sec	192.168.10.99	192.168.10.5	IP
3	6.682	000629644C70	0004AC770BE1	TCP	.AP..., len: 311, sec	192.168.10.99	192.168.10.5	IP
4	6.822	000629644C70	0004AC770BE1	TCP	.A...., len: 0, sec	192.168.10.99	192.168.10.5	IP
5	7.022	000629644C70	0004AC770BE1	TCP	.A...., len: 0, sec	192.168.10.99	192.168.10.5	IP
6	10.65	000629644C70	0004AC770BE1	TCP	.A...F, len: 0, sec	192.168.10.99	192.168.10.5	IP
7	10.66	000629644C70	0004AC770BE1	TCP	.A...., len: 0, sec	192.168.10.99	192.168.10.5	IP
8	0.000	000000000000	000000000000	STATS	Number of Frames Captu			

Figure 106. NT Network Monitor

From the captured data, the Src Other Addr (Source Other Address) and Dst Other Addr (Destination Other Address) columns show that all traffic came from the source IP address 192.168.10.99, which is the WTE proxy server, and went to the destination IP address 192.168.10.5, which is the intranet Web server. Hence, this verifies that the 192.168.10 network's traffic was not sent to the SOCKS server at all.

Similarly, we verified that all requests destined to the Internet Web server passed through the SOCKS server.

Using WTE as a traditional proxy server, without the flexible-client SOCKS configuration, we could verify that all the traffic was sent to the SOCKS server, regardless of its ultimate destination.

Chapter 10. WTE Proxy Chaining

Proxy chaining is a function that enables you to chain proxies together and to identify domains to which requests should be passed without going through a proxy. In this chapter, we examine the proxy chaining support offered by IBM Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2.

10.1 How WTE Proxy Chaining Works

A client sends its request first to the local WTE server. If the local WTE server does not find the file in its cache, it sends a proxy request to another proxy at a higher level. This WTE server, in turn, sends the request to a higher level proxy in the chain, and so forth. The last proxy in the chain will direct the request to the Internet and return the requested content back through the chain to the local WTE server, which returns it to the client.

All the proxies in the chain have the option of caching Web files. This way, when a file is requested, it may be unnecessary for a WTE proxy in the chain to forward the request to the next proxy in the chain or, for the highest proxy in the chain, to forward the request to the content server across the Internet. In fact, if the file is already in the local cache of a proxy, this proxy can serve it directly from the cache.

It is important to note that the more proxies you chain together, the greater the latency you will introduce. Chaining two proxies is the recommended limit. You may not need to chain proxies unless you are in a large organization; in that case, you will place small proxies close to the users (in a branch office for example) and then chain these small proxies to a large corporate proxy that connects directly to the Internet.

The following figure shows a graphical representation of the proxy chaining process:

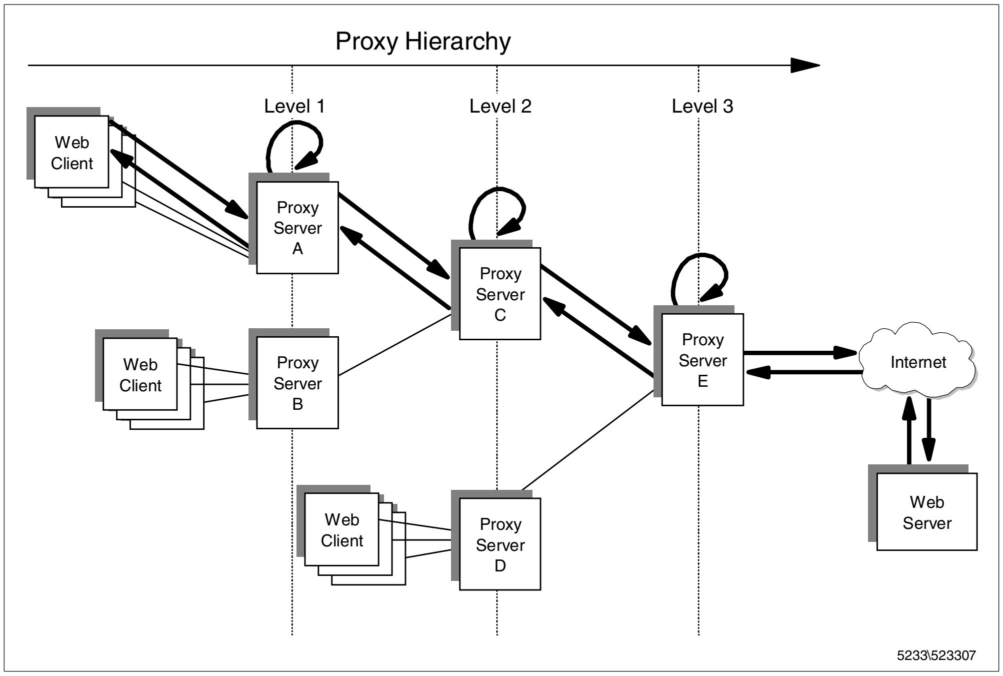

Figure 107. Graphical Representation of Proxy Chaining

Notice that the request goes through the daisy chain of servers and the response passes through the same proxies but in the reverse order. If intermediate proxies have caching enabled, each cache would be searched and would return a cached copy, if found, without forwarding the request to the origin Web server.

These are the advantages that you get when you use proxy chaining:

- Proxies at a lower level, which are closer to the client that originated the request, benefit from the caches of the higher level proxies.
- Proxy chaining reduces the load on the highest level proxy (typically, the proxy closest to the firewall) and ultimately on the content server, since lower level proxies may already have the document cached.
- The larger the number of users, the higher the probability that a proxy server already has a Web document in its cache. Considering that high-level proxies serve a larger number of clients, many requests that cannot be honored by a low-level proxy can be resolved by higher level proxies, since other groups of clients may have already requested the same files.

These are the disadvantages:

- A high-level proxy should have a larger cache, since it has to honor the requests of a large number of users.
- Proxy chaining greatly increases response time for requests, especially for noncacheable files or for those cacheable files that have not been cached yet by any proxies in the hierarchy.
- The risk of failure in a chain increases with each additional node.

WTE allows the creation of proxy chains based on the protocol (HTTP, FTP or Gopher) of the request to be forwarded. In other words, a proxy server can be configured to send all incoming requests with a particular protocol to a higher level proxy server in the chain. At the same time, you can specify the domains to which requests should be passed without going through a proxy.

10.2 Proxy Chaining Scenario

The following figure shows the architecture of the environment where we experimented with proxy chaining:

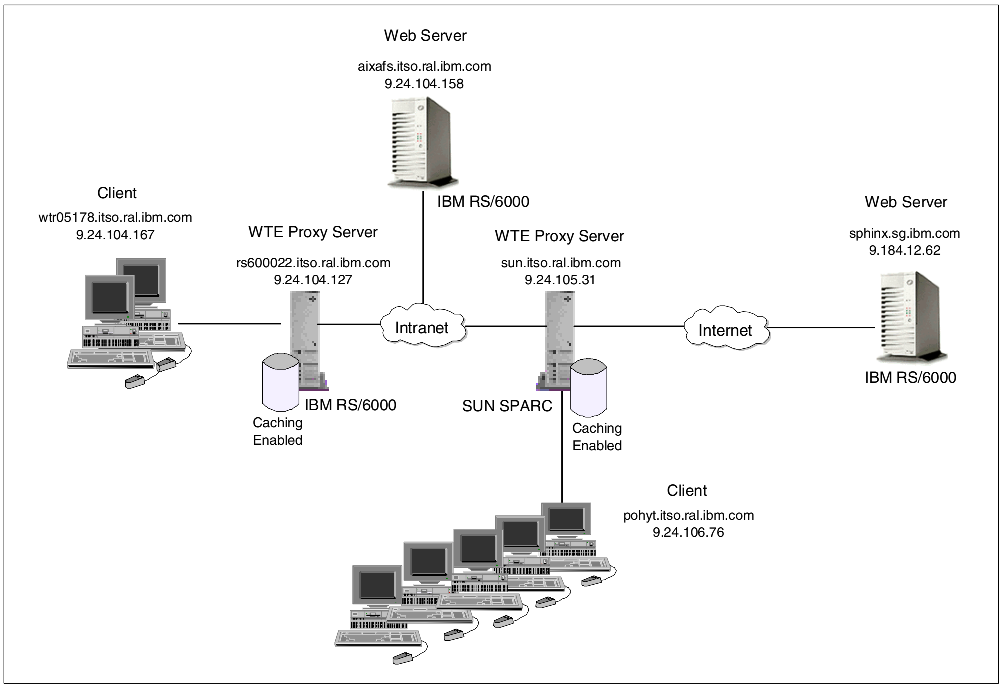

Figure 108. Architecture of the Proxy Chaining Environment

In this scenario, we will configure two WTE proxy servers in the chain. The first level proxy was installed on an IBM RS/6000 whose host name was rs600022 and whose IP address was 9.24.104.127. The second level proxy was installed on a SUN SPARC station whose host name was sun and whose IP address was 9.24.105.31. The first level proxy was configured to forward those requests that its cache could not satisfy to the second level proxy. Caching was enabled on both proxy servers.

On both machines, rs600022 and sun, we defined:

- The root directory for cached files as /wte/cache
- The cache access log file as /wte/logs/httpd-cache

As shown in the above figure, we had a Web browser installed on a client machine, whose host name was wtr05178 and whose IP address was 9.24.104.167. This Web browser, which was configured to use rs600022 as its proxy server, played the role of a first-level client, since it was configured to forward its requests to the first-level proxy installed on rs600022.

Another Web browser was installed on another client machine, pohyt. This Web browser played the role of the second-level client machine, since it was configured to forward its requests to the second-level proxy server, sun.

As you can see from the diagram, the first-level proxy in the architecture that we implemented was to serve a smaller group of clients, while the second-level proxy was to receive requests from a larger group of clients.

If a file requested by the first-level client was not found in the first-level proxy server's cache, the request would be directed to the second-level proxy server, sun. If the file was still not found on this proxy server's cache, the request would be forwarded to the Web content server, placed behind the Internet. The Web server machine was an RS/6000 located in Singapore, running Version 4.3 of the AIX operating system. The Web server function on this machine was obtained with Lotus Domino Go Webserver. The host name of this machine was sphinx.sg.ibm.com and its IP address was 9.184.12.62. On this Web server, the httpd daemon log file was /usr/lpp/internet/server_root/logs/httpd-log.Mar041999.

The following table summarizes the configuration of the environment where we performed our experiment:

Table 9. Environment Configuration for Proxy Chaining

Role	Operating System	Host Name	Domain	IP Address
First-level client	AIX 4.3.1	wtr05178	itso.ral.ibm.com	9.24.104.167
First-level proxy	AIX 4.3.1	rs600022	itso.ral.ibm.com	9.24.104.127
Second-level client	AIX 4.3.1	pohyt	itso.ral.ibm.com	9.24.106.76
Second-level proxy	AIX 4.3.1	sun	itso.ral.ibm.com	9.24.105.31
Web server	AIX 4.3	sphinx	sg.ibm.com	9.184.12.62

10.2.1 How to Set Up Proxy Chaining with WTE

To configure proxy chaining, on the first-level proxy, rs600022, we clicked the **Proxy Configuration** folder on the navigation menu, and selected the **Proxy Chaining and Non-Proxy Domains** form by clicking it. The following figure shows the Proxy Chaining and Non-Proxy Domains form:

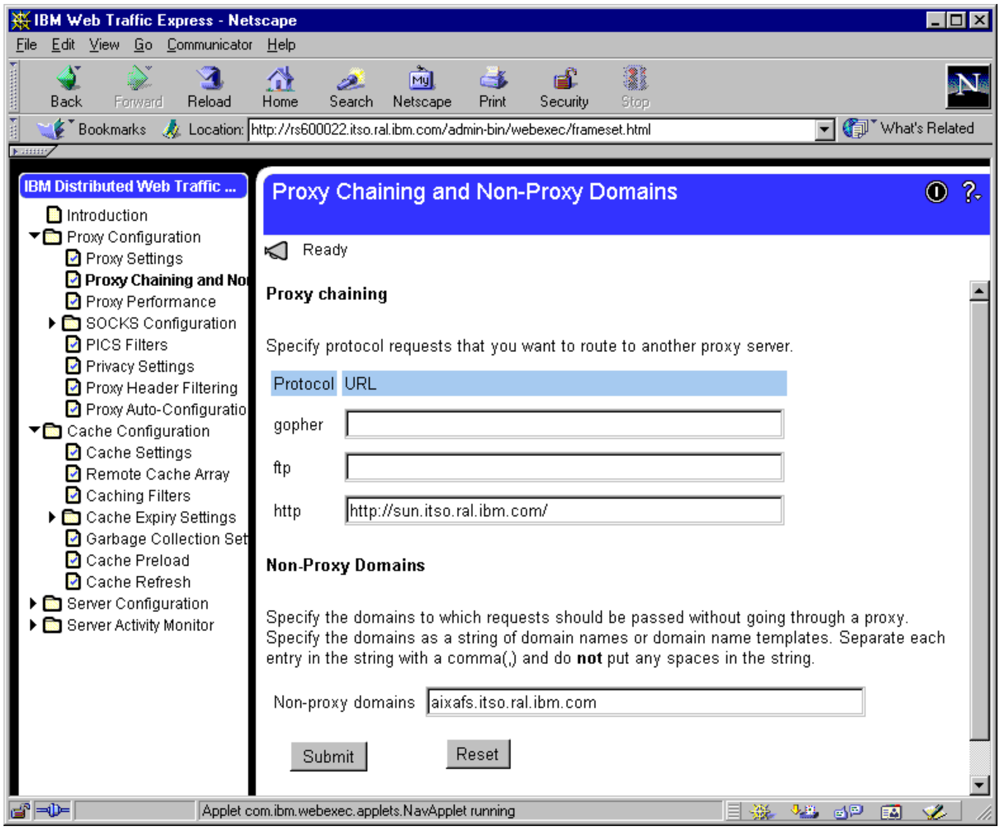

Proxy Chaining and Non-Proxy Domains

Ready

Proxy chaining

Specify protocol requests that you want to route to another proxy server.

Protocol	URL
gopher	
ftp	
http	http://sun.itso.ral.ibm.com/

Non-Proxy Domains

Specify the domains to which requests should be passed without going through a proxy. Specify the domains as a string of domain names or domain name templates. Separate each entry in the string with a comma(,) and do **not** put any spaces in the string.

Non-proxy domains: aixafs.itso.ral.ibm.com

Submit Reset

Figure 109. Proxy Chaining and Non-Proxy Domains Form

WTE supports proxy chaining for FTP, Gopher and HTTP. However, in this section, we only wanted this proxy to forward all HTTP requests that it could not satisfy from its own cache to the second-level proxy, sun. We entered in the http text field the URL `http://sun.itso.ral.ibm.com/`. Notice that the trailing slash must be included in the URL.

We also want to show you an example of non-proxy domain configuration. We entered `aixafs.itso.ral.ibm.com` in the Non-Proxy Domains text field. This was the Web server which we would not need to go through a proxy to access.

Using the Directives

To specify sun as the higher level proxy in the chain, we could also have entered the following directive in the WTE configuration file, ibmproxy.conf:

```
Http_Proxy http://sun.itso.ral.ibm.com/
```

The non-proxy domain directive would have been the following:

```
no_proxy aixafs.itso.ral.ibm.com
```

We clicked the **Submit** button to update the first-level proxy server rs600022 with the changes. To have the changes take effect, we had to restart the WTE server.

Notice that proxy chaining is not enabled by default and *you have to manually configure all the proxy servers that need to recognize a higher level proxy server.*

The proxy server in the highest level of the chain does not have to recognize a higher level proxy and can directly contact the Web servers in the Internet. For this reason, it is the only proxy server in the chain that does not need to be configured for proxy chaining.

10.2.2 Proxy Chaining Experiences

After we performed the steps described in 10.2.1, “How to Set Up Proxy Chaining with WTE” on page 166, we wanted to verify that our proxy chaining implementation was working correctly.

First of all, we stopped the proxy servers on both machines. On the IBM RS/6000 proxy and Web server, we used the command:

```
stopsrc -s httpd
```

On the SUN SPARC server, we used the command:

```
kill -9 HTTP_PID
```

where *HTTP_PID* was the process ID of the WTE server (see 3.5.2, “Starting and Stopping WTE on Solaris” on page 71).

To better follow the test, we removed all the previous logs and caches from all the machines by removing all files and subdirectories under the directories /wte for the WTE and /usr/lpp/internet/server_root/logs for the Web server.

After doing this, we restarted both proxy servers. On the RS/6000 machine, we entered:

```
startsrc -s httpd
```

On the SUN SPARC station, we entered the `httpd` command (see 3.5, “Starting and Stopping WTE” on page 70).

As we wanted to follow the online growth of the log files in consequence of our actions, we used the `tail` command as follows:

1. On rs600022, we entered the following command:

   ```
   tail -f /wte/logs/httpd-cache.Mar051999
   ```

2. On sun, we entered the following command:

   ```
   tail -f /wte/logs/httpd-cache.Mar051999
   ```

3. On the sphinx machine in Singapore (notice the day difference), we entered the following command:

   ```
   tail -f /usr/lpp/internet/server_root/logs/httpd-log.Mar061999
   ```

10.2.2.1 First Access

From the Web browser installed on the second-level client machine, pohyt, on which the proxy was set to the WTE proxy server sun, we request the URL `http://sphinx.sg.ibm.com/sphinx.html`.

What happens to our request?

1. The request from the Web browser is taken by the proxy server sun.
2. The proxy server sun looks at its own cache and does not find the file. No cache access is done on the proxy server sun as shown:

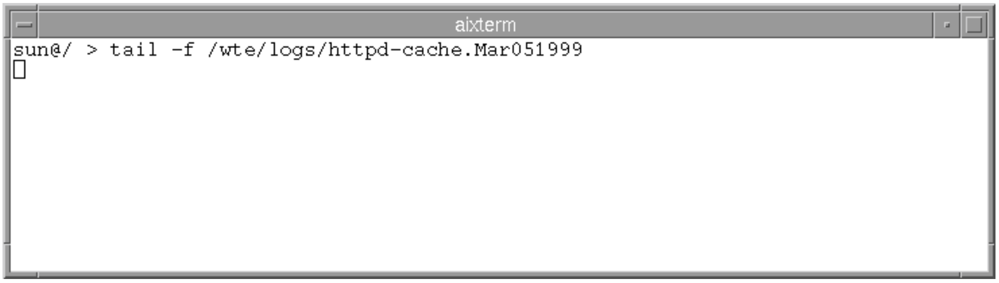

```
sun@/ > tail -f /wte/logs/httpd-cache.Mar051999
```

Figure 110. Cache Access Log File on Proxy Server sun

3. The proxy server sun requests the file from the Web server sphinx.
4. The Web server sphinx registers the access and returns the page sphinx.html to the proxy server sun, whose IP address is 9.24.105.31, as shown:

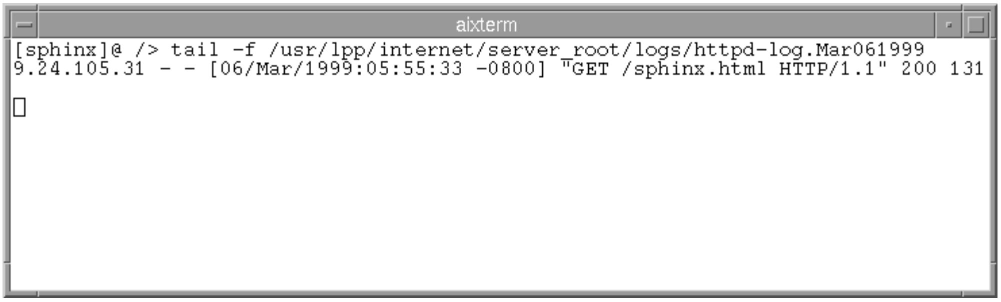

```
[sphinx]@ /> tail -f /usr/lpp/internet/server_root/logs/httpd-log.Mar061999
9.24.105.31 - - [06/Mar/1999:05:55:33 -0800] "GET /sphinx.html HTTP/1.1" 200 131
```

Figure 111. Httpd Log File on Web Server sphinx

5. The proxy server sun caches the page sphinx.html and also sends it to the Web browser.

There was no access to the rs600022 proxy server cache, as shown in the following figure:

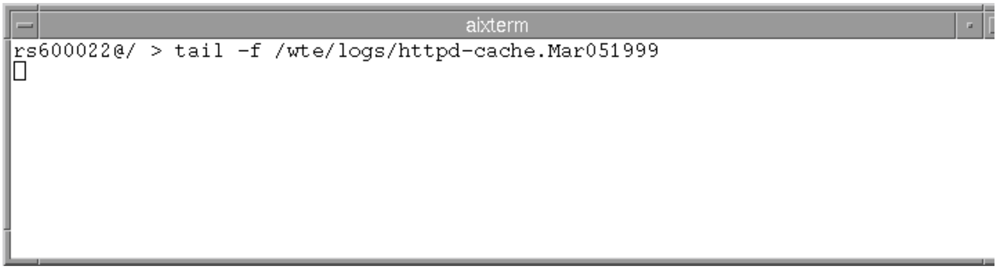

```
rs600022@/ > tail -f /wte/logs/httpd-cache.Mar051999
```

Figure 112. Cache Access Log File on Proxy Server rs600022

This means that, at this point, the file sphinx.html has been cached by sun, but not by rs600022.

10.2.2.2 Second Access

From the Web browser on the first-level client machine wtr05178, which has its proxy server set to rs600022, we request the URL `http://sphinx.sg.ibm.com/sphinx.html`.

What happens to our request this time?

1. The request from the Web browser is taken by the proxy server rs600022.
2. The proxy server rs600022 looks up its own cache and does not find the file sphinx.html. There is no access to the rs600022 proxy server cache as shown:

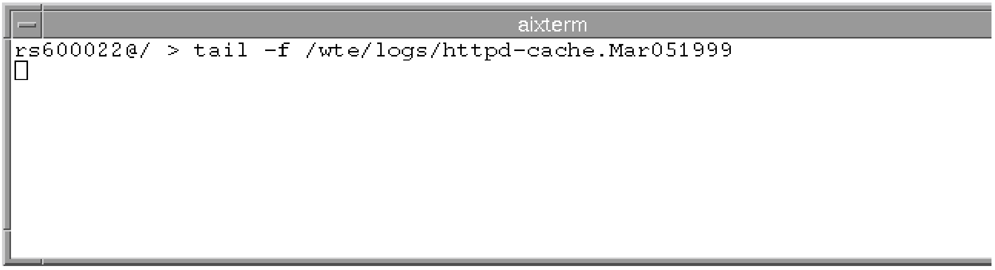

```
rs600022@/ > tail -f /wte/logs/httpd-cache.Mar051999
```

Figure 113. Cache Access Log File on rs600022

3. The proxy server rs600022 requests the file from the next higher level proxy server in the chain, which is the proxy server sun. In turn, the proxy server sun looks up its own cache, where it finds the file sphinx.html. It then sends this file back to the proxy server rs600022. In this scenario, the proxy server rs600022, having IP address 9.24.104.127, has accessed the cache of proxy server sun, as captured by the access log shown below:

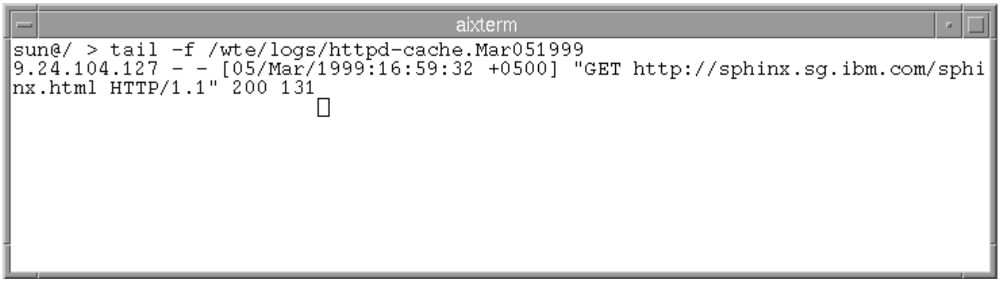

```
sun@/ > tail -f /wte/logs/httpd-cache.Mar051999
9.24.104.127 - - [05/Mar/1999:16:59:32 +0500] "GET http://sphinx.sg.ibm.com/sphi
nx.html HTTP/1.1" 200 131
```

Figure 114. Cache Access Log File on Proxy Server sun

4. The proxy server rs600022 caches the file sphinx.html, and sends it back to the Web browser, wtr05178.

Notice that there was no access to the Web server sphinx as shown in the following figure:

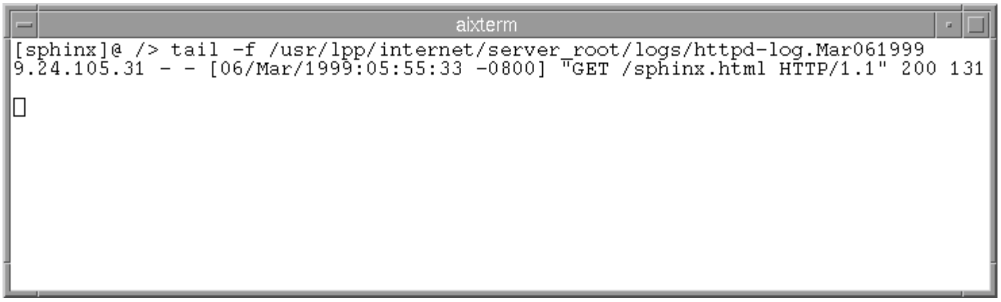

```
[sphinx]@ /> tail -f /usr/lpp/internet/server_root/logs/httpd-log.Mar061999
9.24.105.31 - - [06/Mar/1999:05:55:33 -0800] "GET /sphinx.html HTTP/1.1" 200 131
```

Figure 115. Httpd Log File on Web Server sphinx

Note that, at this point the file sphinx.html has been cached by both the proxy servers, sun and rs600022.

10.2.2.3 Third Access

From the Web browser, on the first-level client wtr05178, which has its proxy set to rs600022, we request again the URL `http://rs600030/afs/webstone/html/sphinx.html`.

What happens to our request this time?

1. The request coming from the Web browser was taken by the proxy server rs600022.
2. The proxy server rs600022 looks up its own cache, finds the file sphinx.html and sends it back to the Web browser. In this scenario, the first-level client machine, having IP address 9.24.104.167, has accessed the cache of rs600022, as captured in the following figure:

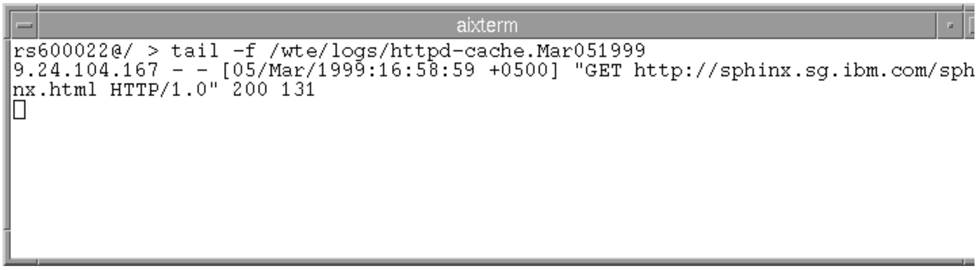

```
rs600022@/ > tail -f /wte/logs/httpd-cache.Mar051999
9.24.104.167 - - [05/Mar/1999:16:58:59 +0500] "GET http://sphinx.sg.ibm.com/sph
nx.html HTTP/1.0" 200 131
```

Figure 116. Cache Access Log File of Proxy Server, rs600022

No accesses were registered on either the second-level proxy, sun, or the destination Web server, sphinx.

10.2.2.4 Fourth Access

Now, from the Web browser on the first-level client machine, wtr05178, which has its proxy set to rs600022, we requested another URL, `http://sphinx.sg.ibm.com/sphinx11.html`. This URL has never been requested before.

What happens to our request this time?

1. The request coming from the Web browser is taken by the proxy server rs600022.
2. The proxy server rs600022 looks up its own cache but does not find the file. Hence the rs600022 proxy server cache is not accessed.
3. The proxy server rs600022 requests the file from the second-level proxy server in the chain, sun.
4. The proxy server sun looks up its own cache. File sphinx11.html is not found there either. Hence, the proxy server sun requests the file from the destination Web server, sphinx.
5. The Web server sphinx registers the access and returns the page sphinx11.html back to the proxy server sun. In this scenario, the proxy server sun, having IP address 9.24.105.31, has accessed the Web server sphinx, as captured in the following figure:

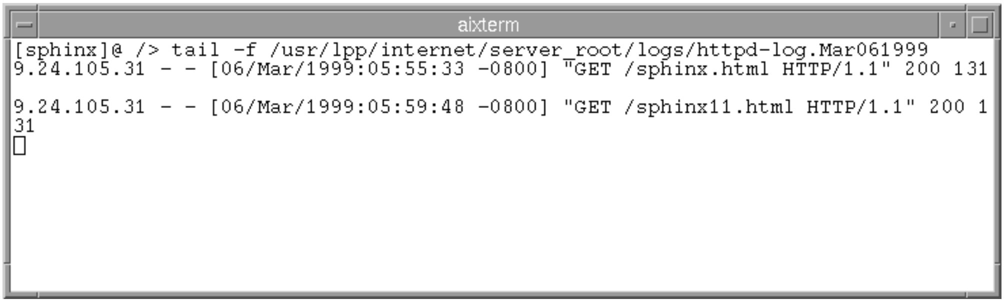

```
[sphinx]@ /> tail -f /usr/lpp/internet/server_root/logs/httpd-log.Mar061999
9.24.105.31 - - [06/Mar/1999:05:55:33 -0800] "GET /sphinx.html HTTP/1.1" 200 131

9.24.105.31 - - [06/Mar/1999:05:59:48 -0800] "GET /sphinx11.html HTTP/1.1" 200 1
31
```

Figure 117. Httpd Log File on the Web Server, sphinx

6. The proxy server sun caches this page and sends it to the first-level proxy server rs600022.
7. The proxy server rs600022 also caches this file and sends it back to the Web browser, wtr05178, that made the request for this page.

10.2.2.5 Fifth Access

From the Web browser installed on the first-level client wtr05178, we request again the URL `http://sphinx.sg.ibm.com/sphinx11.html`.

What happens to our request this time?

1. The request coming from the Web browser is taken by the proxy server rs600022.
2. The proxy server rs600022 looks up its own cache. This time, it finds the file sphinx11.html and sends it back directly to the Web browser. In this scenario, the Web browser on wtr05178 has accessed the cache of rs600022, as captured in the following figure:

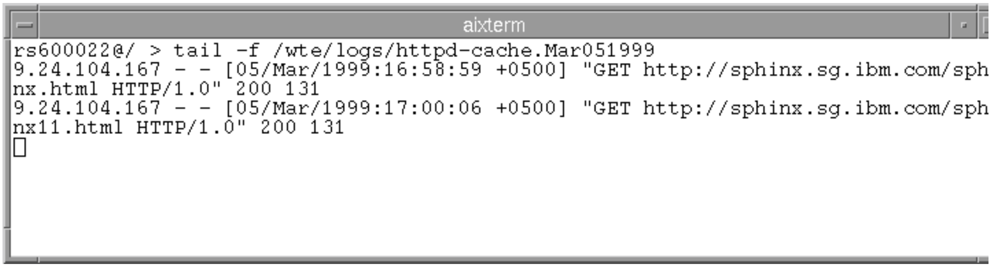

```
rs600022@/ > tail -f /wte/logs/httpd-cache.Mar051999
9.24.104.167 - - [05/Mar/1999:16:58:59 +0500] "GET http://sphinx.sg.ibm.com/sph
nx.html HTTP/1.0" 200 131
9.24.104.167 - - [05/Mar/1999:17:00:06 +0500] "GET http://sphinx.sg.ibm.com/sph
nx11.html HTTP/1.0" 200 131
```

Figure 118. Cache Access Log File on Proxy Server, rs600022

No accesses to the cache on the WTE proxy server sun or the Web server sphinx were registered.

10.2.2.6 Accessing a Machine in the Non-Proxy Domain

To conclude our scenario, we wanted to experiment with proxy chaining and non-proxy domains. Recall that in the Proxy Chaining and Non-Proxy Domains configuration form (see Figure 109 on page 167) we checked the box, we entered the host name of the aixafs Web server machine in the Non proxy domains text box, and that in “Using the Directives” on page 167, we said that this would have been equivalent to entering the following directive in the WTE configuration file:

```
no_proxy aixafs.itso.ral.ibm.com
```

This directive can be used to specify the domains to which the WTE server should directly connect. Notice that this is only used when doing proxy chaining. This directive does not apply when the proxy goes through a SOCKS server; use the socks.conf configuration file for that purpose (see Chapter 9, “WTE Flexible-Client SOCKS Support” on page 155). Also, if you are using neither proxy chaining nor SOCKS, then this directive is not needed.

To demonstrate how the directive above modified the existing configuration, from the Web browser installed on our client machine, we invoked the home page of the aixafs Web server, by pointing the browser to the URL `http://aixafs.itso.ral.ibm.com`.

Following a procedure similar to what we did in 10.2.2.1, “First Access” on page 168 through 10.2.2.4, “Fourth Access” on page 171, we could verify that the request was processed by the Web server rs600022. This Web server did not have the requested page in the cache yet, but this time, instead of forwarding the request to the higher level proxy, sun, it contacted the Web server aixafs directly. The Web server served the response back to the proxy server rs600022, and this, in turn, passed it to the client.

This demonstrates that, in a proxy chaining environment, the `no_proxy` directive allows a low-level proxy to forward requests for a specific Web server to that Web server directly instead of using a higher-level proxy.

Chapter 11. WTE Proxy Server Protection Scenario

This chapter describes how to control access to an IBM Web Traffic Express (WTE) server. WTE is the caching and filtering component of IBM WebSphere Performance Pack.

There are several ways you can protect your server:

- Physically locate the server meant for public access in the external network.
- Disable telnet, rlogin, and finger clients on the system that is running the server.
- Use packet filtering and firewalls.
- Protect access to CGI scripts.

In addition, WTE allows you to control who can access the files on the server and can be configured to accept only requests generated by users belonging to specified domains. When protection is enabled, only authorized users can access the WTE server. A user is required to be authenticated with a user ID and password.

A Note on Transparent Proxy

Do not use protection directives to require user ID and password for access to a transparent proxy (see Chapter 7, “Using WTE as a Transparent Proxy Server” on page 131). In fact, in a transparent proxy scenario, browsers are not aware of the transparent proxy and are not able to handle user ID and password requests from the proxy.

You can set up proxy server protection by editing the ibmproxy.conf configuration file or by using the Configuration and Administration Forms.

In this chapter, we describe our experience with WTE protection. The machine where we installed the WTE was an RS/6000, with AIX 4.3.1 as the operating system. The machine host name was rs600022.itso.ral.ibm.com and its IP address was 9.24.104.127.

11.1 Resource Protection

There are three directives in the ibmproxy.conf configuration file that define the file access protection on the WTE server. The directives are:

- `Protect`
- `DefProt`
- `Protection`

Note

If you are editing the ibmproxy.conf configuration file manually, you must ensure that the `Pass` (to accept document requests) and the `Exec` (to accept CGI program requests) directives accept the requests that the `Protect` directives are activated for.

The `DefProt` and `Protect` directives must be put before any `Pass`, `Exec`, or `Proxy` directives in the WTE configuration file.

When a Protect directive activates protection for a request, it also either identifies the protection setup to use or defines the protection setup as part of the directive. A *protection setup* is a group of protection subdirectives. The subdirectives work together to define how the server should control access to the resources being protected. You can create protection setups in three different ways:

1. You can create protection setups within the configuration file as part of `Protection`, `Protect`, or `DefProt` directives.

   - When you create a protection setup on a `Protection` directive, you give the setup a label that you can point to later from `Protect` and `DefProt` directives. This type of protection setup is called a *named* protection setup.
   - When you create a protection setup on a `DefProt` or `Protect` directive, the protection setup is used only for that directive. The setup cannot be pointed to by other `DefProt` or `Protect` directives. This type of protection setup is called an *in-line* protection setup.

   To include a protection setup as part of a `Protection`, `Protect`, or `DefProt` directive, insert a left brace character ({) as the last character on the line that contains the directive. On each subsequent line, put one Protection subdirective and its value. Indicate the end of the protection setup by putting a right brace character (}) by itself on the line following the last protection subdirective. You cannot use comments within the protection setup.

2. You can create separate protection setup files that you can then point to from `Protect` and `DefProt` directives. A protection setup file is a plain text file. Within the file, each line contains one protection subdirective and the value for that subdirective.

3. You can use the configuration and administration forms to create protection setups within the configuration file. From the Configuration And Administration forms page, click **Server Configuration**, then **Document Protection** and add, change, or delete either an in-line protection setup or a named protection setup.

Within the protection setup, the protection subdirectives control access to the directory or files being protected.

Notice that for the `DefProt` directive, protection is not actually activated for requests matching the template unless the requests also match a template on a subsequent `Protect` directive. We will use the following example to explain the `DefProt` directive briefly:

```
DefProt /secret/* SECRET-PROT
Protect /secret/salesinfo/*
```

The DefProt directive defines what the default protection will be for all the files and directories starting with the template /secret/. The label name SECRET-PROT must match a label name on a Protection directive that needs to be defined before the DefPort directive.

The Protect directive used here demonstrates the use of DefProt. No access file or protection setup names were defined here.

To access a file in /server_root/secret/ no password is required, since DefProt only defines the protection setup but does not apply it. However, to access any files in /secret/salesinfo/ you have to be an authenticated user before being able to access.

11.2 Directives to Protect Access to the Proxy

To define a new protection setup, we can use the configuration and administration forms or simply edit the WTE proxy configuration file, ibmproxy.conf. We use the latter method, since in this case we considered this procedure more useful to understand how the underlying directives work.

To define a new protection setup, PROT-PROXY, we add the following lines after the PROT-ADMIN protection setup in the ibmproxy.conf file to restrict access to the proxy server from users accessing a Web site with the http: method:

```
Protection PROT-PROXY {
      ServerId        ProxyServer
      AuthType        Basic
      GetMask         All@(*)
      PutMask         All@(*)
      PostMask        All@(*)
      Mask            All@(*)
      PasswdFile      /etc/wteusers.passwd
                          }

Protect http:* PROT-PROXY
```

The protection setup is as follows:

- The Protection directive defines a protection setup, named PROT-PROXY.
- The ServerID subdirective specifies the name that identifies the protection setup to requesters. The name does not need to be a real machine name. When the proxy server sends a requester a prompt for a user ID and password, it will also include the name you specify as the value of the ServerID subdirective. Most browsers display this name on the prompt. In this way, the requester is capable of understanding which proxy server sent that prompt, and can decide the user ID and password to send back.
- The AuthType subdirective specifies the type of authentication. We used the value Basic to specify basic authentication, meaning that users are

authenticated through user ID and password. In basic authentication, user ID and password are usually encoded in base64 format, but they are not encrypted (for further details on this, see the IBM redbook *Internet Security in the Network Computing Framework*, SG24-5220.)

- The `PasswdFile` subdirective specifies the path and name of the password file to be used. In our case, we defined this subdirective as `/etc/wteusers.passwd`. A password file contains a list of user IDs and passwords. Each user ID has one valid password defined for it. Note that the user ID in the password file does not have any relation to the addresses or host name machines of the requester, nor with the users defined on the underlying operating system.

  A password file is created in the proxy server machine itself. To create and maintain password files, you can access the configuration and administration forms or use the `htadm` command. We chose to use the `htadm` command, as we show in 11.3.1, “Using the htadm Command” on page 179.

- The `Mask` subdirectives allow you to specify how one or more of the HTTP methods are to be accessed. For example, you can specify different access levels for GET and POST methods using the `GetMask` and `PostMask` subdirectives, or prohibit other methods such as PUT or DELETE from being accessed.

  By modifying the values of the various `Mask` subdirectives, you can specify that access is via user ID and password, IP address, or a combination of both.

  For example:

  - The subdirective

    ```
    GetMask All
    ```

    would specify that any user in the password file specified by the `PasswdFile` subdirective could issue a GET request. The user will then be prompted to enter a valid user ID and password as defined in the password file.

  - The subdirective:

    ```
    GetMask All@96.*.*.*
    ```

    would permit access to all users in the password file, but only if the request came in from an IP address starting with `96`. Any other requests would be rejected.

  Note that `GetMask All@(*)` is equivalent to `GetMask All`.

The `Mask` subdirective allows you to authorize any other enabled method not covered by the other `Mask` subdirectives, such as `GetMask`, `PostMask` and `PutMask`. Notice that these other subdirectives take precedence over the `Mask` subdirective if they are all present in the protection setup. The consequence of these considerations is that we could substitute the protection setup we have just shown with the following:

```
Protection PROT-PROXY {
    ServerId        ProxyServer
    AuthType        Basic
    Mask            All@(*)
    PasswdFile      /etc/wteusers.passwd
}

Protect http:* PROT-PROXY
```

As you can see, the `Protection` directive and its subdirectives are very flexible and will allow you to define your resource access policies in several different ways or combinations.

All you need to do after defining the protection setup is to use it in a `Protect` directive to associate it with your proxy resources. We used the following:

```
Protect http:* PROT-PROXY
```

So, in our case, all client requests starting with `http:` will cause the WTE server to prompt for a user ID and password.

Note that in this situation, client requests that ask to use protocols other than HTTP, such as FTP, would not be protected by the WTE server, since they would start with `ftp:`. To protect this type of request, you should add the following directive:

```
Protect ftp:* PROT-PROXY
```

If you want protect the access to all the WTE server's functions using the protection strategy `PROT-PROXY`, then the only directive to use would be:

```
Protect * PROT-PROXY
```

11.3 Managing Password Files

In WTE, password files can be managed through the command line, using the `htadm` command, or through a Web browser, using the Configuration and Administration forms. In this section, we will show you how to use both these options.

The experience we describe in this section was performed on AIX, but it would be very similar on Solaris and Windows NT. The only difference in the setup between UNIX and Windows NT systems consists in the paths of the directories where the files involved are located.

11.3.1 Using the htadm Command

The `htadm` command allows you to manage the WTE caching and filtering proxy server password files. This command is also used to manage password files in Lotus Domino Go Webserver. Using the `htadm` command, you can add or delete a user ID and password, check a user's password, and create an empty password file. *Passwords are stored encrypted*.

Notice that, on AIX, the htadm executable file is automatically installed in the /usr/sbin directory, which is included in the PATH system environment variable.

On Windows NT, an executable file, named HTAdm.exe, is stored in the Bin directory under the caching and filtering proxy server directory WWW. The Bin directory contains other executable files, and is automatically included in the Path system environment variable during the installation of WTE.

An empty password file must first be created. We created an empty file named wteusers.passwd, in the directory /usr/lpp/internet/server_root/protect of the WTE server. To do this, we entered the following two commands:

```
cd /usr/lpp/internet/server_root/protect
htadm -create wteusers.passwd
```

Then we added to the wteusers.passwd password file a user ID pohyt with password `pohyt` and the identification string:

```
User1: Poh Yee Tiong
```

as a comment. We did this through the following command:

```
htadm -adduser wteusers.passwd pohyt pohyt "User1: Poh Yee Tiong"
```

In a similar way, we added another user ID Marco with password `marco` and the identification string:

```
User2: Marco Pistoia
```

as a comment. We did this through the following command:

```
htadm -adduser wteusers.passwd marco marco "User2: Marco Pistoia"
```

11.3.2 Using the Configuration and Administration Forms

The creation of a password file and users also can be accomplished through the Configuration and Administration forms. In the navigation frame, click the following folders consecutively:

- **Server Configuration**
- **System Management**
- **User Management**

Finally click the **Add User** form.

Note that the Add User form will only be accepted if the Group information is entered. The Group File entry is the PATH and the name of the server group file to which the user belongs. It is used in the protection setup.

The following is an example of the Add User form:

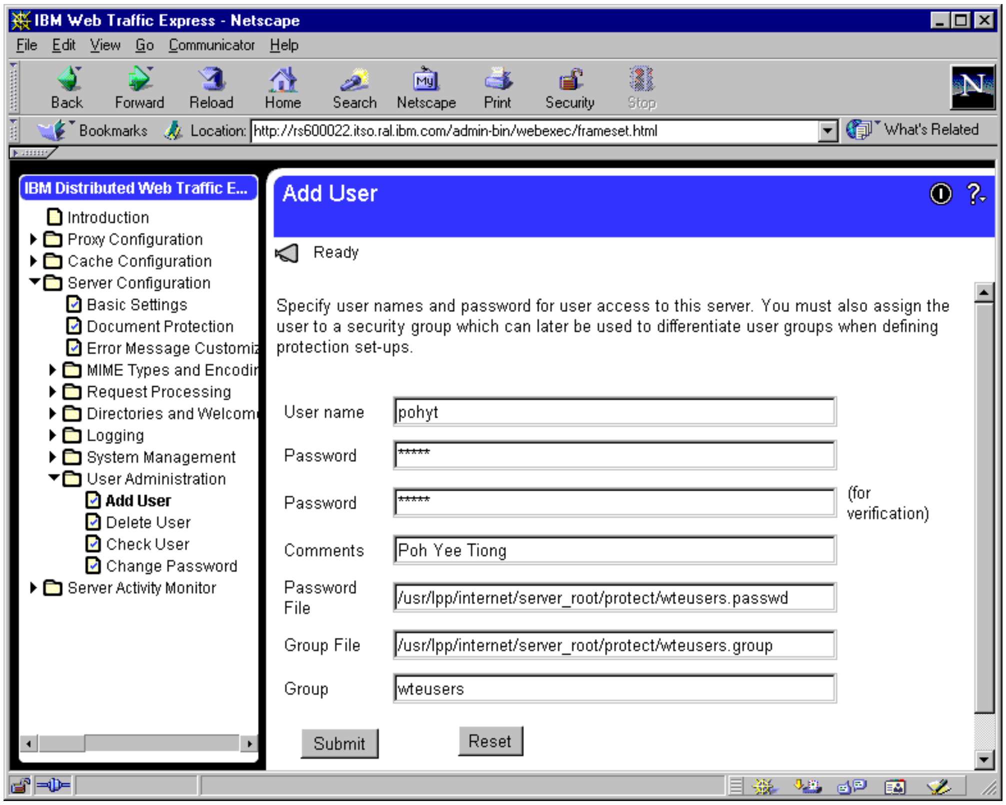

Add User

Specify user names and password for user access to this server. You must also assign the user to a security group which can later be used to differentiate user groups when defining protection set-ups.

User name	pohyt
Password	*****
Password	***** (for verification)
Comments	Poh Yee Tiong
Password File	/usr/lpp/internet/server_root/protect/wteusers.passwd
Group File	/usr/lpp/internet/server_root/protect/wteusers.group
Group	wteusers

Figure 119. IBM WTE GUI - Add User Form

The Configuration and Administration forms also allow us to delete, check and change user passwords.

11.4 Testing the Configuration

To verify the proxy access protection, on a Web client AIX machine, with IP address 9.24.104.167 and host name wtr05178, we configured a Web browser to use the WTE proxy server rs600022.

From the Web browser we requested the home page from the Web server home.netscape.com by entering `http://home.netscape.com` in the Location field of the browser. With the protection setup, we got a prompt from the proxy server rs600022 requesting a user ID and password. Unless we authenticated correctly with basic authentication we could not access the requested URL through the defined proxy server.

When we entered an incorrect password and clicked **OK**, we got the following error authentication question displayed by our Netscape Navigator browser:

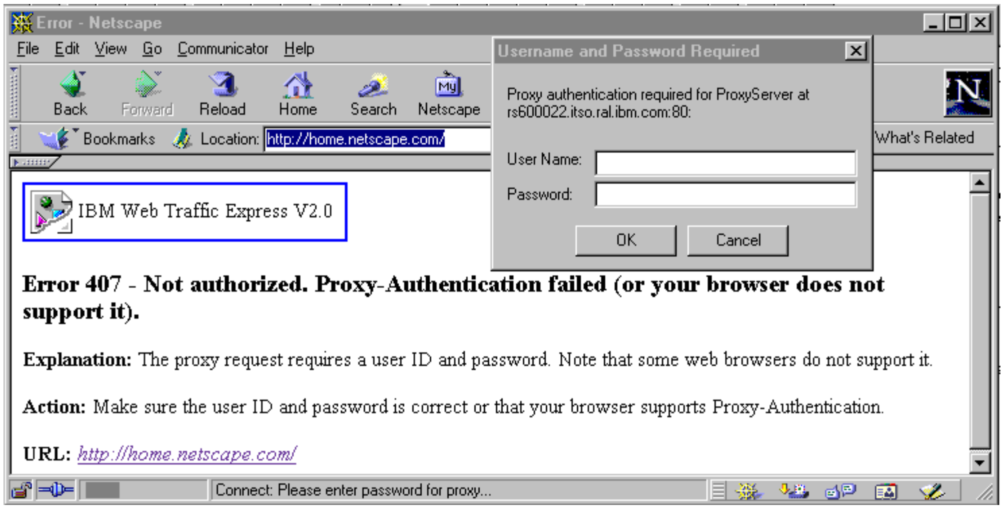
Error - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape
Bookmarks Location: http://home.netscape.com/
What's Related
OK Cancel

Username and Password Required

Proxy authentication required for ProxyServer at rs600022.itso.ral.ibm.com:80:

User Name:

Password:

IBM Web Traffic Express V2.0

Error 407 - Not authorized. Proxy-Authentication failed (or your browser does not support it).

Explanation: The proxy request requires a user ID and password. Note that some web browsers do not support it.

Action: Make sure the user ID and password is correct or that your browser supports Proxy-Authentication.

URL: http://home.netscape.com/

Connect: Please enter password for proxy...

Figure 120. Entering a Wrong User ID and Password

Note that the prompt panel reminds you that the protection setup is named ProxyServer, which is the value we defined for the `ServerID` subdirective.

We clicked **OK** to retry, and then we tried to authenticate with a valid user ID and password. We entered user ID `pohyt` and password `pohyt`, since we had defined such a user ID and password in the wteusers.passwd file. We were then able to successfully access the URL `http://home.netscape.com`.

11.5 Proxy Server Protection in a Hierarchy of Proxies

When creating a hierarchy of proxies, if you wish to require a user ID and password authentication to access the proxies, implement this at the lowest level of the hierarchy. User IDs and passwords are not passed upwards in the hierarchy by proxy servers.

At the higher levels of the hierarchy, it would be reasonable to restrict access to only the IP addresses of the lower-level proxies. For example, if a higher level proxy is only supposed to accept requests from two proxies with IP addresses 5.24.19.7 and 5.23.16.5, you could use the following protection setup:

```
Protection UPPER-LEVEL-PROXY-PROT {
   ServerID        UpperLevelProxy
   AuthType        Basic
   Mask            Any@(5.24.19.7, 5.23.16.5)
}
```

Chapter 12. Proxying FTP URLs with WTE

IBM Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2, can proxy requests for FTP URLs to the appropriate FTP server.

WTE cannot be used to proxy requests from an *FTP client*; WTE will only handle FTP requests received from an *HTTP client* specifying the `ftp://` protocol scheme.

The table below shows the supported HTTP methods for file and directory listing requests:

Table 10. WTE FTP Proxy Supported Methods on Requests

Requests	Supported methods
FTP files	GET, PUT, and DELETE
FTP directory listings	GET

As you can see, only the GET, PUT, and DELETE methods are supported for requests for FTP files. Only the GET method is supported for requests for FTP directory listings.

By default, PUT and DELETE are disabled in the WTE configuration file, ibmproxy.conf. To enable FTP file upload, you must enable PUT. To enable FTP file delete, you must enable DELETE. To see how to enable PUT and DELETE in WTE by manually editing the ibmproxy.conf configuration file, see 12.1, “A Note on the HTTP Methods” on page 183. To see how to enable these two methods through the Configuration and Administration forms, see 12.3, “How to Set Up FTP Proxy” on page 186.

Notice that enabling PUT and DELETE for FTP URLs automatically enables the same methods for HTTP URLs. This has security implications, which we discuss in “A Note on Security with DELETE” on page 184 and “A Note on Security with PUT” on page 185.

Before giving more details on how to configure WTE to handle FTP URLs, it is a good idea to give more details about the HTTP methods, for those readers who are not very familiar with them. This is what we do in the next section.

12.1 A Note on the HTTP Methods

Client requests to a Web server include a method field that indicates the action the Web server is to perform on the specified object. The request identifies the object with a URL.

Following is a list of the HTTP methods that typically Web servers support and, for each method, a description of how the server would respond to a client request invoking that specific method. The description assumes the method is enabled.

- DELETE

With this method, the Web server deletes the object identified by the URL. After the object is deleted, the URL is not valid. Enabling the DELETE method in WTE is optional.

In WTE, to enable DELETE, add this directive to the ibmproxy.conf file:

```
enable delete
```

A Note on Security with DELETE

If you have enabled the DELETE method for FTP file deletion, you should define FTP proxy protection for DELETE requests, to prevent unauthorized file deletion at your FTP server (see 12.6, “Protecting a WTE FTP Proxy Server” on page 190 and Chapter 11, “WTE Proxy Server Protection Scenario” on page 175).

- GET

  With this method, the Web server returns whatever data is identified by the URL. If the URL refers to an executable program, the server returns the output of the program.

- HEAD

  With this method, the Web server returns only HTTP document headers without the document body.

- OPTIONS

  Using this method, the request returns information about the communications options on the request/response chain identified by the URL. This method allows a client to determine the options and requirements associated with an object, or the capabilities of a server, without having to act on or retrieve the object.

- POST

  With the POST method, the request contains data and a URL. The Web server accepts the data enclosed in the request as a new subordinate of the resource identified in the URL. The resource, which may be a data-accepting program, a gateway to some other protocol, or a separate program that accepts annotations, processes the enclosed data.

  The POST method is designed to handle annotation of existing resources; posting of a message to a bulletin board, newsgroup, mailing list, or similar group of articles; providing a block of data, such as data from a form to a data-handling program; or extending a database through an append operation.

  In WTE, the POST method is used to process the Configuration and Administration forms.

- PUT

  The request contains data and a URL. The Web server stores the resource identified in the URL. If the resource already exists, PUT replaces it. If the resource does not exist, PUT creates it.

  For more information on configuring your WWW server to support the PUT method see `http://www.w3.org/Amaya/User/Put.html`.

  Enabling the PUT method allows files to be written to the server using HTTP and FTP.

To enable PUT in WTE, edit the ibmproxy configuration file, adding the following line:

```
enable put
```

A Note on Security with PUT

Because PUT allows clients to write to your server, you should use protection setups to define who can use this method on which files. To achieve file upload security, define a protection scheme that does not allow PUT (or allows it for only specific IP addresses) or require a user ID and password to access it. Then apply the protection scheme to FTP using the `Protect` directive, for example:

```
Protect ftp:* noput-prot
```

See 12.6, “Protecting a WTE FTP Proxy Server” on page 190 and Chapter 11, “WTE Proxy Server Protection Scenario” on page 175 for more details on WTE protection.

- TRACE

  With this method, the Web server echoes the request message sent by the client. This method allows the client to see what is being received at the other end of the request chain and use that data for testing or diagnostic information. The `Content-Type` of the response is `message/http`.

For information on how the Apache HTTP Server supports the GET, POST and PUT methods, see `http://www.apacheweek.com/features/put`.

12.2 Caching of FTP Files

When a request is made to the FTP server to retrieve a file, WTE will first send the FTP server a LIST request for the file to obtain FTP directory information about the file. If the FTP server responds to the LIST request with a positive completion reply and the directory information for the file, WTE will use the file date from the FTP directory information to create an HTTP `Last-Modified` header with the date parsed from the FTP directory information. The caching function of WTE will then use this `Last-Modified` header, together with the `CacheLastModifiedFactor` directive in the configuration file, to determine the length of time that the FTP file should remain in the cache before expiring.

12.2.1 Date and Time

Because the FTP protocol does not have as strict definitions about date and time information as the HTTP protocol does, there are several factors that can cause the `Last-Modified` header generated by WTE for FTP files to be off slightly from the actual file date. These factors are listed below:

- Unlike the HTTP protocol, the FTP protocol does not specify that returned dates must be in Greenwich Mean Time (GMT). For this reason, the date returned by the FTP server is most likely to be in the local time of the FTP server. WTE has no way of determining what time zone the FTP server is running in, so it assumes that the FTP server is in the same time zone as the WTE server. The exception to this is the Windows NT FTP server, which appears to return dates in GMT. If WTE detects that the FTP server is running on Windows NT, it will assume the directory date is in GMT.

- Some FTP servers specify the date in the returned directory information in the format of `Month Day Year` only, and will not include the actual hours or minutes information for the date specified. If the FTP server does not return the hour and minute information for the file, then WTE will assume that the file was last modified on the latest possible hour/minute of the date returned by the FTP server. For example, if the FTP server returned directory information for a file that indicated the file was last modified on October 13, 1998, but did not include information on the hours or minutes, WTE would assume the file was modified at 11:59:59 p.m. on October 13, 1998. Then, assuming that the FTP server was not a Windows NT FTP server, this date would be converted from the local time zone of the WTE server to the corresponding GMT.

When an FTP file expires from the cache, WTE will simulate the HTTP `if modified since` revalidation process for the FTP file. It does this by again issuing the FTP LIST command for the requested file, parsing the file date from the response returned by the FTP server, and comparing this date with the date generated for the `Last-Modified` header when the file was initially retrieved. If the file date has not changed, then WTE will mark the FTP file in the cache as having been re-validated, set a new expiration time for the file, and serve the file from the cache rather than re-retrieving it from the FTP server. If the two file dates do not match, then WTE will go ahead and re-retrieve the file from the FTP server, and store the new copy of the file in the cache with the new file date.

Due to the wide variety of FTP servers available, and the idiosyncrasies of the individual FTP servers, it is not always possible to obtain the directory information for the file from the FTP server. If WTE is unable to determine the file date for the FTP file, it will not generate a `Last-Modified` header for the file. Instead, it will use the value specified for the `CacheDefaultExpiry` directive that matches the URL to determine the length of time to keep the file in the cache. When this time period expires, WTE will always re-retrieve the file from the FTP server. If you notice specific FTP files in your cache that seem to be using the `CacheDefaultExpiry` directive and are constantly being retrieved (generating a high volume of network traffic), consider specifying a more granular `CacheDefaultExpiry` that covers the specific files and holds them in the cache for a longer period of time, for example:

```
CacheDefaultExpiry ftp://ftpserver1/archive/*.exe 2 weeks
```

FTP files which are retrieved for a specific user ID, rather than by anonymous login, are considered to be *private* files and are not cached.

12.3 How to Set Up FTP Proxy

FTP proxy can be configured using the Configuration and Administration forms or manually editing the ibmproxy.conf configuration file.

We will configure the WTE as an FTP proxy by using the Configuration and Administration forms. On the navigation menu:

1. Click **Proxy Configuration** folder.
2. Click **Proxy Settings**.

You will be presented with the WTE Proxy Settings form, as shown:

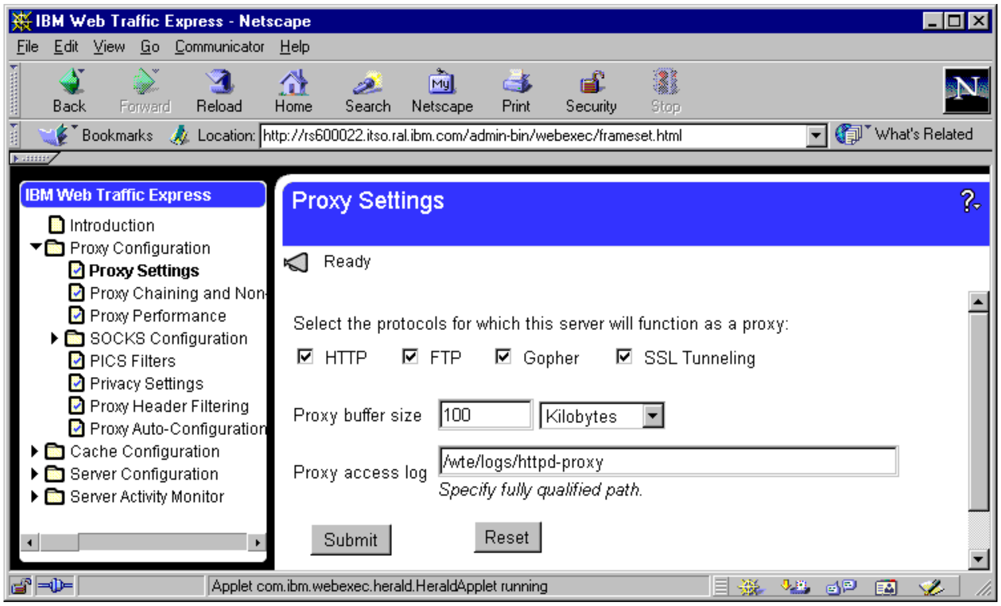

Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html

Proxy Settings

Ready

Select the protocols for which this server will function as a proxy:

☑ HTTP ☑ FTP ☑ Gopher ☑ SSL Tunneling

Proxy buffer size 100 Kilobytes

Proxy access log /wte/logs/httpd-proxy

Specify fully qualified path.

Applet com.ibm.webexec.herald.HeraldApplet running

Figure 121. FTP Proxying - Proxy Settings

To enable FTP proxy support:

1. Click the **FTP** box.
2. Click the **Submit** button to update the changes.

By default, PUT and DELETE are disabled in the ibmproxy.conf configuration file. To enable FTP file upload, you must enable PUT. To enable FTP file deletion, you must enable DELETE. On the navigation menu:

1. Click the **Server Configuration** folder.
2. Click the **Request Processing** folder.
3. Click **HTTP Methods.**

You will be presented with the WTE HTTP Methods form, as shown:

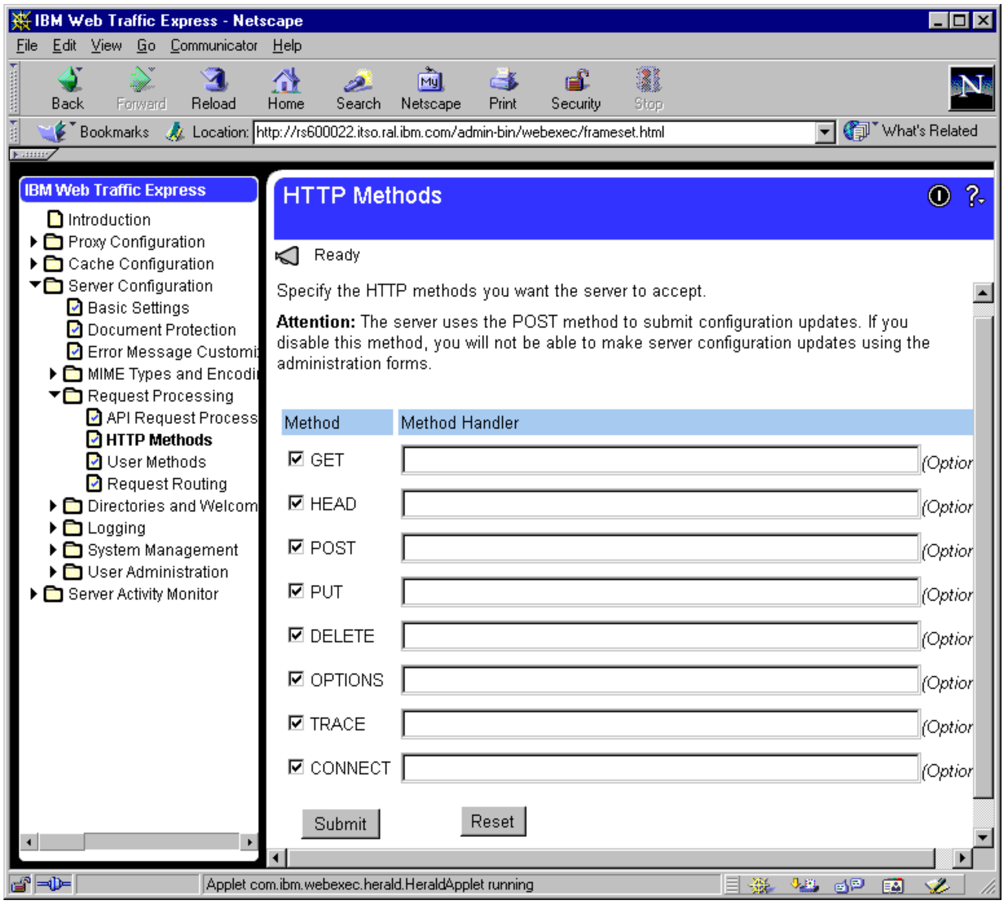

HTTP Methods

Ready

Specify the HTTP methods you want the server to accept.

Attention: The server uses the POST method to submit configuration updates. If you disable this method, you will not be able to make server configuration updates using the administration forms.

Method	Method Handler
☑ GET	(Optior
☑ HEAD	(Optior
☑ POST	(Optior
☑ PUT	(Optior
☑ DELETE	(Optior
☑ OPTIONS	(Optior
☑ TRACE	(Optior
☑ CONNECT	(Optior

Figure 122. FTP Proxying - HTTP Methods

To enable upload (PUT) and deletion (DELETE) of FTP files:

1. Click the **PUT** box.
2. Click the **DELETE** box.
3. Click the **Submit** button to update the changes.

12.4 Identifying the User ID and Password for FTP Server Access

If no user ID or password is specified in the request URL, WTE attempts to login to the requested FTP server anonymously (using the user ID anonymous). Most FTP servers require an e-mail address as the password for anonymous FTP. If the FTP server asks for a password for the anonymous login, WTE will send the e-mail address specified by the `WebMasterEMail` directive in the ibmproxy.conf configuration file.

If the FTP server in the request URL requires a specific user ID and password to login, you can enter the user ID and password in the request URL, according to the following syntax: `ftp://`*`user_ID`*`:`*`password@ftp_server_host`*`/`.

If you do not wish to specify the password for the FTP user ID in the request URL, you can enter just the user ID in the URL. The syntax in this case is: `ftp://`*`user_ID@ftp_server_host`*. WTE will first attempt to login to the FTP server using the specified user ID with no password. If the login with no password is unsuccessful, then the browser will prompt you for the password to be used with the specified user ID.

At least the user ID must be specified in the URL to use anything other than anonymous login. If the user ID is not specified, anonymous login will be used and you will not be prompted for the user ID.

12.5 Specifying the Directory Path Mode for FTP URLs

WTE allows you to specify whether you want the path names in FTP URLs to be interpreted as being *relative* (relative to the logged in user's working directory) or *absolute* (relative to the root directory).

To change how FTP URLs should be interpreted, on the navigation menu:

1. Click the **Proxy Configuration** folder.
2. Click **Proxy Performance**.

You will be presented with the WTE Proxy Performance form, as shown below.

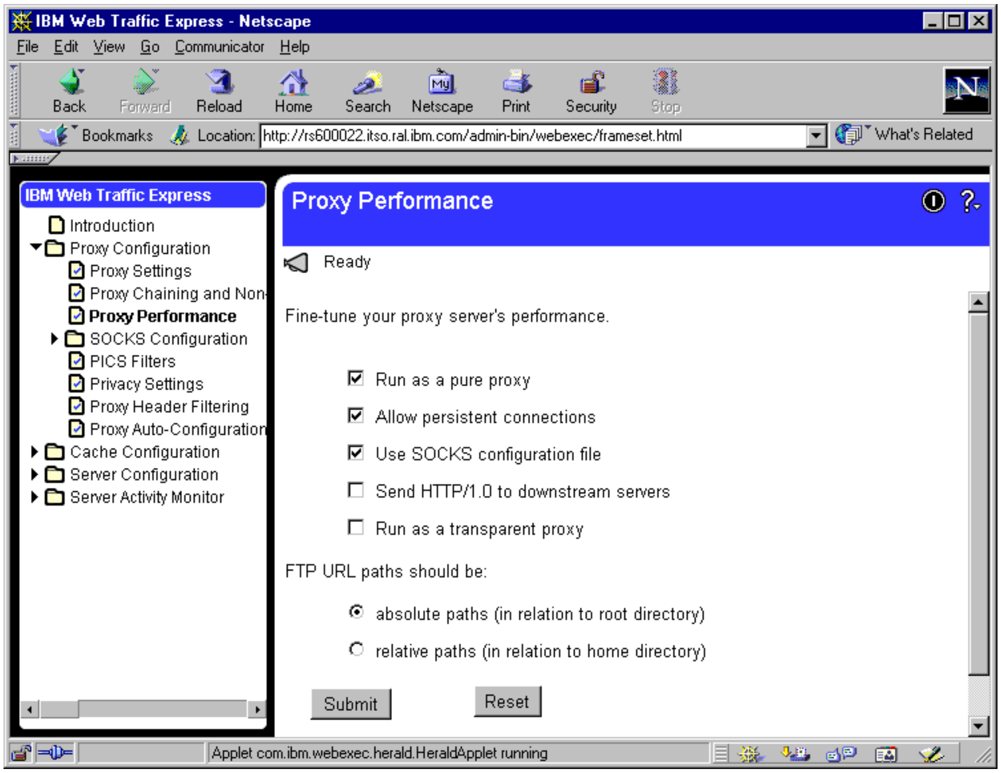

Proxy Performance

Ready

Fine-tune your proxy server's performance.

☑ Run as a pure proxy

☑ Allow persistent connections

☑ Use SOCKS configuration file

☐ Send HTTP/1.0 to downstream servers

☐ Run as a transparent proxy

FTP URL paths should be:

● absolute paths (in relation to root directory)

○ relative paths (in relation to home directory)

Submit Reset

Figure 123. FTP Proxying - FTP URLs

3. Click either the **absolute paths** or **relative paths** box to specify how the FTP URL paths should be interpreted.
4. Click the **Submit** button to make the changes to the ibmproxy.conf configuration file.

For example, if the user ID user1 with password user1pw had a default working directory of /export/home/user1 in the FTP server FTPhost, and in this directory was a subdirectory named test containing a file named test1.exe, the URL to retrieve this file from the FTP server would be specified as shown in the following table:

Table 11. FTP URL Paths

FTP URL path	URL Specification
absolute	ftp://user1:user1pw@FTPhost/export/home/user1/test/test1.exe
relative	ftp://user1:user1pw@FTPhost/test/test1.exe

If relative FTP URL paths are used, you can still access files specifying an absolute path name by using the convention of escaping the initial slash (`/`) with its URL codification `%2F` to indicate the root directory. For example, if user1 wanted to access a file test2.exe in user2's working directory, /export/home/user2, this file could still be accessed using the URL ftp://user1:user1pw@FTPhost/%2Fexport/home/user2/test2.exe. This URL would be interpreted as being relative to the root directory / even if relative FTP URL path names were being used.

12.6 Protecting a WTE FTP Proxy Server

In this section, we describe how you can configure WTE FTP proxy server protection. A more general discussion on WTE protection can be found in Chapter 11, “WTE Proxy Server Protection Scenario” on page 175.

To set up FTP proxy server protection, a protection scheme must be specified previously in the configuration file. If configured, the `Protect` directive must appear before its corresponding `Proxy` directive. Our sample configuration directive is as shown:

```
Protect ftp://*    Protection_Scheme_Name
```

You can also specify a more detailed request template to use different protection schemes for different FTP servers or files, such as:

```
Protect ftp://ftpserverhost1/Testfiles/*   FTPHOST1-PROXY-PROT
Protect ftp://ftpserverhost2/*             FTPHOST2-PROXY-PROT
```

Security with PUT and DELETE

As we mentioned in 12.1, “A Note on the HTTP Methods” on page 183, if you have enabled the PUT method for FTP file upload or the DELETE method for FTP file deletion, it is recommended to define also the FTP proxy protection for at least PUT and DELETE requests. This is to prevent unauthorized file updating at your FTP server.

After editing the ibmproxy.conf configuration file, we need to restart the WTE server to make the changes effective (see 3.5, “Starting and Stopping WTE” on page 70).

12.7 Proxy Chaining FTP Requests

If you are using multiple WTE proxy servers chained together, you can specify that requests containing FTP URLs be sent to a chained proxy server rather than directly to the FTP server. The following figure shows an FTP proxy chaining scenario:

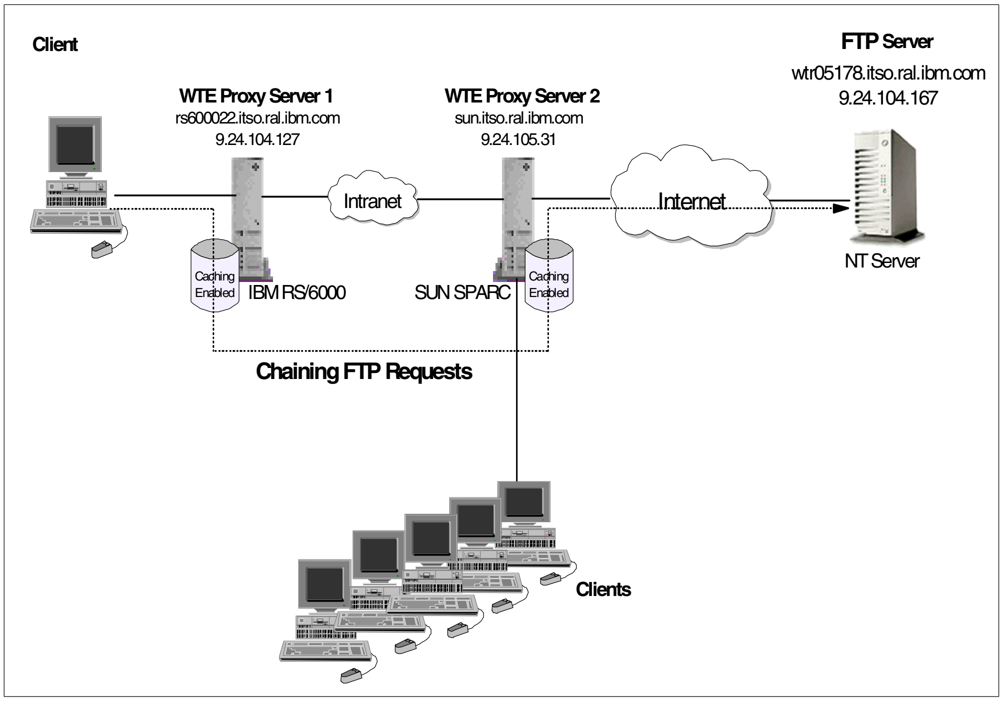

Figure 124. Chaining FTP Requests

When the file is not found in the WTE Proxy Server 1, WTE Proxy Server 1 will contact WTE Proxy Server 2 for the file, instead of contacting the FTP server named in the request URL directly. If WTE Proxy Server 2 does not have the file, WTE Proxy Server 2 will then contact the FTP server, fetch the file, cache it and send the file to WTE Proxy Server 1, which also caches a copy of the file, and serves the browser client. If WTE Proxy Server 2 already has the file in its cache (which may be probable because WTE Proxy Server 2 is higher than WTE Proxy Server 1 in the proxy chain, and a larger number of clients use it as their proxy), then WTE Proxy Server 2 will serve the page back to WTE Proxy Server 1 and requests will not be sent across the Internet.

When specifying the URL of the chained proxy in the `ftp_proxy` directive, the `http://` protocol scheme is used even though this is chaining an `ftp://` protocol scheme request.

See Chapter 10, “WTE Proxy Chaining” on page 163 for more details on WTE proxy chaining.

Chapter 13. Using WTE as a Reverse Proxy Server

Suppose a company wants to allow customers to order products over the Internet, for example, using IBM WebSphere Application Server. Connecting IBM WebSphere Application Server directly to the Internet exposes the entire ordering system to potential attacks by outsiders. To eliminate that exposure, this company should put IBM WebSphere Application Server behind a firewall. Then, Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2, should be placed between the firewall and the Internet, and configured as a *reverse proxy server* (see 2.9, “Reverse Proxy” on page 45). In this case, WTE (configured to be a reverse proxy) will represent the Web server. Now, only WTE needs to get through the firewall. This means that the customer ordering system is protected from potential attacks because outsiders only access the WTE server.

The following figure offers a graphical representation of a reverse proxy environment:

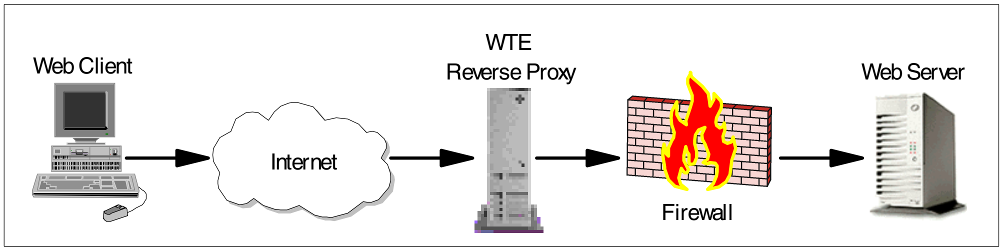

Figure 125. Reverse Proxy Basic Scenario

Notice that Figure 125 on page 193 shows only a basic reverse proxy scenario. In reality, IBM WebSphere Performance Pack allows you to implement much more complex architectures, where multiple WTE reverse proxy servers are load-balanced by a SecureWay Network Dispatcher (ND) server and share the same cache using the AFS Enterprise File System (AFS).

13.1 How Reverse Proxy Works

In a reverse proxy scenario, a client machine does not need to be aware that a reverse proxy exists. The user on the client machine points its browser to the WTE reverse proxy server’s URL, and the WTE reverse proxy server actually goes through its designated SOCKS server to get the page from the specified Web server behind the firewall as defined in its proxy configuration file. This is transparent to the end user, who only has to invoke a Web page from a publicized URL. It is WTE that returns the requested pages to the client. Therefore, from the client’s perspective, the WTE reverse proxy server represents the Web server.

Notice that, if `Caching` is set to `On`, cacheable pages are cached by the WTE reverse proxy server. This way, the Web server is invoked only for pages that are not yet cached and for dynamic pages, which are not cacheable. Therefore, putting a WTE reverse proxy between the Internet and the firewall has also the advantage of reducing the workload on the Web server machine, which in this way is mainly used for dynamic content only.

13.2 How to Set Up Reverse Proxy

To configure WTE as a reverse proxy server, we need to edit the proxy configuration file manually. By default, on AIX and Solaris, this file is /etc/ibmproxy.conf, while on Windows NT, it is C:\WWW\ibmproxy.conf.

The first parameter of the `Proxy` directive specifies the template to be used by the client to access the content on the origin server. The second parameter identifies the template for the actual URL of the content to be retrieved.

A reverse proxy server can be protected using the `PROTECT` directive.

In the proxy configuration file, change the default Pass configuration directive in the Mapping rules section to the Proxy directive, mapping the server_root to the designated Web Server.

Table 12. Reverse Proxy Configuration

Operating System	Changes in the Proxy Configuration File

AIX	Change:
	`# *** ADD NEW MAPPING RULES HERE ***`
	`Pass       /*             /usr/lpp/internet/server_root/pub/*`
	To:
	`# *** ADD NEW MAPPING RULES HERE ***`
	`Proxy      /*             http://www.ibm.com/*`

Solaris	Change:
	`# *** ADD NEW MAPPING RULES HERE ***`
	`Pass      /*             /usr/internet/server_root/pub/*`
	To:
	`# *** ADD NEW MAPPING RULES HERE ***`
	`Proxy    /*              http://www.ibm.com/*`

Windows NT	Change:
	`# *** ADD NEW MAPPING RULES HERE ***`
	`Pass            /*`
	To:
	`# *** ADD NEW MAPPING RULES HERE ***`
	`Proxy    /*              http://www.ibm.com/*`

13.3 Reverse Proxy Scenario

We tested the reverse proxy scenario using the setup shown in Figure 126 on page 195. Here, we will use an example with the objective to protect IBM Corporation’s Web server, `www.ibm.com`.

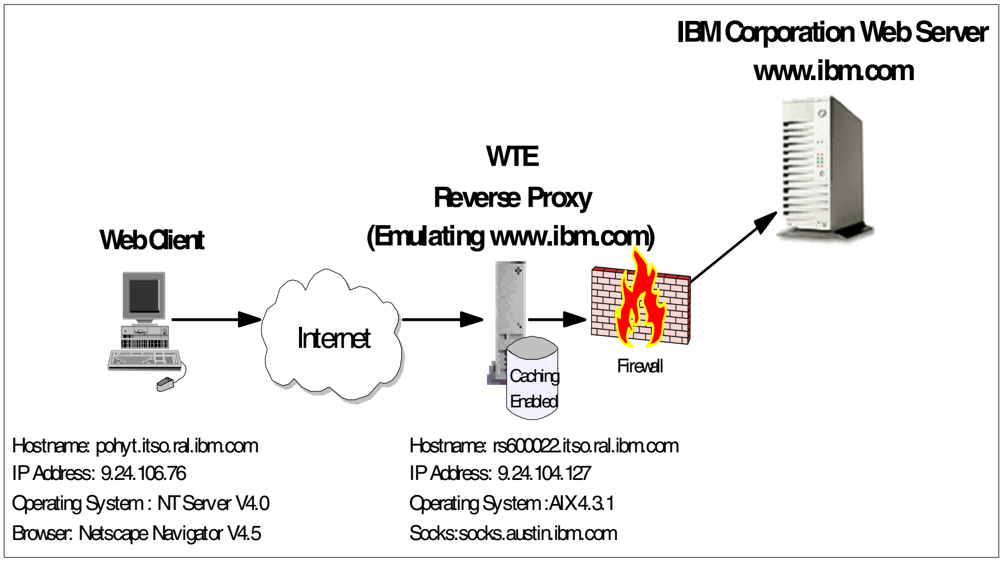

Hostname: pohyt.itso.ral.ibm.com
IP Address: 9.24.106.76
Operating System : NT Server V4.0
Browser: Netscape Navigator V4.5

Hostname: rs600022.itso.ral.ibm.com
IP Address: 9.24.104.127
Operating System : AIX 4.3.1
Socks:socks.austin.ibm.com

Figure 126. Reverse Proxy Basic Scenario

We first deployed a WTE reverse proxy server, whose host name was rs600022.itso.ral.ibm.com, outside the IBM firewall. On the AIX machine rs600022.itso.ral.ibm.com, we modified the /etc/ibmproxy.conf file as described above. After the changes, we restarted the WTE server.

We then announced to the world that the IBM Corporation home page is at `http://rs600022.itso.ral.ibm.com`.

From the Web client machine, whose host name is pohyt.itso.ral.ibm.com, we point our Web browser to the URL `http://rs600022.itso.ral.ibm.com`, and we get the IBM Corporation home page as shown below in Figure 127 on page 196. We need not change the browser proxy configuration.

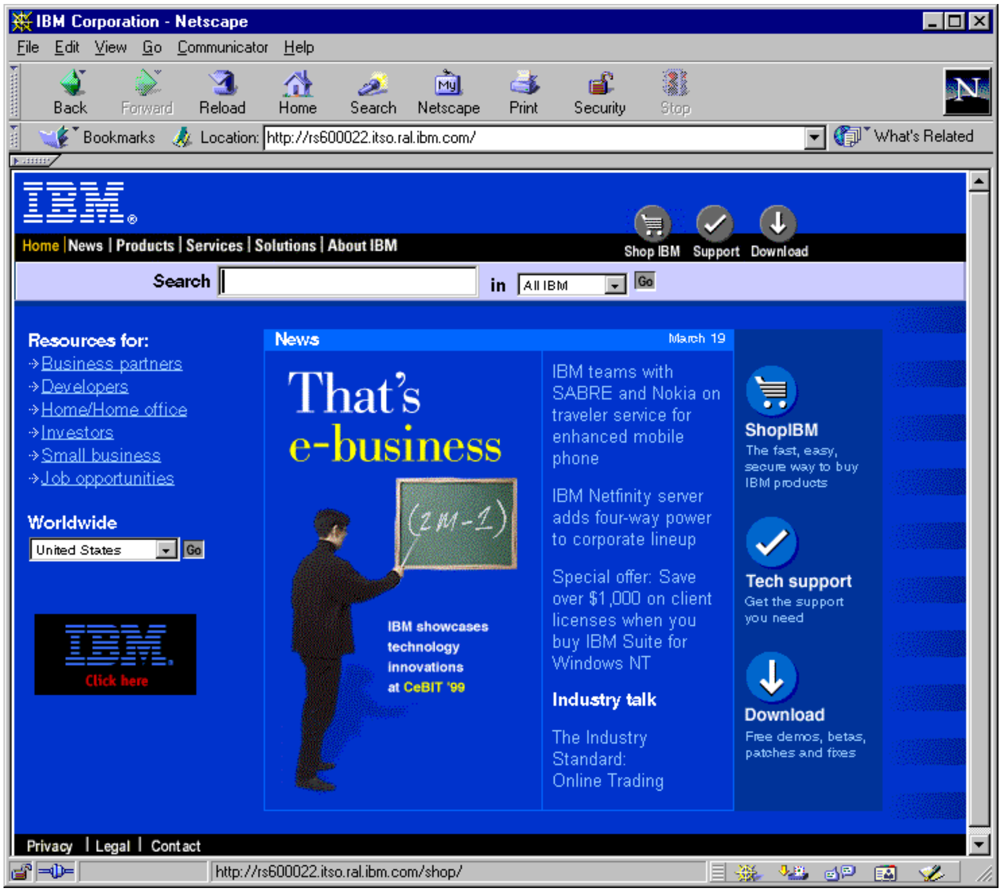
IBM Corporation - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/
What's Related
IBM
Home | News | Products | Services | Solutions | About IBM
Shop IBM Support Download
Search
in All IBM Go
Resources for:
Business partners
Developers
Home/Home office
Investors
Small business
Job opportunities
Worldwide
United States Go
IBM Click here
News March 19
That's e-business
(2M-1)
IBM showcases technology innovations at CeBIT '99
Industry talk
ShopIBM
Tech support
Download
Privacy | Legal | Contact
http://rs600022.itso.ral.ibm.com/shop/

IBM teams with SABRE and Nokia on traveler service for enhanced mobile phone

IBM Netfinity server adds four-way power to corporate lineup

Special offer: Save over $1,000 on client licenses when you buy IBM Suite for Windows NT

The Industry Standard: Online Trading

The fast, easy, secure way to buy IBM products

Get the support you need

Free demos, betas, patches and fixes

Figure 127. http://www.ibm.com is now http://rs6000.itso.ibm.com

This completes our example on the WTE reverse proxy scenario.

Reverse Proxy Configuration and Administration Forms

Once a WTE machine has been configured to act as a reverse proxy, you cannot access the WTE Configuration and Administration forms in the usual way. For example, once rs600022.itso.ral.ibm.com has become reverse proxy to www.ibm.com, you can no longer configure WTE by pointing a Web browser to the URL `http://rs600022.itso.ral.ibm.com`. In fact, because of the `Pass` statement modification described in Table 12 on page 194, the URL above would display the home page of the IBM Web site, as shown in Figure 127 on page 196.

In order to access the Configuration and Administration forms, you should directly point your browser to the URL `http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html`.

13.4 Protecting a Reverse Proxy Server

In this section, we demonstrate how to protect a reverse proxy server by manually editing the WTE proxy configuration file, ibmproxy.conf. We will configure WTE protection for a WTE reverse proxy server using the Protect directive.

To set up reverse proxy protection, a protection scheme must be specified previously in the configuration file. If configured, the Protect directive must appear before its corresponding Proxy directive. Our sample configuration directives is as shown in the following screen:

```
Protection   PROT-REVPROXY {
ServerID     Proxy_Authorization
AuthType     Basic
GetMask      All@(*)
      PutMask      All@(*)
      PostMask     All@(*)
      Mask         All@(*)
      PasswdFile   /etc/revproxy.passwd

}

Protect /*    PROT-REVPROXY
Proxy   /*    http://www.ibm.com/*
```

After we edited the proxy configuration file, we needed to restart the WTE server to make the changes effective (see 3.5, “Starting and Stopping WTE” on page 70). We also created a password file /etc/revproxy.passwd and added the user pohyt, as shown in the following session screen:

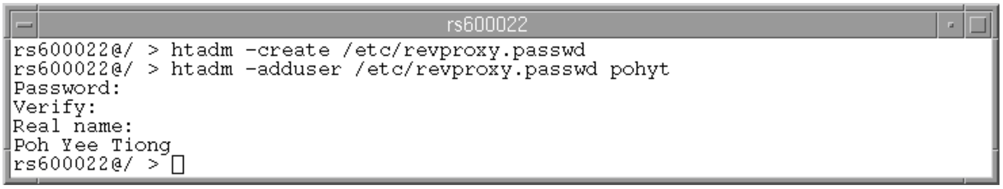

```
rs600022@/ > htadm -create /etc/revproxy.passwd
rs600022@/ > htadm -adduser /etc/revproxy.passwd pohyt
Password:
Verify:
Real name:
Poh Yee Tiong
rs600022@/ > 
```

Figure 128. Creating the Password File and a New User

When a client first tries to access the IBM Corporation home page, the user will be required to enter an authorized User Name and Password, as shown below:

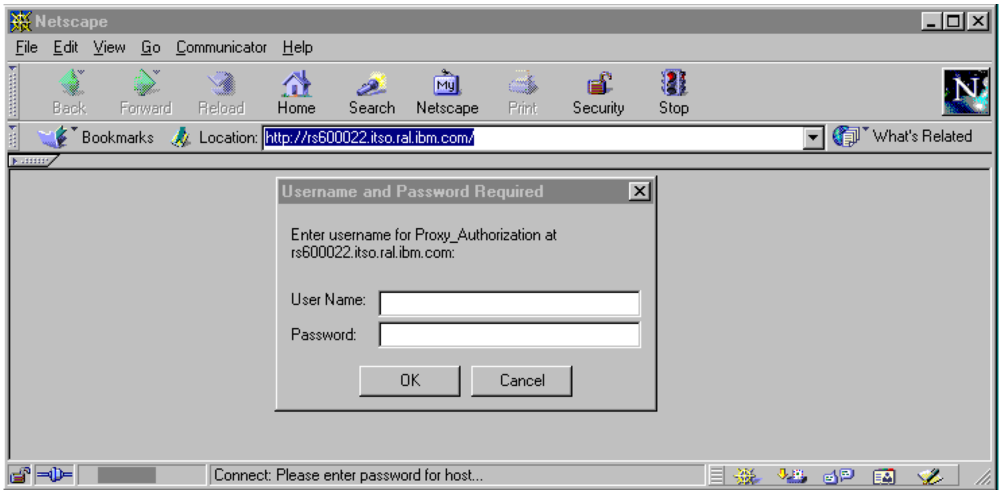

Location: http://rs600022.itso.ral.ibm.com/

Username and Password Required

Enter username for Proxy_Authorization at rs600022.itso.ral.ibm.com:

User Name:

Password:

Connect: Please enter password for host...

Figure 129. Protecting the Reverse Proxy - Authentication 1

When configured as a protected reverse proxy server, WTE returns the following status code when a user attempts to access it without the necessary credentials:

```
Error 401 - Not Authorized - Authentication Failed
```

This status code is normally used to indicate protected files, and differs from the status code:

```
Error 407 - Not Authorized. Proxy-Authentication Failed
```

normally returned by a protected proxy server.

The client does not know the request is being proxied and therefore does not understand a 401 status code in response to the request.

If the content at the origin server is protected, the origin server also returns the status code:

```
Error 401 - Not Authorized -Authentication Failed
```

indicating that a correct User Name and Password are required to access the content. When the client receives a 401 status code response, it resends the request with an `Authorization` header containing the requested User Name and Password. The following is a screen when authentication failed because the User Name and Password entered were not correct:

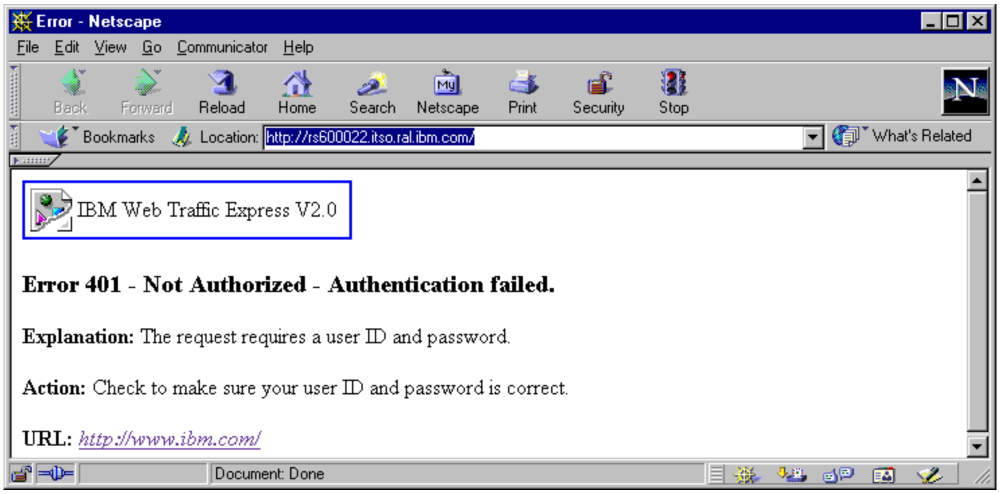
Error - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/
What's Related
IBM Web Traffic Express V2.0
Document: Done

Error 401 - Not Authorized - Authentication failed.

Explanation: The request requires a user ID and password.

Action: Check to make sure your user ID and password is correct.

URL: *http://www.ibm.com/*

Figure 130. Protecting the Reverse Proxy - Authentication Failed

We re-entered the correct user name and password:

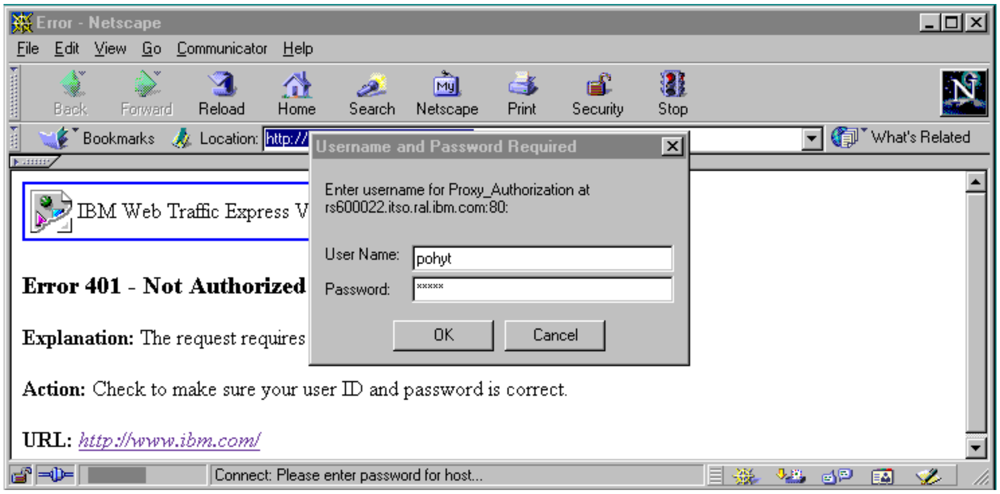
Error - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://
What's Related
OK
Cancel

Username and Password Required

Enter username for Proxy_Authorization at rs600022.itso.ral.ibm.com:80:

User Name: pohyt

Password: *****

IBM Web Traffic Express V

Error 401 - Not Authorized

Explanation: The request requires

Action: Check to make sure your user ID and password is correct.

URL: http://www.ibm.com/

Connect: Please enter password for host...

Figure 131. Protecting the Reverse Proxy - Authentication 2

Because there is no way to send an `Authorization` header for both the reverse proxy server and the content at the origin server on the same request, the reverse proxy server and the content at the origin server cannot both be protected at the same time (this can be done with a regular proxy server). If this occurs, the reverse proxy server will return to the client the following response status code:

```
Error 403 - Access Forbidden
```

To summarize, *the reverse proxy server and the content cannot both be protected simultaneously; only either the reverse proxy server or the content may be protected.*

After a successful authentication, we were able to access the IBM Corporate home page. This completes our example on the WTE reverse proxy server protection.

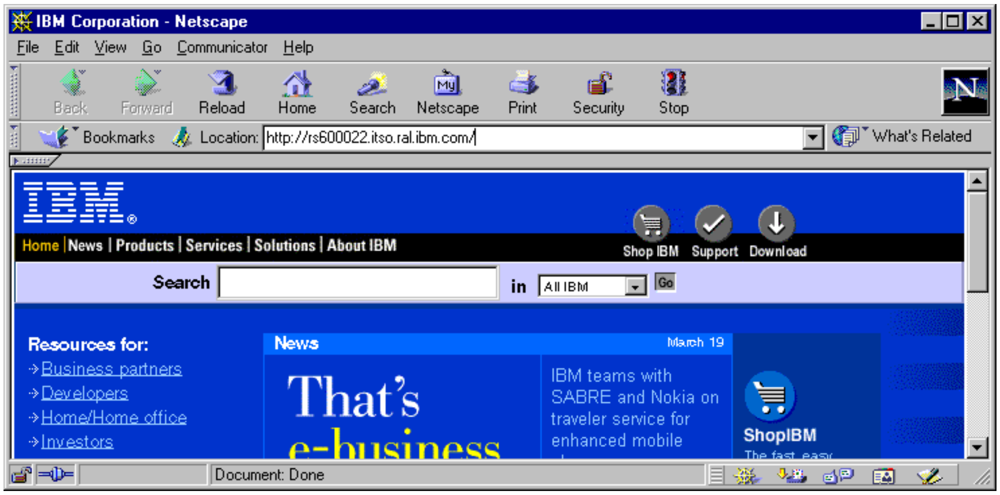
IBM Corporation - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location:
What's Related
IBM
Home | News | Products | Services | Solutions | About IBM
Shop IBM Support Download
Search
in All IBM Go
Resources for:
Business partners
Developers
Home/Home office
Investors
News March 19
That's e-business
ShopIBM
The fast, easy

Location: http://rs600022.itso.ral.ibm.com/

IBM teams with SABRE and Nokia on traveler service for enhanced mobile

Document: Done

Figure 132. Protecting the Reverse Proxy - Authentication Successful

Chapter 14. WTE PICS Support

Although some browsers allow you to restrict which sites users can view, this configuration can be easily changed by a browser administrator. Parents that want to prevent their children from accessing offensive content on the Web know that their children are usually very smart and knowledgeable. Although this kind of browser protection is password-protected, it is enough to uninstall the browser and reinstall it again to overcome content filters at the browser level.

Another limitation of content filtering at the browser level is that enforcing restrictions and management of a large pool of these Web browsers would be a very difficult task. For this reason, an enterprise that wishes to apply Web content filtering to traffic accessing the intranet would face the difficulty of configuring all of the browsers in the company, and still run the risk that employees could overwrite the configuration as explained above.

The solution to the problems just described comes with IBM Web Traffic Express, the caching and filtering component of IBM WebSphere Performance Pack V2. WTE allows you to apply the same content filtering control *at the proxy level* based on Platform for Internet Content Selection (PICS) labels. WTE can also be configured to selectively pass or fail certain URLs. When content filtering is applied at the proxy server level, the filtering configuration can be modified and adjusted only by the proxy server administrator, and it is easily possible to centrally manage content filtering without having to enforce restrictions on each machine.

14.1 What Is the PICS Protocol?

The PICS protocol is a technological standard that builds up a content labeling and filtering system for Web information. It is formed by a set of specifications used to create and manage ratings for the information published on any given Web site. The idea behind PICS is relatively simple. Since people have all kinds of different preferences and values, there should be a ratings and labeling standard that people can use to choose what content they would like to see, or wish to exclude, based on certain parameters.

PICS development began in mid-1995, when the computing and online industries became sensitized to the possibility of online content censorship by the United States government and other governments around the world. Working under the aegis of the World Wide Web Consortium (W3C), the interested parties came together to create a technological solution that would support different rating systems. The PICS standard was adopted in May 1996. For the most up-to-date PICS information, see the PICS Web page of the World Wide Web Consortium (W3C) at the URL `http://www.w3.org/PICS/`, partially shown in the following figure:

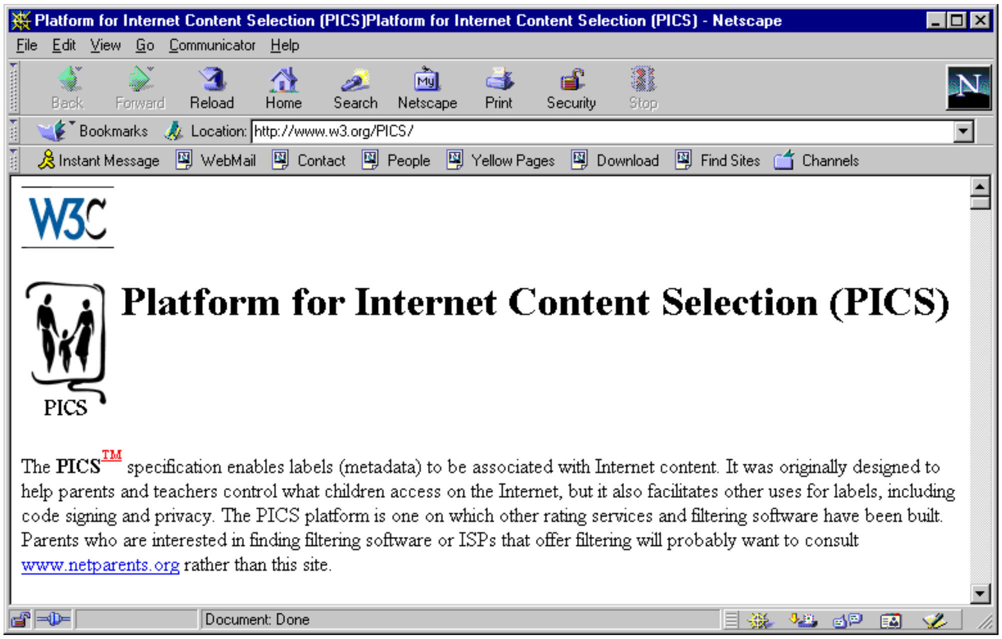

Platform for Internet Content Selection (PICS)Platform for Internet Content Selection (PICS) - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://www.w3.org/PICS/
Instant Message WebMail Contact People Yellow Pages Download Find Sites Channels
W3C
PICS
Document: Done

Platform for Internet Content Selection (PICS)

The **PICS**[TM] specification enables labels (metadata) to be associated with Internet content. It was originally designed to help parents and teachers control what children access on the Internet, but it also facilitates other uses for labels, including code signing and privacy. The PICS platform is one on which other rating services and filtering software have been built. Parents who are interested in finding filtering software or ISPs that offer filtering will probably want to consult www.netparents.org rather than this site.

Figure 133. PICS Web Page at the W3C Web Site

14.1.1 PICS Content Labels

To have a common content rating and filtering system (such as the one described by PICS) effectively working, it is necessary that the main actors in the Web world (Web clients on one side, and Web servers on the other) cooperate in a synergistic way. Web clients are required to use PICS-compliant browsers or PICS-compliant browser-side software to check the content of a Web site, and Web site administrators are required to label their documents using the PICS standard.

Documents should be labeled according to the PICS specifications for the labels, as they are published at `http://www.w3.org/TR/REC-PICS-labels` by the W3C Consortium. A PICS *content label* associated with a document is a piece of information that describes rating information about that document. It contains the rating of the document in reference to various dimensions or categories and specifies who has provided the label and when, along with other information.

14.1.2 Rating Systems

The PICS protocol defines a standard about how the ratings should be specified and transmitted, but it does not standardize the criteria by which you can rate the document’s content. What PICS means to standardize is the method to rate the information. Such a method is also known as a *rating system*. A rating system specifies the dimensions used for labeling, the scale of allowable values on each dimension, and a description of the criteria used in assigning values.

For example, a certain rating system can rate the humorous content in a document based on a single dimension or category named Humor, with allowable values 0, 1, 2, 3, 4, and 5.

14.1.3 Rating Services

PICS does not intend to define an objective system of preferences. In other words, anyone can establish their own criteria to rate the information on the Web.

If you are a webmaster, you can decide on a rating system, rate your documents by creating the rating labels and then publish the documents to the world. What you should do is make your documents compliant to the PICS standard, so that your ratings can be understood by PICS-compliant programs on any platform, anywhere in the world.

Ratings can be assigned to a specific document, a group of documents or even to an entire Web site. The only problem now is that Web clients need to be sure that the ratings assigned correspond to what they expect. How can they trust you? How can they make sure that your rating criteria, meaning, your own rating system, meet their own preferences? The answers to these questions are in the following subsections.

14.1.3.1 How Labels Are Provided

Typically, Web sites do not rate themselves, although they could, considering that the PICS specifications do not forbid them to do so. What Web sites usually do is ask to be rated by a third party, called a *rating service*. A rating service is an entity that evaluates Web content according to their own published, well-known criteria. As a webmaster, you can contact one of the rating services to request assistance in assessing and labeling your own site and documents. The W3C Consortium publishes a list of PICS self-rating services at the URL `http://www.w3.org/pub/WWW/PICS/selfrat.htm`.

Typically, after choosing a self-labeling service, you can connect to its Web site and describe the content of a single document, a group of documents or the entire Web site you want them to rate. You do this by filling out an online form. After submitting the form, the service provides you with a text label. At this point there are three possibilities:

1. The text label provided by the rating service can be dynamically embedded in the HTTP headers of the documents to be rated when clients request those documents.

   The Web server administrator should determine whether their Web server software supports transmitting PICS labels in HTTP headers. If such support is confirmed, the text information can be used to create PICS-compliant rating labels embedded in the HTTP header. The Web server administrator should store the labels in the Web server file system and use the PICS configuration file to manage and transmit them.

   This is the preferred method for transmitting PICS labels, since it can be done automatically and does not add a significant overhead. The labels are sent along with the Web pages when a client requests them.

   Notice that Lotus Domino Go Webserver supports this technology. It allows you to store the rating labels for all the documents on your Web site and manage them from a central file. In this way rating labels can be embedded dynamically in the HTTP headers of the requested documents.

   Another advantage of the technology we have described here is that security mechanisms such as message digests and digital signatures can be incorporated in the label creation to grant label validity.

2. The text label provided by the rating service can be statically embedded in the HTML headers of the documents to be rated. This operation must be accomplished when documents are created, or before clients request them.

   If the Web server software does not yet support the HTTP extensions for PICS, the Web server administrator can use `<META>` tags in the HEAD section of all the HTML files to store the PICS labels. This means that the Web server administrator should edit each of the HTML files that need to be rated and insert the rating information in the HEAD sections.

   This process is entirely manual and therefore time-consuming, error-prone and difficult to maintain. It does not allow you to incorporate any of the security mechanisms that could guarantee the validity of the labels, such as message digests and digital signatures.

3. Independent rating services may make use of some procedures or software tools to examine particular URLs and create labels describing those URLs. The rating service stores the labels and distributes them from a separate server, called the *label bureau*. In this case a rating server also operates a label bureau server. Filtering software residing at the client-side needs to be instructed to check at that label bureau to find the labels.

   A label bureau server provides labels dynamically when a client sends a specific request. In other words, if the user's profile asks for labels from a label bureau server, those labels will be requested in parallel with the content.

   Notice that some clients might accept rating information that is embedded in the file, but others might require a separate label from a registered rating service and a guarantee that it was created by that service. In other cases, Web clients might decide to contact the rating service only if the label information is not embedded with the requested document from the server. Then, it might send a subsequent request directly to the rating service asking for the label information for that document.

   Although this technology allows you to incorporate security mechanisms such as message digests and digital signatures to guarantee the validity of the labels, it requires a second connection, which takes longer and can discourage future visits to a specific Web site. The browser needs to wait until the label information is returned before it displays any data. Faster response data is the main reason why rating labels for a site should reside at the site itself.

No restrictions are applied to any of the three processes described above, in that anyone can label any document from any site. The PICS specification does not determine who can or will act as a rating service.

Many new rating systems were created in the last few years by organizations with no interest in providing filtering software. Web users can elect to trust one or more of those organizations, and can decide to accept only ratings supplied by the organizations they trust.

A rating service can choose any criteria on which to rate Web sites. While some might rate Web sites for their violence or sexual content, others could choose to rate on educational content, political correctness, or even how *cool* the site is. Also, a rating service can rate any and all Web sites it wants to rate.

The two leading rating systems, Recreational Software Advisory Council on the Internet (RSACi) and SafeSurf, are based on PICS. The RSACi implements PICS labels using a number rating scale. The RSACi is an independent, non-profit organization based in Washington, D.C., that empowers the public, especially parents, to make informed decisions about electronic media by means of an open, objective, content advisory system. IBM is a corporate sponsor of RSACi and is represented on the RSACi Board of Directors.

So RSACi acts as a rating service which operates also as a label bureau. The RSACi rating system rates parameters such as nudity, graphic language, sexual content, and violence on a 0 (lowest) to 4 (highest) scale. For example, if you want to block sites with violence, but think some foul language is acceptable, you can have a filter accept sites with a violence maximum of 2, and a language maximum of 3.

14.1.3.2 URL Identifiers

Each rating system is identified by a valid URL. This enables several services to use the same rating system and refer to it by its identifier. The URL names that a rating system can be accessed to obtain a human-readable description of the rating system, such as the categories, scales, and intended criteria for assigning ratings. The URL that identifies a rating system is also used to advertise the rating services that use that particular rating system.

A rating service is identified by a URL as well. This identifier is included in all the labels assigned by the rating service. Since the service identifier is a URL, it can be used to retrieve a document. That document can be in any format, but PICS specifications recommend that it be in HTML format and give a description not only of the rating service, but also of the rating system. That document should at least provide a link to another document describing the rating system.

The rating service assigns labels according to some rating systems, and then distributes them, mainly via a label bureau. A label bureau is also identified by a URL, which the PICS-compliant client software should be able to contact to request the labels. In this sense, you can think of a label bureau as a computer system that supplies, via a computer network, ratings of documents.

14.2 PICS Filtering at the Proxy Server Level

PICS is a standard that lists rules for rating the information contained on any given Web site. Rating decisions are made based on particular categories, usually violence and pornography. Some Web browsers are PICS-compliant and are able to filter content information as content is received from the destination server. However, relying on browser settings may not be a safe solution, because such settings can be adjusted at the browser, and may be easily compromised.

PICS-Compliant Web Browsers

Microsoft Internet Explorer Version 4.0 or higher supports the PICS protocol and can effectively block access to specific sites, depending on software configuration. Netscape Navigator has not been made PICS-compliant yet.

With WTE, PICS filtering can be implemented at the proxy server level. This removes the responsibility from the client. The proxy server administrator can directly prevent certain types of information from being served to specific browsers (or to groups of browsers). Browsers with PICS filter settings support can perform further filtering of the Web pages that they receive. This method ensures that clients will only get the level of content specified at the proxy. Interaction between the client and the proxy administrator would be required to change the sensitivity of the filter.

The diagram below shows the logical flow of the PICS filtering at a proxy server level:

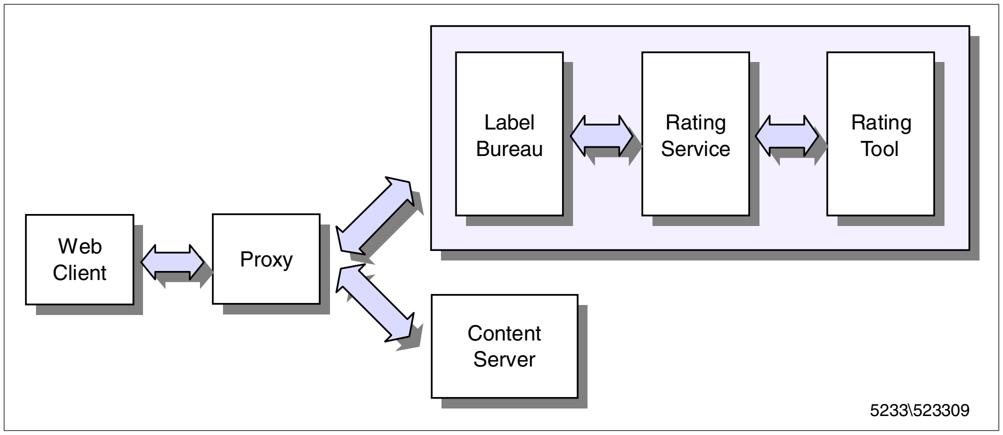

Figure 134. PICS Filtering at the Proxy Server Level

Let’s say the proxy server administrator has set up and enabled a filtering profile on the proxy server. To decide whether a particular document will be passed or blocked, the proxy filtering profile uses the values contained in the PICS labels supplied by a label bureau server, which is in turn managed by a rating service. The rating service might own a rating tool to examine particular URLs and create labels describing those URLs. This rating tool can implement a procedure to discover a new site as soon as it goes online, examine the site and store a PICS label at the label bureau server.

Later, when content from this site is requested by the user's client software, the PICS label is requested from the label bureau server by the proxy. If the profile establishes that the values in this PICS label indicate that the content is not wanted, then the proxy will send an HTML page back to the client, explaining why the content is not being delivered. It may or may not include instructions on how to bypass the blocking or how to initiate the correction of a faulty rating if necessary. However, in most cases, the values in the label mean the content is acceptable, and the content is fetched from the content server to the proxy and then forwarded to the user.

If proxy caching is enabled, during further fetches, the user sees no delay, since the proxy server caches only pages that have previously been accepted, and then serves them directly from its cache.

The label bureau server might tell the proxy that this site has not yet been rated, and the profile determines whether the user is sent the content anyway, or sent a `Not Yet Rated` page by the proxy. The reviewer's tool is notified by the label bureau that this site should be rated as soon as possible. Some label bureau servers in the future will fetch the content and run a program to create an interim rating that will be returned, pending the reviewer's more accurate site evaluation.

The diagram shown in Figure 134 on page 206, which we have just explained, is simplified if the Web server is enabled to embed PICS labels in the Web documents it serves (either in the HTTP header or in the HTML header), and the proxy server is configured to accept such labels without the need to contact an external label bureau. In the next section we show you how to implement a working PICS environment, and discuss both the possibilities where the PICS labels are provided by an external label bureau or by the Web server itself.

14.3 PICS Scenario

In this section we describe how to build a PICS environment where the content filtering is performed at the proxy server level.

We show you a working scenario where the PICS labels generated by a rating service are provided by an associated label bureau (see Point 1 on page 203). Then, we will show you another example where the PICS labels are embedded in the Web documents at the Web server level (see Point 2 on page 204 and Point 3 on page 204).

The purpose of this section is to demonstrate the capability of the Caching and Filtering component of filtering Web content. At the end of the section, it will be clear that content filtering, when applied at the proxy server level rather than at the Web client level, guarantees security, reliability, and centralized management.

14.3.1 Components and Network for the PICS Scenario

The following table summarizes the components of the scenario:

Table 13. PICS Scenario - Hardware, Software, and Network Configuration

Host Name	IP Address	Operating System	Service
pohyt	9.24.106.76	Windows NT Server 4.0	**Web Client** Netscape Navigator V4.5
rs600022	9.24.104.127	AIX 4.3.1	**Caching and Filtering Proxy Server** IBM Web Traffic Express V2.0
wtr05178	9.24.104.167	Windows NT Server 4.0	**Web Server** Lotus Domino Go Webserver Version 4.6.2.5
rs600030	9.24.104.97	AIX 4.3.1	**Rating Service and Label Bureau** Lotus Domino Go Webserver Version 4.6.25

All the machines shown in the table above belonged to the itso.ral.ibm.com domain.

The architecture of the PICS environment we built is shown in the following diagram, which you can easily compare with the general architecture diagram shown in Figure 134 on page 206:

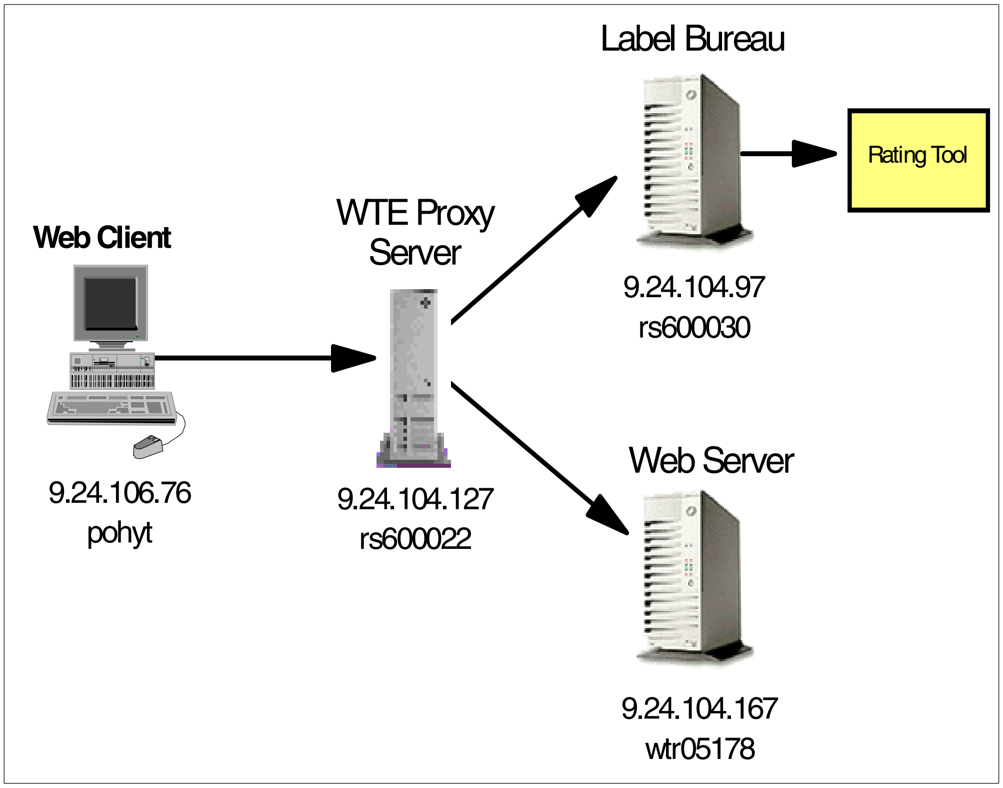

Figure 135. PICS Scenario Environment

The flow in a PICS environment where content filtering is applied at the proxy server level is described in 14.2, “PICS Filtering at the Proxy Server Level” on page 205.

To demonstrate how PICS works, we built up an architecture including all the components involved in a PICS flow.

14.3.2 Configuration of the Web Server

The Web server was configured on the Windows NT Server machine wtr05178, using Lotus Domino Go Webserver Version 4.6.2.5. On the Web server we decided to publish, among others, the following four HTML documents at the URL `http://wtr05178.itso.ral.ibm.com/PICSxmp`:

1. age1.html
2. age2.html
3. age3.html
4. age4.html

These four files are a subset of the sample files distributed with Lotus Domino Go Webserver. After a default installation of Lotus Domino Go Webserver V4.6.2.5 on Windows NT, they are located in the directory \WWW\HTML\PICSxmp.

Acting as webmasters of this site, we chose to have the mentioned HTML pages be rated by a rating service.

14.3.3 Configuration of the Rating Service

A rating service evaluates Web content based on its own published criteria. It then distributes the labels through a label bureau.

We will set up the rating service and the label bureau on an AIX machine, rs600022, using Lotus Domino Go Webserver since it can be configured to act as rating service and label bureau. Lotus Domino Go Webserver can also store rating labels for other Web sites and serve them in response to client requests.

First we must define our own rating system and specify the rating criteria. This is done in a PICS-compliant rating system description file (known as a RAT file). A RAT file describes the rating system used to rate the documents. A rating service must provide a RAT file along with the rating labels for the HTML documents.

To learn more about RAT files, check the W3C Consortium's PICS specifications for rating services and rating systems at the URL `http://www.w3.org/pub/WWW/PICS/services.html`. It includes the syntax for the machine-readable format of RAT files.

The Lotus Domino Go Webserver fileset includes a sample RAT file, called coolness.rat and located in the directory /usr/lpp/internet/server_root/labels. We created the directory /usr/lpp/internet/server_root/wteRatings and copied the coolness.rat file into this new directory. We modified the original file, only changing the second, third and fourth lines to customize the file according to our requirements. The entire file is shown in the following two figures:

```
((PICS-version 1.1)
 (rating-system "http://rs600030.itso.ral.ibm.com/wteRatings/newEnzo.html")
 (rating-service "http://rs600030.itso.ral.ibm.com/wteoRatings/V1-0.html")
 (name "The new Enzo Rating System")
 (description "This rating system is based on sample provided with Domino Go
  WebServer, It categorizes three important criteria: coolness, age-range,
  and amount of graphics.")
 (category
  (transmit-as "Coolness")
  (name "Coolness Index")
  (label
   (name "Way Cool")
   (description "This site is majorly cool")
   (value 1))
  (label
   (name "cool")
   (description "Pretty cool")
   (value 2))
  (label
   (name "Mediocre coolness")
   (description "Tries to be cool but falls a little short")
   (value 3))
  (label
   (name "Not cool")
   (description "Like the name, not cool")
   (value 4))
  (label
   (name "Totally Uncool")
   (description "This site is a waste of time")
   (value 5)))
 (category
  (transmit-as "Age-range")
  (name "age-range")
  (label
   (name "All ages")
   (description "This site is suitable for everyone")
   (value 1))
  (label
   (name "Teenager and older")
   (description "Teenagers and older")
   (value 2))
  (label
   (name "Adult only")
   (description "For adults only")
   (value 3))
  (label
   (name "No one")
   (description "No one should see this site")
   (value 4)))
 (category
  (transmit-as "Graphics")
  (name "number of graphics")
  (label
```

Figure 136. (Part 1 of 2). coolness.rat File

```
   (name "0 to 1")
   (description "Hardly any graphics")
   (value 1))
  (label
   (name "1 to 5")
   (description "A few graphics")
   (value 2))
  (label
   (name "10 to 20")
   (description "A lot of graphics")
   (value 3))
  (label
   (name "20+")
   (description "Plan to be here all day")
(value 4)))
```

Figure 137. (Part 2 of 2). coolness.rat File

Now, we explain the meaning of the sample RAT file.

1. The URL identifier of the rating system is `http://rs600030.itso.ral.ibm.com/wteRatings/newEnzo.html` as shown in the following figure. The document available at that URL should be a human-readable description of the categories, scales, and intended criteria for the assigned ratings.

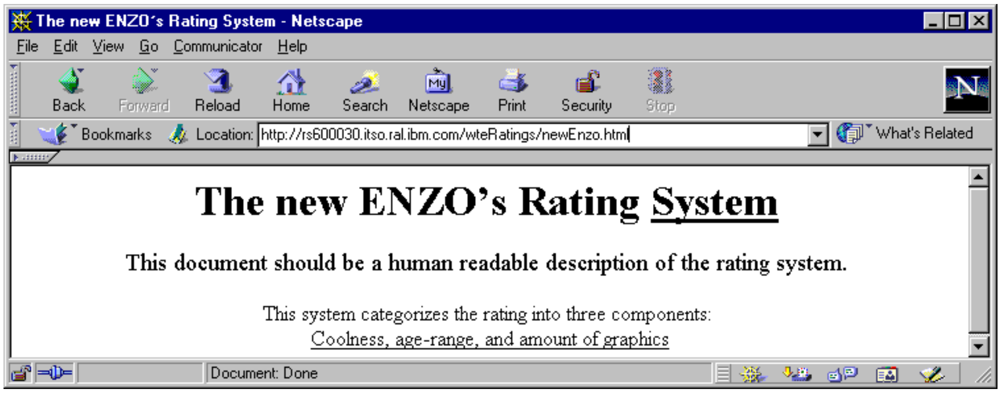
The new ENZO's Rating System - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600030.itso.ral.ibm.com/wteRatings/newEnzo.html
What's Related
Document: Done

The new ENZO's Rating System

This document should be a human readable description of the rating system.

This system categorizes the rating into three components:
Coolness, age-range, and amount of graphics

Figure 138. Web Page Pointed to by the Rating System URL Identifier

2. The URL identifier of the rating service is `http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html` as shown in the following figure. The labels themselves will have this URL in them to identify the rating service that created them. The document available at this URL should be a human-readable description of the rating service.

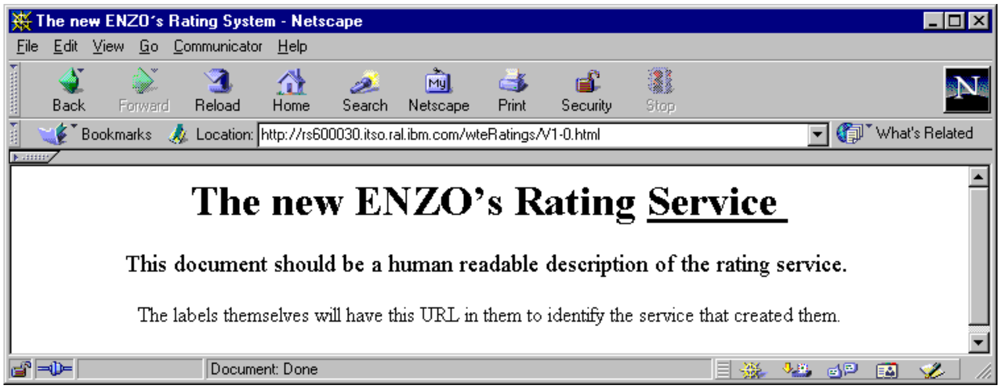
The new ENZO's Rating System - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html
What's Related
Document: Done

The new ENZO's Rating Service

This document should be a human readable description of the rating service.

The labels themselves will have this URL in them to identify the service that created them.

Figure 139. Web Page Pointed by the Rating Service URL Identifier

URL Identifiers Accessibility

In order to make the URLs specified for the rating-service and rating-system parameters accessible from Web browsers, the following `Pass` directive needs to be inserted into the httpd.conf configuration file of the label bureau and rating service machine, rs600030:

```
Pass /wteRatings/*   /usr/lpp/internet/server_root/wteRatings/*
```

3. The name of the rating service is `The new Enzo Rating System`.
4. There are three top-level categories in this rating system. Each category has two names: a short transmission name to be used in labels, and a name that is more easily understood and usually longer than the transmission name.

   For example, the first one has transmission name `Coolness` and the longer name is `Coolness Index`. The second one has a transmission name `Age-range` and name `age-range`. The third one has transmission name `Graphics` and name `number of graphics`.
5. All the categories only allow the integer values 1, 2, 3, and 4. Each label attribute provides a label with a value and associates a name and a description to the value. In a content filtering operation you can use the value or the name of a label, but it is recommended to use the value rather than the name of a label, in order to be independent from the language and to make it easier to write rules expression filtering.

Once we modified and personalized the RAT file describing the rating system, we needed to create the labels for the documents to be rated. We then stored the labels in the label bureau server's directory. In our case, we installed the labels into the wteRatings directory. Actually, since in this case we were acting as rating service administrators, we should have used some rating tools to rate the documents, as explained in 14.2, "PICS Filtering at the Proxy Server Level" on page 205.

Let's assume that we used our own rating tool to produce the following label files for the documents age1.html, age2.html, age3.html, and age4.html respectively:

1. age1.lbl
2. age2.lbl
3. age3.lbl

4. age4.lbl

In reality, we did not need to use any rating tools because Lotus Domino Go Webserver already provides such files, which on AIX, come in the directory /usr/lpp/internet/server_root/labels after a default installation of the product.

We copied those label files into the directory where we had already copied the RAT file, which was /usr/lpp/internet/server_root/wteRatings. Then we made some small changes to the label files, in that we essentially changed the first line of each file to personalize the label according to our system configuration. The label files age1.lbl, age2.lbl, age3.lbl, and age4.lbl are shown in the following four figures respectively:

```
(PICS-1.1 "http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html"
l for %%URL%%
comment "This site is suitable for everyone"
on "1996.04.04T08:15-0500"
r (Coolness 0 Age-range 1 Graphics 0))
```

Figure 140. PICS Label File - age1.lbl

```
(PICS-1.1 "http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html"
l for %%URL%%
comment "For teenagers and older"
on"1996.04.04T08:15-0500"
r (Coolness 0 Age-range 2 Graphics 0))
```

Figure 141. PICS Label File - age2.lbl

```
(PICS-1.1 "http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html"
l for %%URL%%
comment "For adults only"
on "1996.04.04T08:15-0500"
r (Coolness 0 Age-range 3 Graphics 0))
```

Figure 142. PICS Label File - age3.lbl

```
(PICS-1.1 "http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html"
l for %%URL%%
comment "No one should see this site"
on "1996.04.04T08:15-0500"
r (Coolness 0 Age-range 4 Graphics 0))
```

Figure 143. PICS Label File - age4.lbl

The above label files have a very similar structure:

1. Through the string `PICS-1.1`, they specify that the label structure complies to the PICS label specifications Version 1.1, contained in the document PICS Label Distribution Label Syntax and Communication Protocols Version 1.1, found at the URL `http://www.w3.org/TR/REC-PICS-labels`.

2. They report the URL identifier `http://www.w3.org/TR/REC-PICS-labels` of the rating service that has created the label. We remind you that this URL is the same URL specified as a value of the `rating-service` parameter in the RAT file.
3. They show which document the label was created for. In this case, the short version `l` for the keyword `labels` has been used. The value after the keyword `for` should be the URL of the document that the label refers to. So, for example, in the first label file age1.lbl, which refers to the age1.html document, the full line should be:

   ```
   labels for "http://wtr05178.itso.ral.ibm.com/PICSxmp/age1.html
   ```

   Lotus Domino Go Webserver has added extensions to the format of the labels specified at `http://www.w3.org/TR/REC-PICS-labels`. One of these extensions allows you to insert variables in some label files, such as `URL`; the value of this variable is indicated with `%%URL%%`, and the current URL will be substituted for this value. This means that when the label bureau server receives a request for a rating label that contains the string `%%URL%%`, it replaces this string with the correct current URL before sending the label.
4. They may specify some comments for the label by using lines starting with the `comment` statement. These comment lines are sent back to the clients.

   Another extension of Lotus Domino Go Webserver allows you to insert another type of comment for your own use into the label files, by beginning the comment lines with the pound character (`#`). Lines beginning with this character are not sent back to the clients, but are for the use of the rating service and label bureau administrator. This type of comment is an addition to the `comment` statement used inside labels.
5. They report the date on which this rating label was issued.
6. Finally, they specify the value of the rating assigned to the document for each category specified in the RAT file. Notice that it is possible to use the short version `r` for the key word `ratings`.

If we decided to use the complete keywords and replace the `%%URL%%` string with the actual value of the `URL` variable, a label file (for example, age1.lbl) should appear similar to the following:

```
(PICS-1.1 "http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html"
labels for "http://wtr05178.itso.ral.ibm.com/PICSxmp/age1.html
comment "This site is suitable for everyone"
on "1996.04.04T08:15-0500"
ratings (Coolness 0 Age-range 1 Graphics 0))
```

Figure 144. PICS Label File - age1.lbl - Complete Version

14.3.4 Configuration of the Label Bureau

At this point, by creating the RAT file coolness.rat, we established the rating system we intended to apply. We published that rating system through the URL `http://rs600030.itso.ral.ibm.com/wteRatings/newEnzo.html`. We also published the rating service at the URL `http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html`. Finally, we rated some documents of the Web site `http://wtr05178.itso.ral.ibm.com`.

As we are installing the rating service and label bureau on the same machine, we continue to customize the label bureau on the machine rs600030. The function of the label bureau is to serve the stored rating labels in response to client requests.

First, we had to inform the server that it would act as a PICS label bureau, and we had to specify where to direct the PICS rating label requests. We added the following `Service` directive to the configuration file httpd.conf of our rating service and label bureau AIX machine, rs600030:

```
Service /wteRatings INTERNAL:PICS-Ratings
```

We entered `/wteRatings`, so clients who wanted to request labels from our label bureau should request the URL `http://rs600030.itso.ral.ibm.com/wteRatings`. In your configuration, you will replace `/wteRatings` with the relative path of the URL you will use on the label bureau server for label requests. For example, if for label requests you publish the URL `http://www.coolratings.com/CoolSite`, you would only include `/CoolSite` in the `Service` directive.

Next, we had to tell our server which documents had been rated, what host would serve them, and where the labels could be found in the label bureau server machine file system. We specified all of these by editing the PICS configuration file /etc/ics_pics.conf, which is used to associate the rated documents to the specific label files. The /etc/ics_pics file on our system is shown in the following figure:

```
#
# COMPONENT_NAME: web ics_pics.conf
#
# FUNCTIONS:
#
# ORIGINS: 10  26  27
#
# (C) COPYRIGHT Lotus Development Corporation 1997
# (C) COPYRIGHT International Business Machines Corporation 1997
# All Rights Reserved
#
# This software is subject to the Lotus Software Agreement
# Restricted Rights for U.S. government users, and applicable export
# regulations. Lotus is a registered trademark and Lotus Domino Go Webserver
# is a trademark of Lotus Development Corporation.

#
#          Sample PICS Configuration file for
#              Lotus Domino Go Webserver
#
# Please use the remote administration interface to add labels
# to your documents.

DefineLBService "http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html" "The Enzo Rating System"
/usr/lpp/internet/server_root/wteRatings/coolness.rat  {
LABELFILE    /usr/lpp/internet/server_root/wteRatings/age1.lbl   This site is suitable for everyone
LABELFILE    /usr/lpp/internet/server_root/wteRatings/age2.lbl   Teenagers and older
LABELFILE    /usr/lpp/internet/server_root/wteRatings/age3.lbl   For adults only
LABELFILE    /usr/lpp/internet/server_root/wteRatings/age4.lbl   No one should see this site
LABELFILE    /usr/lpp/internet/server_root/wteRatings/cool1.lbl   This site is majorly cool
LABELFILE    /usr/lpp/internet/server_root/wteRatings/cool2.lbl   Pretty cool
LABELFILE    /usr/lpp/internet/server_root/wteRatings/cool3.lbl   Tries to be cool but falls a little short
LABELFILE    /usr/lpp/internet/server_root/wteRatings/cool4.lbl   Like the name, not cool
LABELFILE    /usr/lpp/internet/server_root/wteRatingswteRatings/cool5.lbl   Totally Uncool
LABELFILE    /usr/lpp/internet/server_root/wteRatings/graphics1.lbl   Hardly any graphics
LABELFILE    /usr/lpp/internet/server_root/wteRatings/graphics2.lbl   A few graphics
LABELFILE    /usr/lpp/internet/server_root/wteRatings/graphics3.lbl   A lot of graphics
LABELFILE    /usr/lpp/internet/server_root/wteRatings/graphics4.lbl   Plan to be here all day
}

LabelsFor http://wtr05178.itso.ral.ibm.com http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html  {
    /PICSxmp/age1.html  /usr/lpp/internet/server_root/wteRatings/age1.lbl
    /PICSxmp/age2.html  /usr/lpp/internet/server_root/wteRatings/age2.lbl
    /PICSxmp/age3.html  /usr/lpp/internet/server_root/wteRatings/age3.lbl
    /PICSxmp/age4.html  /usr/lpp/internet/server_root/wteRatings/age4.lbl
    /PICSxmp/cool1.html  /usr/lpp/internet/server_root/wteRatings/cool1.lbl
    /PICSxmp/cool2.html  /usr/lpp/internet/server_root/wteRatings/cool2.lbl
    /PICSxmp/cool3.html  /usr/lpp/internet/server_root/wteRatings/cool3.lbl
    /PICSxmp/cool4.html  /usr/lpp/internet/server_root/wteRatings/cool4.lbl
    /PICSxmp/cool5.html  /usr/lpp/internet/server_root/wteRatings/cool5.lbl
    /PICSxmp/graphics1.html  /usr/lpp/internet/server_root/wteRatings/graphics1.lbl
    /PICSxmp/graphics2.html  /usr/lpp/internet/server_root/wteRatings/graphics2.lbl
    /PICSxmp/graphics3.html  /usr/lpp/internet/server_root/wteRatings/graphics3.lbl
    /PICSxmp/graphics4.html  /usr/lpp/internet/server_root/wteRatings/graphics4.lbl
}
```

Figure 145. ics_pics.conf File

The above ics_pics.conf configuration file contains two types of paragraphs:

1. The `DefineLBService` paragraph lists local label files associated with our own local bureau and rating service.

   The first line of the paragraph consists of the keyword `DefineLBService`, the rating service URL, the quoted name of the rating service, the fully qualified name of the service's RAT file that describes the rating system, and an opening brace.

   The body of the paragraph lists the fully qualified names of the label files associated with this service. Each file name is introduced by the keyword `LABELFILE`.

The paragraph ends with a closing brace.

The correct syntax is summarized in the following screen:

```
DefineLBService servicename "name-of-service" ratingfile {
     LABELFILE  /path/LabelFile1 "description"
     LABELFILE  /path/LabelFile2 "description"
     ...
}
```

where:

- *servicename* is the name or URL identifier of the rating service.
- *name-of-service* is a text string representing the name of the rating service.
- *ratingfile* is the fully qualified name of the service's RAT file in the label bureau server machine file system.
- */path/LabelFile1*, */path/LabelFile2*, *...* are the fully qualified names of the label files in the label bureau server machine file system.
- *description* is a text description of the label.

2. The second paragraph in the ics_pics.conf file is `LabelsFor`. It specifies the ratings given by the rating service for documents on a given Web server.

   The first line of the paragraph consists of the keyword `LabelsFor`, the name of the Web server on which the rated documents are found, the name of the rating service and an opening brace.

   Case Sensitiveness

   We noticed that in Lotus Domino Go Webserver V4.6.2.5, the name of the Web server must match exactly for the filtering to be successful. That is, if we requested `http://WTR05178.itso.ral.ibm.com/PICSxmp/age4.html`, it will pass through as the case does not match what was defined (see the highlighted text in Figure 145 on page 216).

   As our purpose here is just to show that we can use a label bureau and how we can integrate WTE in a PICS scenario, we will not pursue this further.

   The body of the paragraph specifies labels for a set of documents. The paragraph ends with a closing brace.

   The following screen summarizes the correct syntax:

```
LabelsFor servername servicename {
    /WebPath1/document1    /path/LabelFile1
    /WebPath2/document2    /path/LabelFile2
    ...
}
```

   where:

   - *servername* is a fully qualified URL of the remote servers on which the documents being rated are found. Note that the fully qualified URL must not end with a trailing slash; thus, `http://rs600030.itso.ral.ibm.com` is acceptable as a value of the *`servername`* variable hostname on a `LabelsFor` line, but `http://rs600030.itso.ral.ibm.com/` is not.

- *`servicename`* is the fully qualified URL to which clients will send their label requests.
- *`/WebPath1/document1`*, *`/WebPath2/document2`*, *`...`* are the Web paths and names of the documents being rated.

  Notice that the path a Web client would use when requesting, for example, document1 is given by adding *`/WebPath1/document1`* at the end of the *`servername`* value, that is *`servername/WebPath1/document1`*.
- *`/path/LabelFile1`*, *`/path/LabelFile2`*, *`...`* are the fully qualified names of the label files in the label bureau machine file system.

This completes the description of the label bureau configuration.

14.3.5 WTE Proxy Server Configuration

PICS rules in a WTE proxy server can be created through the Configuration and Administration Forms, by editing the `DefinePicsRule` directives in the ibmproxy.conf configuration file, or by accessing the Common Configuration utility, provided in WebSphere Performance Pack Version 2 (see Chapter 20, “Common Configuration” on page 293).

Notice that in previous versions of WTE, the PICS filtering configuration had to be performed in a separate configuration file, called javelin.conf on AIX and Solaris, and javelin.cnf on Windows NT.

14.3.5.1 PICS Filtering Directives

In this section, we describe the basic syntax for defining PICS rules in steps by manually editing the ibmproxy.conf configuration file. If you decide to use either the Configuration and Administration Forms or the Common Configuration utility, you would find the ibmproxy.conf configuration file automatically edited according to the syntax we explain in this section.

1. **Specify the filter name**

   The filter name is specified as shown in the following screen:

```
DefinePicsRule "Filter by Age" {
      (PicsRule-1.0
      )
}
```

PicsRule-1.0

All the statements that define the PICS rule are proceeded by the keyword `PicsRule-1.0` and enclosed in parentheses. WTE V2.0 supports Version 1.0 of the PICS Rules Specifications, and not the last version, 1.1. For this reason, we recommend that you use the syntax defined in Version 1.0 of the PICS Rules Specification, since the syntax specified in Version 1.1 would not be recognized.

2. **passURL and failURL directives**

   Any pass or fail rules for specific URLs are added using the `passURL` and `failURL` directives. A `passURL` statement will always pass the specified URL and a `failURL` will always fail it. Wildcards can be used and multiple pass and fail instructions are allowed. Note that the order of these statements is very important as the checking process stops when it finds the first match. Here is an example:

   ```
   passURL ("http://w3.ibm.com")
   passURL ("http://*.lotus.com")
   failURL ("http://sphinx.sg.ibm.com")
   ```

3. **passURL for the Service Information URL**

   When you want to conditionally filter using PICS labels from a service, service information must be inserted into the rule, as we will see in Step 4. You must insert a `passURL` statement to ensure the service information URL is passed. For example:

   `passURL("http://rs600030.itso.ral.ibm.com/wteRatings")`

4. **Service Information**

   The service information is preceded by the statement `serviceinfo` and is enclosed between parentheses. Notice the following:

   - The statement `name` allows you to specify the rating service URL identifier.
   - The statement `shortname` allows you to specify a common name for the rating service. You will use this common name to refer to the rating service in the `Filter` statement, which is the last statement of the directive.
   - The `bureauURL` statement allows you to specify the label bureau URL, which is the URL used on the label bureau server to serve label requests.
   - The statement `ratfile` allows you to specify the RAT file.
   - The `available-with-content` variable informs the rule if the label is provided within the HTML `<META>` tag of the requested page (see Point 2 on page 204).

   Here is an example:

   ```
   serviceinfo (
                 name "http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html"
                 shortname "newEnzo"
                 bureauURL "http://rs600030.itso.ral.ibm.com/wteRatings"
                 ratfile "coolness.rat"
                 available-with-content "NO"
               )
   ```

5. **Filter name information**

   The next addition to the rule is the filter name information. The rule name is what the user sees when the page is failed or blocked. The description is a text description for the rule writer. The filter name and rule name must be identical, as shown in the following example:

```
name (
       rulename "Filter by Age"
       description "newEnzo filter, filters on anyday, anytime, for URLs based on age range"
     )
```

6. **Filters based on users, days and time**

   The `ibm-javelin-extensions` directive allows you to apply filters based on users, days and time:

   ```
   ibm-javelin-extensions (
                              applies-to (
                                           "pohyt@(*)"
                                         )
                              active-days "1111111"
                              start-time "00:00:01"
                              end-time "23:59:59"
   ```

   Notice that the ibm-javelin-extensions directive is not defined by the PICS Rules Specification, but is a special directive implemented in WTE. This directive is used by the WTE server administrator to better define the filtering behavior. In the example above, we are instructing the WTE caching and filtering proxy server to consider the filtering rule active all seven days of the week (`1111111`), each day from 00:00:01 to midnight 23:59:59.

7. **Filter statement**

   The RAT file in the rating service contains information on the criteria, or categories, of the content being filtered. These categories are then given allowable ranges of values.

   The PICS filter in the WTE proxy server describes the operation done on the category to be filtered in order to be passed to the requester. The operator can be: `>`, `<`, `>=`, `<=`, `!=`, `==`, `||`, `&&`, `AND`, or `OR`.

   This is the sequence of the operations that are performed:

   1. The label bureau uses the categories and rates the sites with respect to those categories. The result is a PICS label that has information on the different categories.
   2. The PICS filter in the WTE proxy server reads the label and uses the value in the label in an operation performed within the `Filter` statement. Two operations can be performed in the `Filter` statement: `Pass` or `Block`. Based on the operation specified within the `Filter` statement, the URL is then passed or failed respectively.

   The following screen shows two sample `filter` statements:

   ```
   Filter (
             Pass ("newEnzo.Age-range <= 3")
          )

   Filter (
             Block("newEnzo.Age-range >3)"
          )
   ```

Notice, that both filters in the above screen will bear the same results when a rating label on the `Age-range` category is provided. The difference in their behavior is clear when an `Age-range` rating for the page is unavailable.

In the `Pass` statement, the comparison must return `true` for the URL to be passed. Instead, in the `Block` statement, the comparison must return `true` for the URL to be blocked. Because the comparison returns `false` anyway when a Web page is not rated on the `Age-range` category, the `Pass` statement would block the URL, while the `Block` statement would pass the URL.

Putting together in the right order all the configuration lines shown in Step 1 on page 218 through Step 7 on page 220, you will get the section of the ibmproxy.conf configuration file that regulates PICS filtering. This section is shown in Figure 159 on page 231.

14.3.5.2 Using the Configuration and Administration Forms

In 14.3.5.1, “PICS Filtering Directives” on page 218, we showed how to configure PICS filtering by editing the ibmproxy.conf file. In this section, we demonstrate how the same results can be obtained by accessing the Configuration and Administration Forms. This second approach (or the Common Configuration utility, presented in Chapter 20, “Common Configuration” on page 293) is recommended because it excludes any syntax errors that might be generated when manually editing the configuration file.

First of all, you have to access the Configuration and Administration Forms, as shown below:

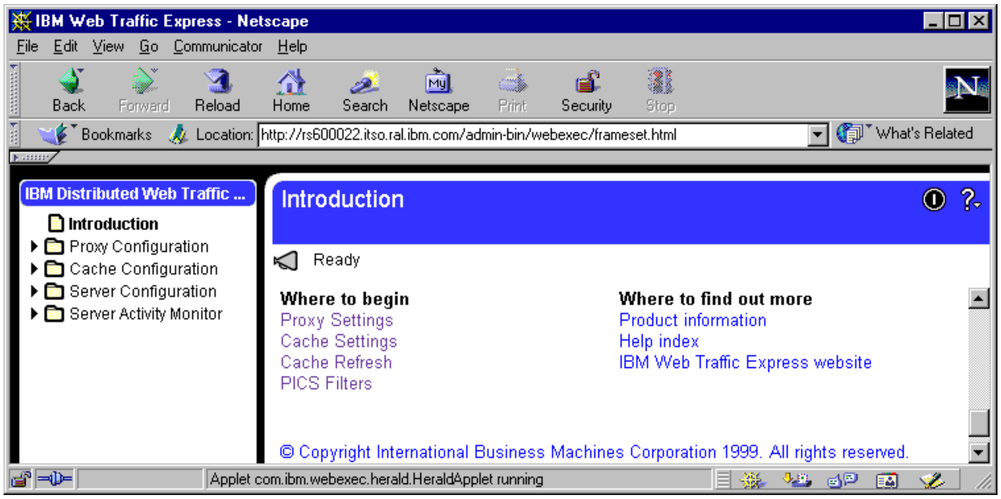

http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html

Where to begin
Proxy Settings
Cache Settings
Cache Refresh
PICS Filters

Where to find out more
Product information
Help index
IBM Web Traffic Express website

© Copyright International Business Machines Corporation 1999. All rights reserved.

Applet com.ibm.webexec.herald.HeraldApplet running

Figure 146. Configuration and Administration Forms - PICS

Then, click the **PICS Filters** link. You will be presented with the PICS Filters form, as shown in the following figure:

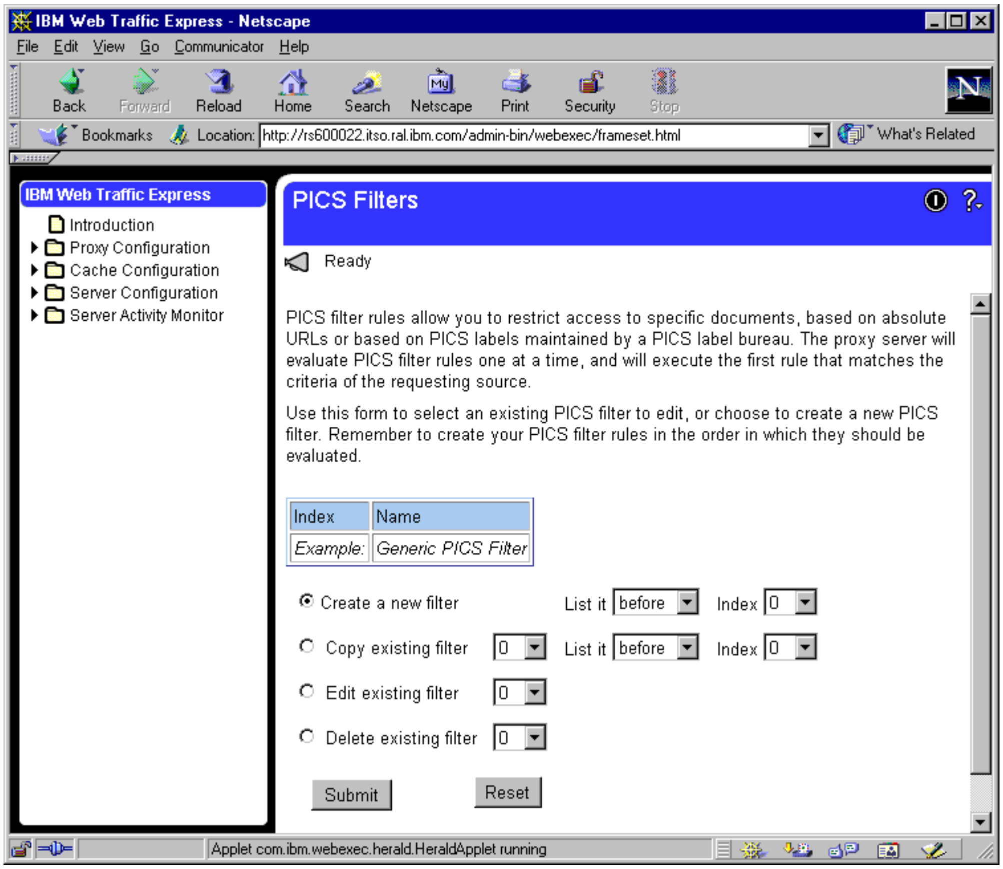
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
Ready
Applet com.ibm.webexec.herald.HeraldApplet running

PICS Filters

PICS filter rules allow you to restrict access to specific documents, based on absolute URLs or based on PICS labels maintained by a PICS label bureau. The proxy server will evaluate PICS filter rules one at a time, and will execute the first rule that matches the criteria of the requesting source.

Use this form to select an existing PICS filter to edit, or choose to create a new PICS filter. Remember to create your PICS filter rules in the order in which they should be evaluated.

Index	Name
Example:	*Generic PICS Filter*

● Create a new filter List it before Index 0

○ Copy existing filter 0 List it before Index 0

○ Edit existing filter 0

○ Delete existing filter 0

Submit Reset

Figure 147. PICS - Create a New Filter

We will begin by creating a new filter:

1. Click **Create a new filter.**
2. Click the **Submit** button.

You will be presented with the next form requesting a filter name, description, days and duration. We will name the new filter `newEnzo` and will select to perform filtering on any day at any time.

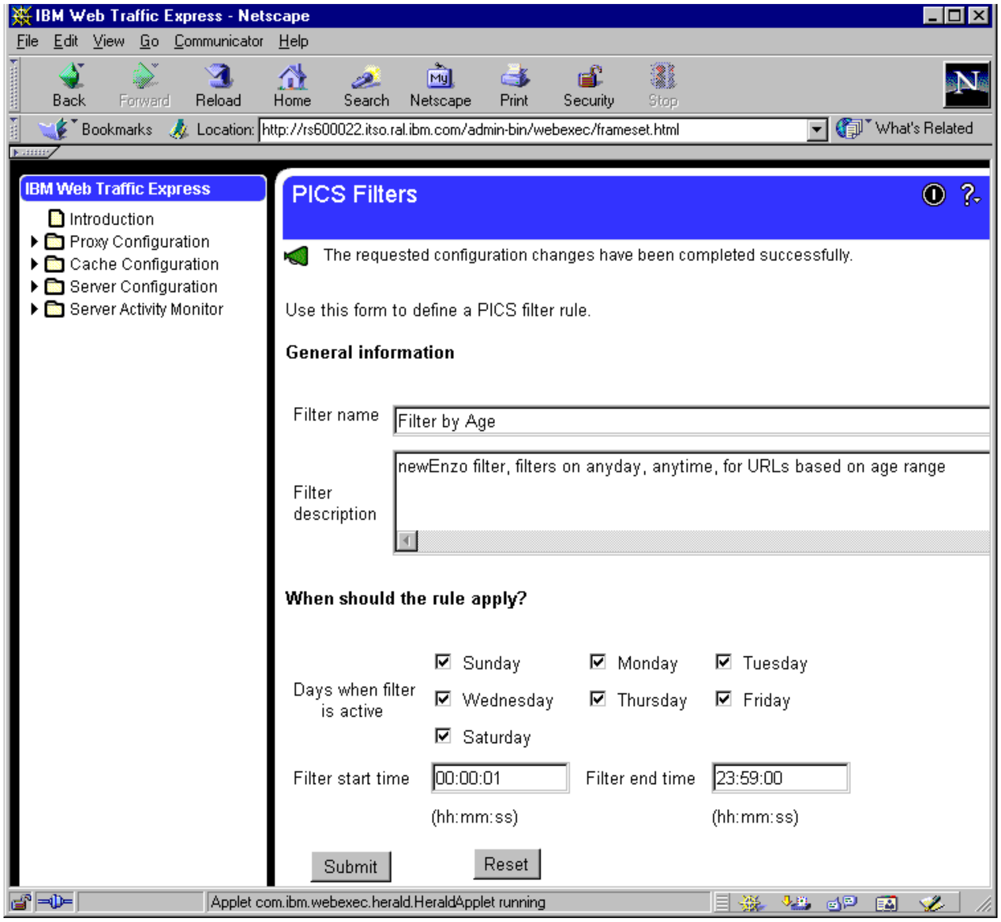
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filters
Applet com.ibm.webexec.herald.HeraldApplet running

The requested configuration changes have been completed successfully.

Use this form to define a PICS filter rule.

General information

Filter name: Filter by Age

Filter description: newEnzo filter, filters on anyday, anytime, for URLs based on age range

When should the rule apply?

Days when filter is active: ☑ Sunday ☑ Monday ☑ Tuesday ☑ Wednesday ☑ Thursday ☑ Friday ☑ Saturday

Filter start time: 00:00:01 (hh:mm:ss) Filter end time: 23:59:00 (hh:mm:ss)

Submit Reset

Figure 148. PICS - Filter Name, Date and Time

If you want to configure PICS filtering in the same way, this is the list of actions you should take:

1. Type in the new Filter name as `newEnzo`.
2. Type in a description for the purpose of this filter.
3. Click the days you want the filtering to be active.
4. Type in the duration of the filter by entering the start time and end time.
5. Click the **Submit** button.

After clicking the **Submit** button, you will be presented with a summary of the data you have entered:

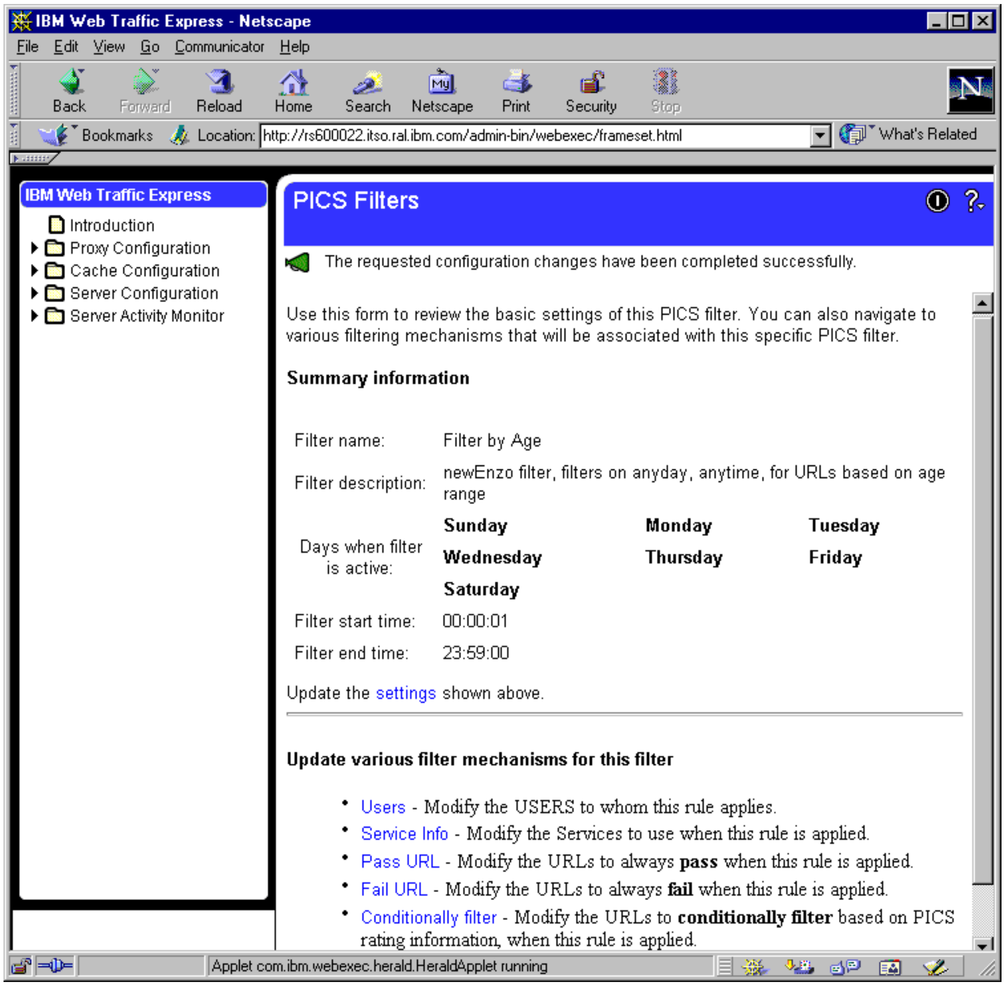
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
Applet com.ibm.webexec.herald.HeraldApplet running

PICS Filters

The requested configuration changes have been completed successfully.

Use this form to review the basic settings of this PICS filter. You can also navigate to various filtering mechanisms that will be associated with this specific PICS filter.

Summary information

Filter name:	Filter by Age		
Filter description:	newEnzo filter, filters on anyday, anytime, for URLs based on age range		
Days when filter is active:	**Sunday**	**Monday**	**Tuesday**
	Wednesday	**Thursday**	**Friday**
	Saturday		
Filter start time:	00:00:01		
Filter end time:	23:59:00		

Update the settings shown above.

Update various filter mechanisms for this filter

- Users - Modify the USERS to whom this rule applies.
- Service Info - Modify the Services to use when this rule is applied.
- Pass URL - Modify the URLs to always **pass** when this rule is applied.
- Fail URL - Modify the URLs to always **fail** when this rule is applied.
- Conditionally filter - Modify the URLs to **conditionally filter** based on PICS rating information, when this rule is applied.

Figure 149. PICS - Summary Information

We can now proceed to enter the filtering policies. We begin by clicking **Users**:

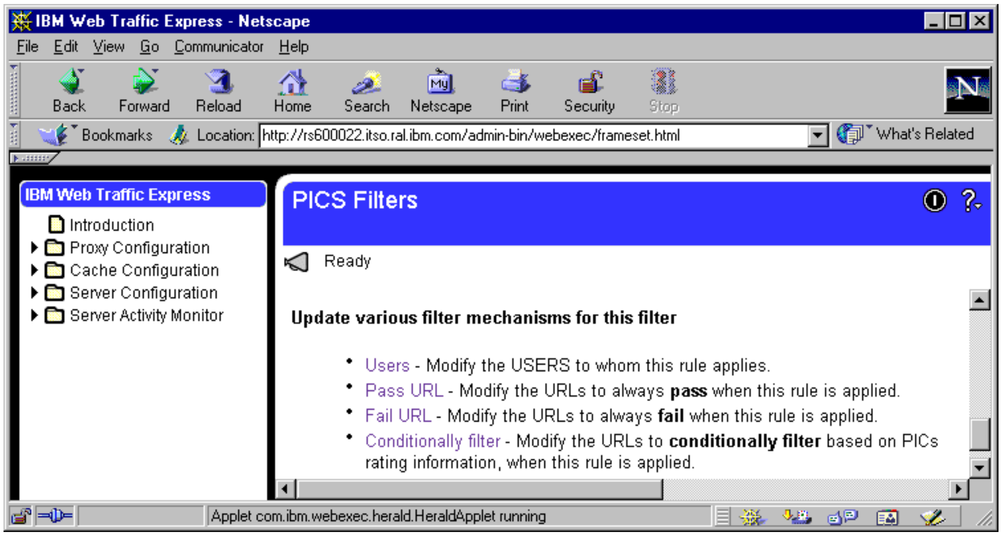
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
Applet com.ibm.webexec.herald.HeraldApplet running

PICS Filters

Ready

Update various filter mechanisms for this filter

- Users - Modify the USERS to whom this rule applies.
- Pass URL - Modify the URLs to always **pass** when this rule is applied.
- Fail URL - Modify the URLs to always **fail** when this rule is applied.
- Conditionally filter - Modify the URLs to **conditionally filter** based on PICs rating information, when this rule is applied.

Figure 150. PICS Filters Menu

The filter for users enables us to restrict to whom this filter rule applies. You may leave it empty if it applies to everybody. The following is the form for users:

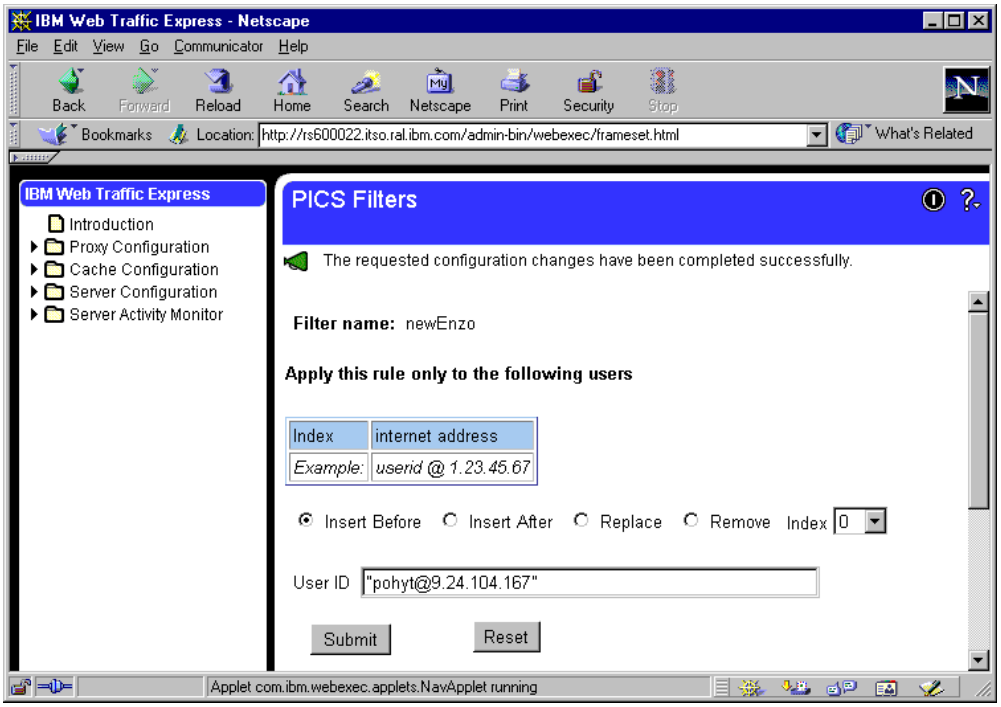
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filters
Submit
Reset
Applet com.ibm.webexec.applets.NavApplet running

The requested configuration changes have been completed successfully.

Filter name: newEnzo

Apply this rule only to the following users

Index	internet address
Example:	*userid @ 1.23.45.67*

● Insert Before ○ Insert After ○ Replace ○ Remove Index 0

User ID "pohyt@9.24.104.167"

Figure 151. PICS - Users

This is the list of actions you should take in the form above:

1. Type in the user ID that you want to apply this filter to.
2. Click the **Submit** button to add the changes to the WTE proxy configuration file, ibmproxy.conf.
3. Continue with this form if you have more users to add.

Notes

When entering the user ID, you have to include the quotes, as shown in Figure 151. Otherwise, you will get the following error message when you restart the WTE proxy server:

```
[22/Mar/1999:13:42:34 +0500] [OK] [host: ] PICS Rule starting at line 2411
could not be parsed.
The error was detected near the text "(
".
```

Moreover, some problems were reported in WebSphere Performance Pack V2 about the fact that the user ID was inserted into the wrong line in the ibmproxy.conf configuration file. Hence, if you need to filter by user, we recommend that you check the ibmproxy.conf configuration file after submitting the form above, and verify that the user ID line is inserted correctly, as shown in Step 6 on page 220. You should manually edit the file if you see that the user ID line has been inserted in the wrong line.

After this, we proceed to enter the service information. Scroll down the PICS Filters form to configure the other filter mechanisms. Then, click the **Service Info** URL link. You will be presented with the following configuration form:

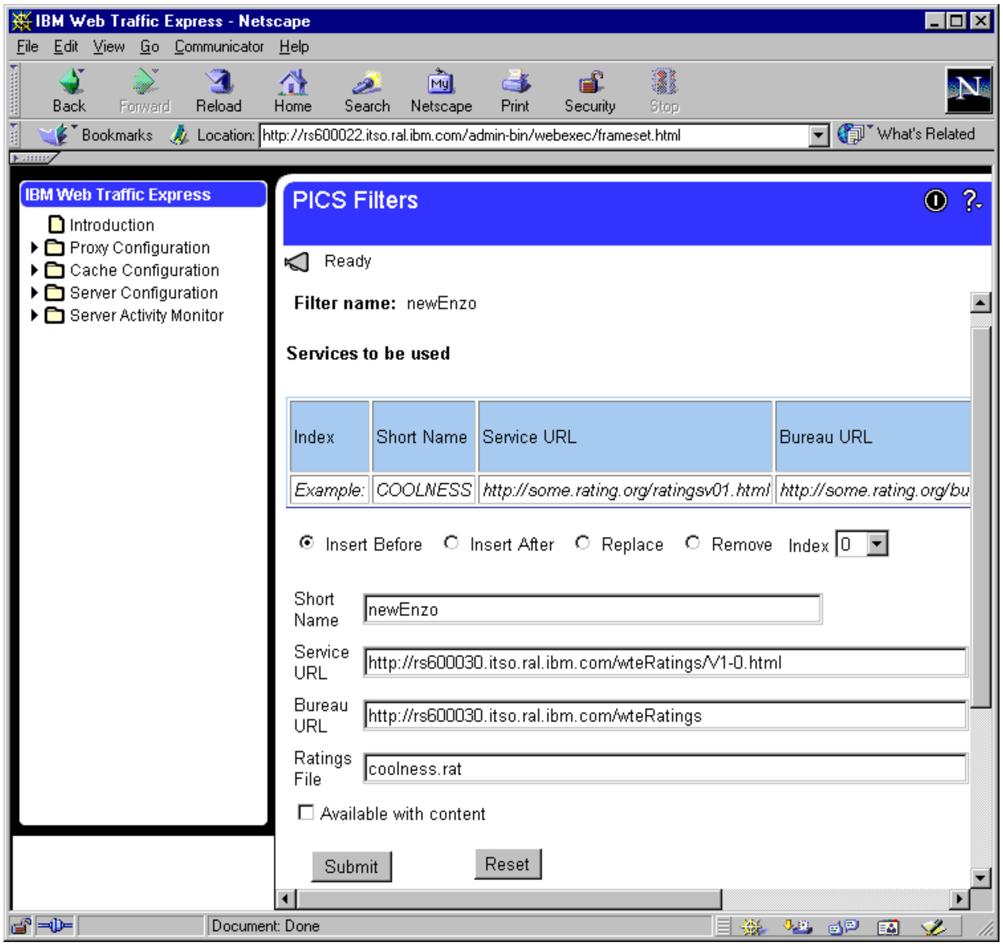
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filters
Ready
Document: Done

Filter name: newEnzo

Services to be used

Index	Short Name	Service URL	Bureau URL
Example:	*COOLNESS*	*http://some.rating.org/ratingsv01.html*	*http://some.rating.org/bu*

● Insert Before ○ Insert After ○ Replace ○ Remove Index 0

Short Name	newEnzo
Service URL	http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html
Bureau URL	http://rs600030.itso.ral.ibm.com/wteRatings
Ratings File	coolness.rat

□ Available with content

Submit Reset

Figure 152. PICS - Services to Be Used

On the form above:

1. Type in the short name for this service.

   We typed `newEnzo`.

2. Type in the rating service URL.
3. Type in the label bureau URL.
4. Type in the name of the RAT file.
5. Indicate whether the labels are available with the content.

   In our case, we did not mark the **Available with content** check box as we were using a label bureau and a rating service residing on a machine different from the Web server itself.

6. Click the **Submit** button to add the changes to the WTE proxy configuration file, ibmproxy.conf.

After completing the form above, we proceeded to enter the URL we wanted to specifically pass. To do this, you should scroll down the PICS Filters form and select **Pass URL** from the PICS Filters menu, as shown in the following figure:

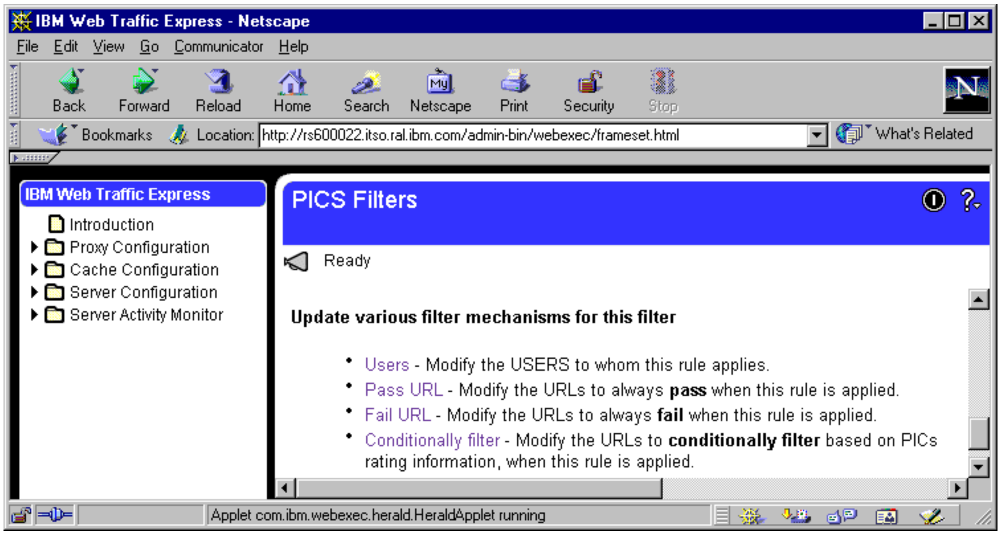

IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filters
Ready
Applet com.ibm.webexec.herald.HeraldApplet running

Update various filter mechanisms for this filter

- Users - Modify the USERS to whom this rule applies.
- Pass URL - Modify the URLs to always **pass** when this rule is applied.
- Fail URL - Modify the URLs to always **fail** when this rule is applied.
- Conditionally filter - Modify the URLs to **conditionally filter** based on PICs rating information, when this rule is applied.

Figure 153. PICS - Update Other Filters

You will be presented with the following configuration form:

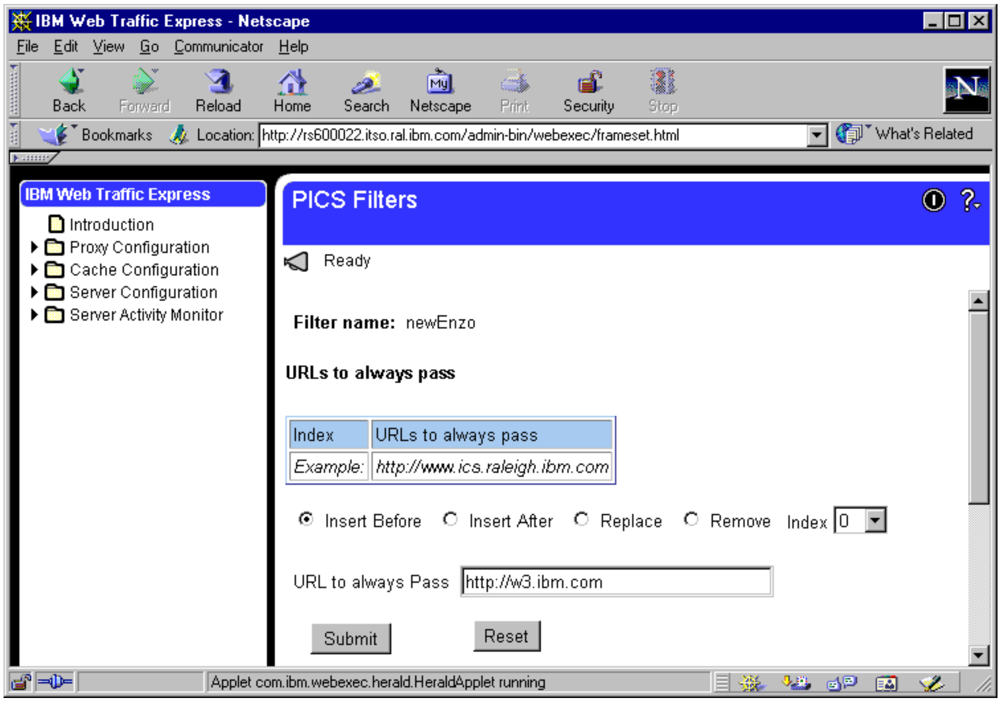

IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filters
Ready
Applet com.ibm.webexec.herald.HeraldApplet running

Filter name: newEnzo

URLs to always pass

Index	URLs to always pass
Example:	*http://www.ics.raleigh.ibm.com*

● Insert Before ○ Insert After ○ Replace ○ Remove Index 0

URL to always Pass http://w3.ibm.com

Submit Reset

Figure 154. PICS - URLs to always Pass

On the form above, perform the following sequence of actions:

1. Enter the URL that you want to pass in this filter.

   We entered `http://w3.ibm.com`.

2. Click the **Submit** button to add the changes to the WTE proxy server configuration file, ibmproxy.conf.

Notice that you can use the form above to specifically pass as many URLs as you want.

After completing the form above, we proceeded to enter the URL we wanted to specifically fail. Scroll down to the bottom of the PICS Filters form, and click **Fail URL**, as shown:

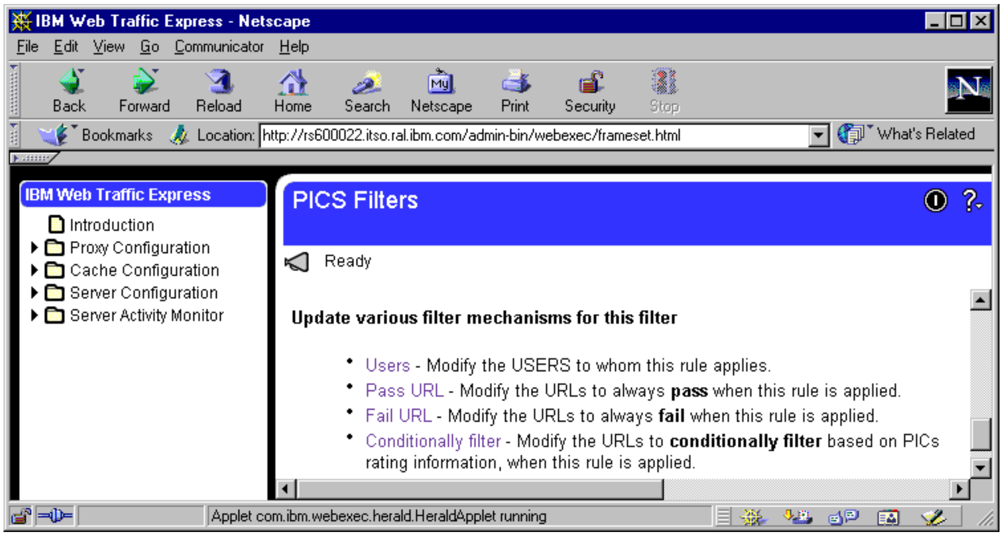
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filters
Ready
Applet com.ibm.webexec.herald.HeraldApplet running

Update various filter mechanisms for this filter

- Users - Modify the USERS to whom this rule applies.
- Pass URL - Modify the URLs to always **pass** when this rule is applied.
- Fail URL - Modify the URLs to always **fail** when this rule is applied.
- Conditionally filter - Modify the URLs to **conditionally filter** based on PICs rating information, when this rule is applied.

Figure 155. PICS - Update Other Filters

You will be presented with the following configuration form, which you can use to specifically block one or more URLs:

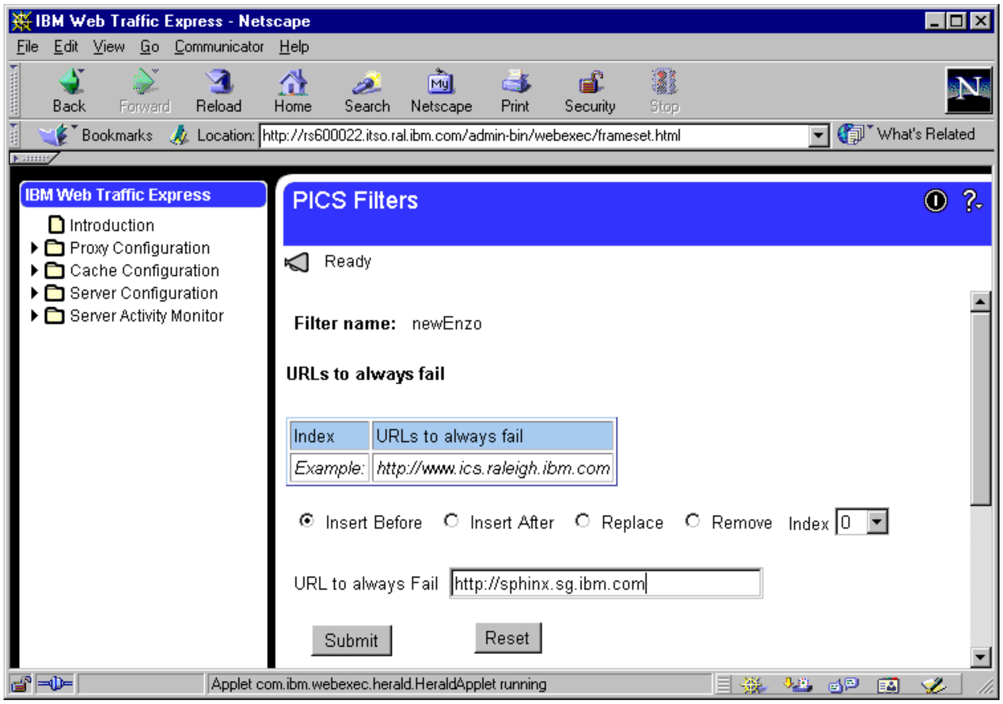
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filters
Ready

Filter name: newEnzo

URLs to always fail

Index	URLs to always fail
Example:	*http://www.ics.raleigh.ibm.com*

● Insert Before ○ Insert After ○ Replace ○ Remove Index 0

URL to always Fail http://sphinx.sg.ibm.com

Submit
Reset
Applet com.ibm.webexec.herald.HeraldApplet running

Figure 156. PICS - URLs to always Fail

On the form above:

1. Enter the URL which you want to fail in this filter.

   We entered `http://sphinx.sg.ibm.com`.

2. Click the **Submit** button to add the changes to the WTE proxy server configuration file.

After completing the form above, we proceeded to enter the conditions for this filter based on the PICS rating information. Scroll down to the bottom of the PICS Filters form and select **Conditionally filter** from the PICS Filters menu, as shown in the following figure:

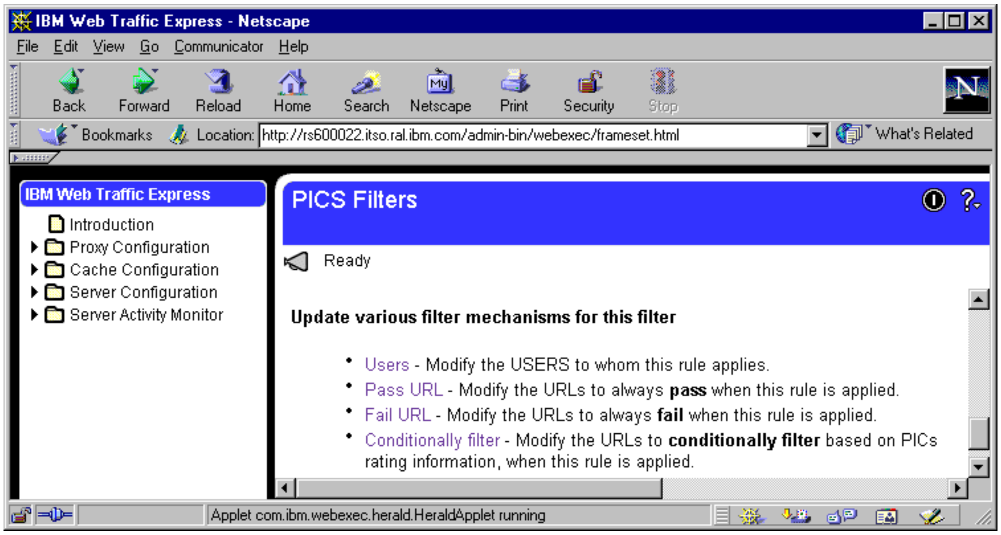
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filters
Ready
Applet com.ibm.webexec.herald.HeraldApplet running

Update various filter mechanisms for this filter

- Users - Modify the USERS to whom this rule applies.
- Pass URL - Modify the URLs to always **pass** when this rule is applied.
- Fail URL - Modify the URLs to always **fail** when this rule is applied.
- Conditionally filter - Modify the URLs to **conditionally filter** based on PICs rating information, when this rule is applied.

Figure 157. PICS - Update Other Filters

You will be presented with the following configuration form:

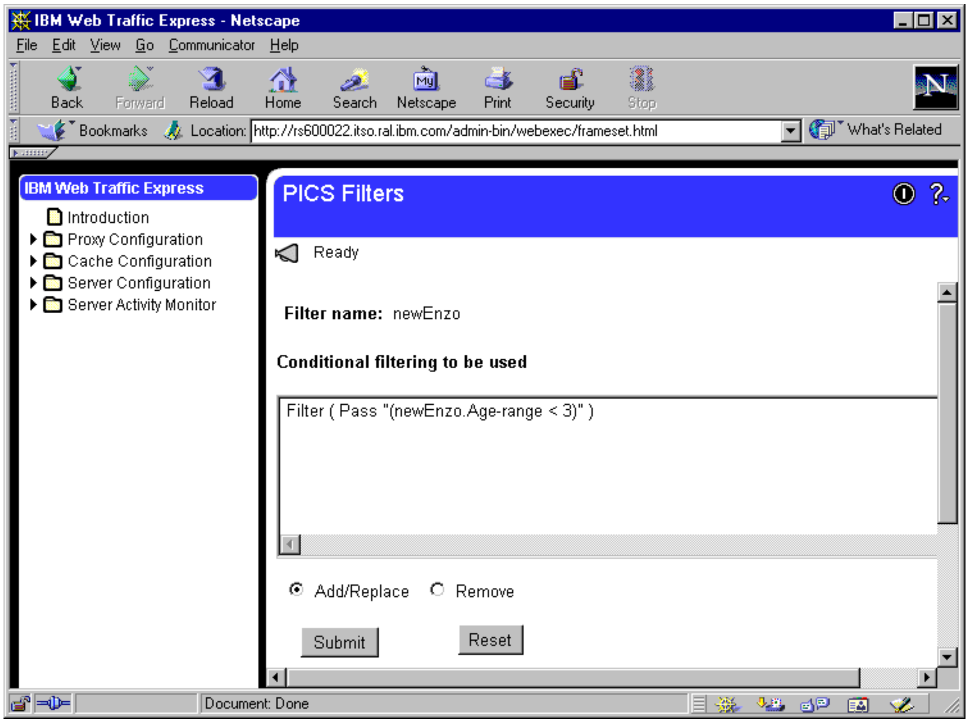
IBM Web Traffic Express - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
What's Related
IBM Web Traffic Express
Introduction
Proxy Configuration
Cache Configuration
Server Configuration
Server Activity Monitor
PICS Filters
Ready
Submit
Reset
Document: Done

Filter name: newEnzo

Conditional filtering to be used

```
Filter ( Pass "(newEnzo.Age-range < 3)" )
```

● Add/Replace ○ Remove

Figure 158. PICS - Conditional Filtering

On the form above:

1. Enter the conditions based on the PICS rating information, as shown in Figure 158.

Conditional Filter Format

Note that if the format entered is wrong, the conditional filter will not get updated into the WTE proxy server configuration file, ibmproxy.conf, and no errors will be reported.

2. Click the **Submit** button to add the changes to the WTE proxy server configuration file.

Finally, we restart the WTE server to activate the changes. Figure 159 shows how the WTE proxy server configuration file gets automatically updated after submitting the forms above.

14.3.5.3 A Look at the WTE Configuration File

After making and activating all the changes (either manually or automatically, using the Configuration and Administration Forms or the Common Configuration utility), the following are the lines inserted into the ibmproxy.conf file:

```
# ===================================================================== #
#
#        PICS Filtering directives
#
# ===================================================================== #
DefinePicsRule "Filter by Age" {
     (PicsRule-1.0
          (
          passURL ( "http://w3.ibm.com" )
          failURL ( "http://sphinx.sg.ibm.com" )
          passURL ( "http://rs600030.itso.ral.ibm.com/wteRatings")
          serviceinfo (
                      shortname "newEnzo"
                      name "http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html"
                      bureau "http://rs600030.itso.ral.ibm.com/wteRatings"
                      ratfile "coolness.rat"
                      available-with-content "NO"
                      )
          name     (
              rulename "Filter by Age"
             description "newEnzo filter, filters on anyday, anytime, for URLs based on age range"
                   )
          ibm-javelin-extensions (
               applies-to (
                           "pohyt@(*)"
                          )
               active-days "1111111"
               start-time "00:00:01"
               end-time "23:59:59"
                                )
           Filter ( Pass "(newEnzo.Age-range < 3)" )
         )
    )
}
```

Figure 159. Addition to the ibmproxy.conf File

14.3.6 Working with PICS Filtering at the Proxy Server Level

Adopting the configuration we have done so far, we performed some tests. Before proceeding, we stopped and restarted the Lotus Domino Go Webserver on the two machines wtr05178 and rs600030 where it was running. We also stopped and restarted the WTE server, rs600022.

On the Web client machine, we configured the Web browser to access the Internet through the proxy server rs600022.itso.ral.ibm.com. Then, we requested the URL `http://sphinx.sg.ibm.com/`. The filter failed this page, since we had explicitly blocked it in the PICS filtering configuration. The following is the error message sent to the client by the WTE server:

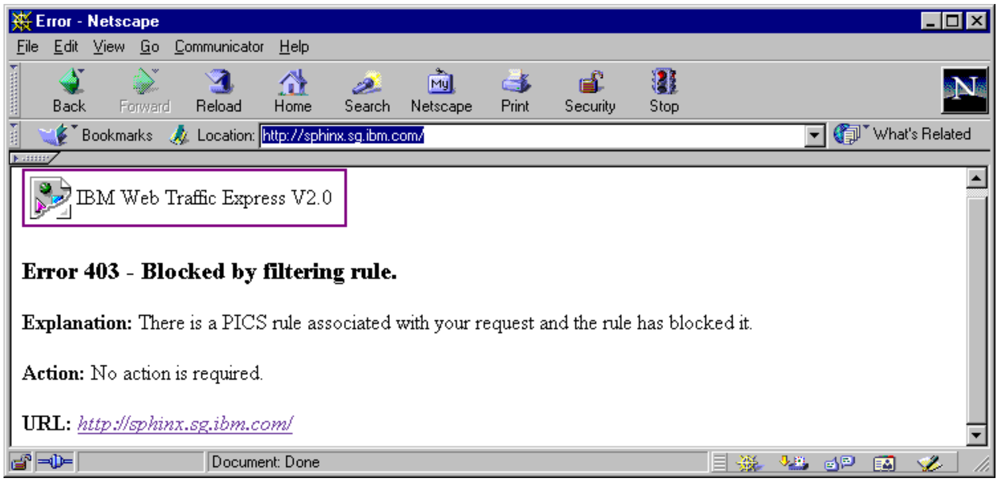
Error - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://sphinx.sg.ibm.com/
What's Related
Document: Done

IBM Web Traffic Express V2.0

Error 403 - Blocked by filtering rule.

Explanation: There is a PICS rule associated with your request and the rule has blocked it.

Action: No action is required.

URL: *http://sphinx.sg.ibm.com/*

Figure 160. failURL - http://sphinx.sg.ibm.com

Then we requested the page `http://w3.ibm.com`, which we had explicitly permitted, and this was passed by the WTE server, as shown in the following figure:

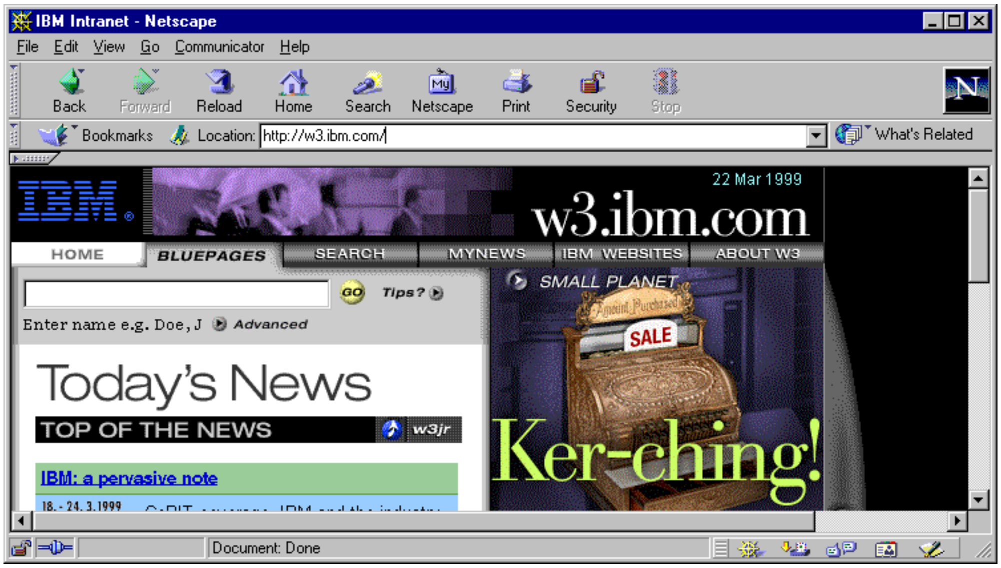

Figure 161. passURL - http://w3.ibm.com

We then tested the way the WTE server interacts with the label bureau and rating service machine. We did this by requesting the pages age1.html, age2.html, age3.html, and age4.html from the Web server wtr05178.itso.ibm.com. The following is the reset for the four requests:

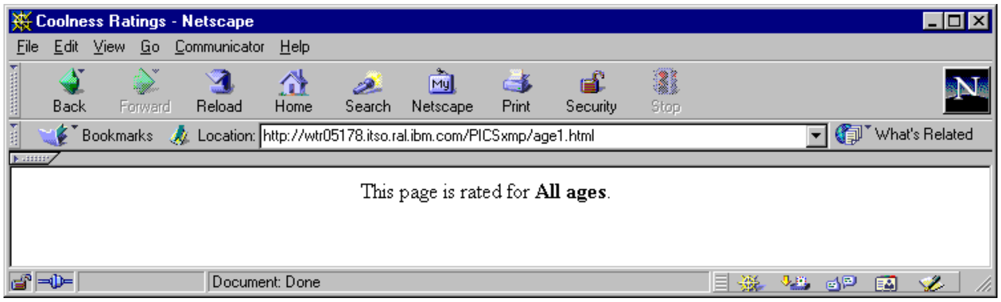
Coolness Ratings - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/PICSxmp/age1.html
What's Related

This page is rated for **All ages**.

Document: Done

Figure 162. Pass - http://wtr05178.itso.ral.ibm.com/PICSxmp.age1.html

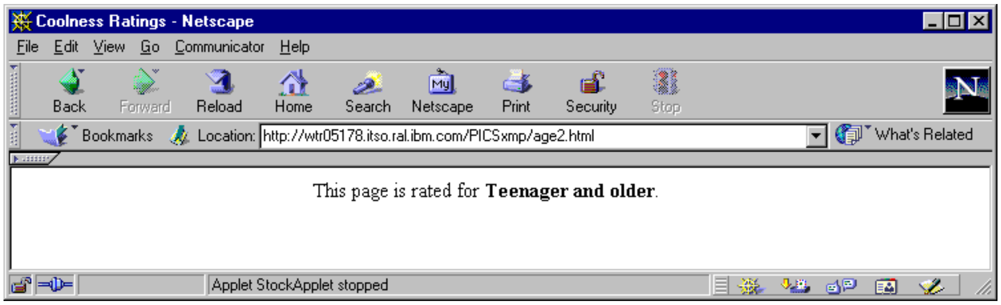
Coolness Ratings - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/PICSxmp/age2.html
What's Related

This page is rated for **Teenager and older**.

Applet StockApplet stopped

Figure 163. Pass - http://wtr05178.itso.ral.ibm.com/PICSxmp/age2.html

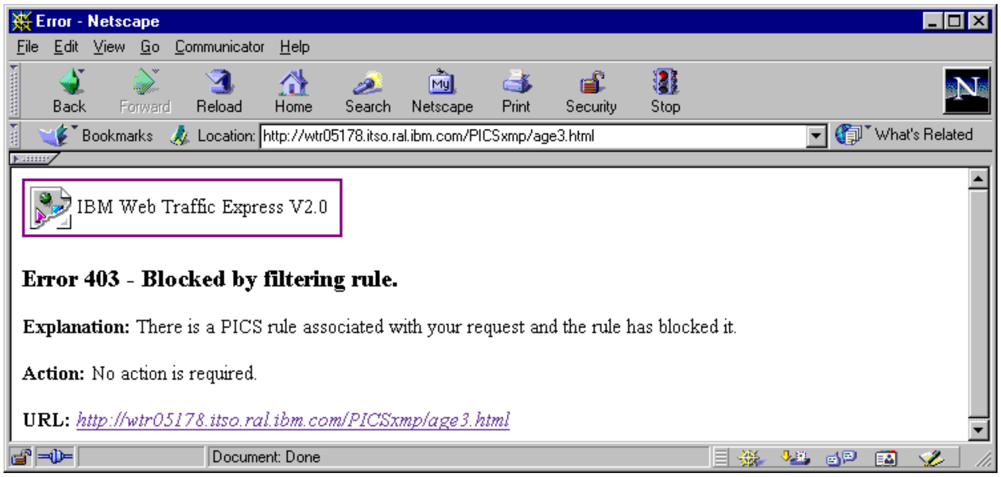
Error - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/PICSxmp/age3.html
What's Related

IBM Web Traffic Express V2.0

Error 403 - Blocked by filtering rule.

Explanation: There is a PICS rule associated with your request and the rule has blocked it.

Action: No action is required.

URL: http://wtr05178.itso.ral.ibm.com/PICSxmp/age3.html

Document: Done

Figure 164. Block - http://wtr05178.itso.ral.ibm.com/PICSxmp/age3.html

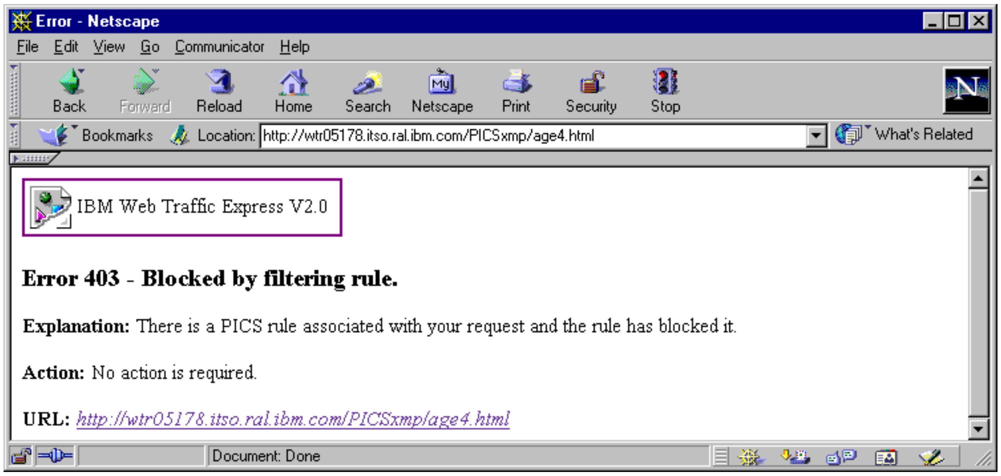
Error - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/PICSxmp/age4.html
What's Related
IBM Web Traffic Express V2.0
Document: Done

Error 403 - Blocked by filtering rule.

Explanation: There is a PICS rule associated with your request and the rule has blocked it.

Action: No action is required.

URL: http://wtr05178.itso.ral.ibm.com/PICSxmp/age4.html

Figure 165. Blocked - http://wtr05178.itso.ral.ibm.com/PICSxmp/age4.html

As we had set the filter rule as `Pass` for `Age-range` less than 3, we only expected age1.html and age2.html to pass. Hence, the above examples demonstrate that our configuration was successful.

14.3.7 Embedding the PICS Labels in the Web Documents

The last experience we wanted to try was to eliminate the label bureau server from the PICS architecture, as shown in the following diagram:

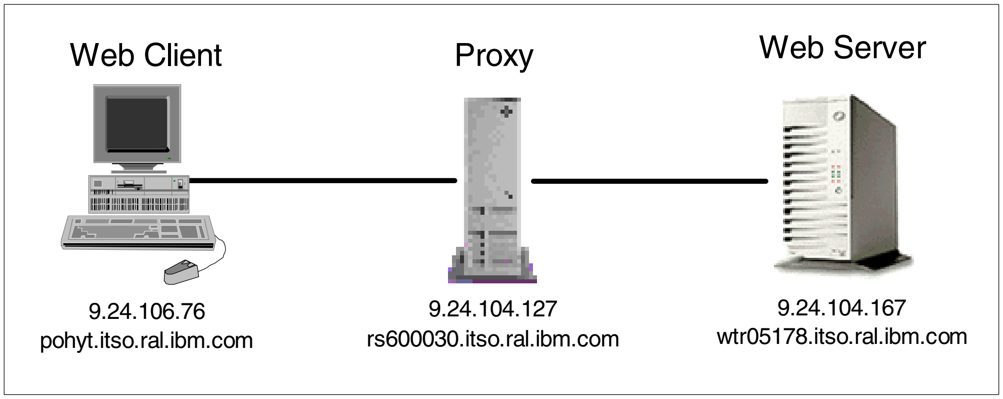

Figure 166. PICS Environment without the Label Bureau Server

We have mentioned that it is possible, and even preferable, that when the proxy server forwards the client's request to the Web server, it is the Web server itself that provides the PICS labels to the proxy (see Point 1 on page 203 and Point 2 on page 204). In this way, only one connection is required, and the user on the client machine will not experience a long wait.

This operation can be performed on the server-side dynamically, by embedding the PICS labels in the HTTP headers of the Web documents while they are being sent back to the proxy, or statically, by manually embedding the PICS labels in the HTML headers of the Web documents when these are created. We implemented the solution to embed the PICS labels in the HTML headers.

A Web server administrator can accomplish this configuration by including the PICS labels in the <META> tags of the HTML document header. This method has three heavy drawbacks. First, if you use this method, you will be able to send labels only with HTML documents, not with images, video, or anything else. Second, this process is entirely manual and therefore time-consuming, error-prone, and difficult to maintain. Third, as we have already said, it does not incorporate any of the security mechanisms (message digest, digital signature, etc.) that would guarantee the validity of the labels, if this is a requirement specified on the client-side. So, the static choice, in general, is not recommended, and the dynamic choice is certainly preferable. However, we only wanted to see what changes must be done on the WTE server configuration when this different approach is adopted, so we implemented this solution.

Notice that for the above to work, it is necessary to specify in the ibmproxy.conf configuration file of the WTE server that the PICS labels are embedded in the HTML document. To do this, the directive `available-with-content` must be set to `"YES"` (see Step 4 on page 219 and Figure 159 on page 231). You can set the value of the `available-with-content` directive to `"YES"` by either manually editing the ibmproxy.conf configuration file or checking the **Available with content** check box in the PICS Filters configuration form (see Figure 152 on page 226).

To implement this PICS scenario, we first prepared an HTML file, called testage4.html, and copied it into the directory \WWW\HTML\PICSxmp of the Web server machine wtr05178. This file had to contain the PICS label in the <META> tag of its HTML header. The file testage4.html is shown in the following figure:

```
<HTML>
<HEAD>
<META http-equiv="PICS-Label" content=
(PICS-1.1 "http://rs600030.itso.ral.ibm.com /wteRatings/V1-0.html"
labels on "1994.11.05T08:15-0500"
for "http://wtr05178.itso.ral.ibm.com/PICSxmp/testage4.html"
ratings (Coolness 0 Age-range 2 Graphics 0)) >
</HEAD>

<BODY>
<CENTER>
<H2>The label for this page is embedded in the META tag of the header</H2>
This page is rated for <B>No one</B>.
</CENTER>
</BODY>
</HTML>
```

Figure 167. testage4.html

As you can see from the string highlighted in the figure above, we first set the category `Age-range` for this Web page to `2`.

This HTML page passed the WTE PICS filters we set earlier and displayed the page successfully as shown:

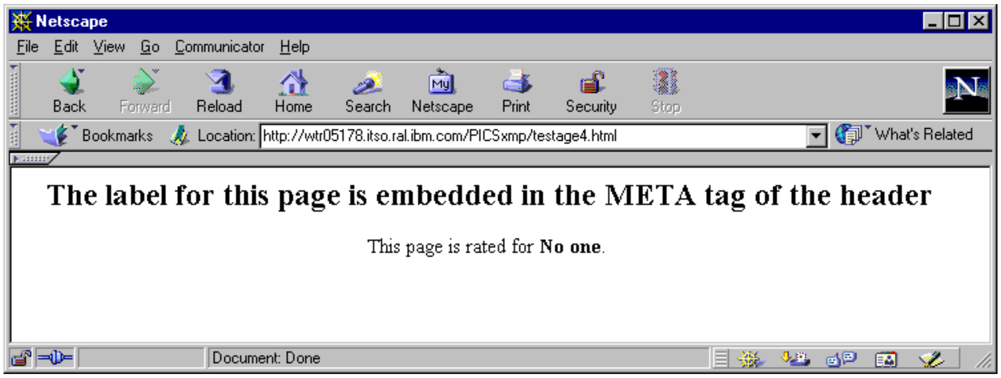
Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/PICSxmp/testage4.html
What's Related
Document: Done

The label for this page is embedded in the META tag of the header

This page is rated for **No one**.

Figure 168. testage4.html Passed the WTE PICS Filter

We then modified the HTML code shown in Figure 167 on page 235 and set the `Age-range` to `4`. As per the filter rule, it failed:

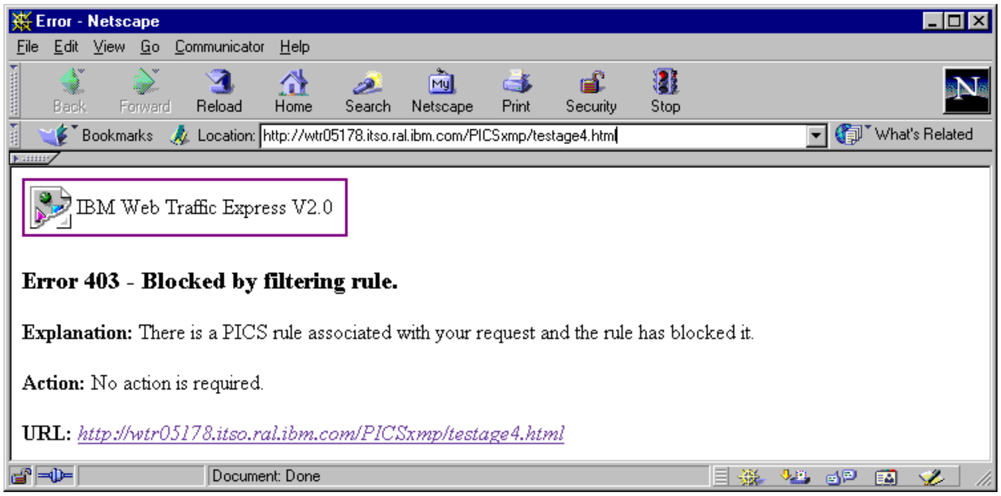
Error - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/PICSxmp/testage4.html
What's Related
IBM Web Traffic Express V2.0
Document: Done

Error 403 - Blocked by filtering rule.

Explanation: There is a PICS rule associated with your request and the rule has blocked it.

Action: No action is required.

URL: *http://wtr05178.itso.ral.ibm.com/PICSxmp/testage4.html*

Figure 169. testage4.html Failed the WTE PICS Filter

This completes the scenario on PICS filtering.

Note

When changes are made to the ibmproxy.conf file, retest a scenario only after you have cleared memory cache and disk cache in the browser and also the cache on the WTE proxy server. This is to prevent getting cached data.

Chapter 15. WTE Cancel Control

Cancel control is a new function, implemented in IBM Web Traffic Express (WTE) V2, the caching and filtering component of IBM WebSphere Performance Pack Version 2. This function lets the administrator set a parameter that allows the proxy server to cache certain documents even if the user disconnects or selects **Stop** before the entire document has been retrieved. The parameter value represents a percentage of the file being received. If the amount of the file already received is equal to or greater than the value, the whole file will be received and stored in the cache regardless of the fact that the user disconnected or clicked **Stop**.

15.1 How to Configure Cancel Control

Cancel control is defined using the `ContinueCaching` directive, which specifies the point at which a file being received from a content server will continue to be received from the content server and stored in the cache even if the connection to the client which requested the file has been terminated. The value specified represents a percentage of the size of the file being transferred. If less than this percentage of the file has been transferred from the content server at the time the client connection is terminated, file transfer from the content server will be terminated and the partial cache file will be removed from the cache.

The syntax for the `ContinueCaching` directive is:

```
ContinueCaching percentage
```

For example, if:

```
ContinueCaching 75
```

is specified, WTE will continue transferring the file from the content server and generate the cache file provided that at least 75% of the file has already been transferred before WTE detects the client connection has terminated.

15.2 Cancel Control Pros and Cons

In the following list, we explain the effects of turning cancel control on or off, and why and when you should or should not use it:

1. Cancel control makes very little difference for small objects – and the vast majority of objects moving through a proxy are small.
2. Where cancel control does matter is for large objects – documents over 100 KB (plus or minus, depending on how much network bandwidth you have; if your clients have high bandwidth connections, this number is higher). Enabling cancel control will generally allow more large objects to enter the cache.
3. Cancel control makes a trade-off of cache space vs. bandwidth savings. The intent is that if WTE has already loaded a significant part of a document, then it might as well continue to load the rest of it, so the cached file can be used when another client requests the document.
4. If your cache is limited, or you are using the caching algorithm that maximizes the cache to improve user response time (see Point 1 on page 38), we would

not suggest using cancel control (or if you do, set it to a high value, such as 90%). Enabling cancel control does not make much sense when using this type of caching algorithm as it gives preference to small files, and tends to remove larger files.

5. If you have a significant amount of cache (for example, if your cache is not running close to capacity), and are using the caching algorithm to minimize network bandwidth (see Point 2 on page 38), we would suggest turning on cancel control.
6. When using cancel control, it usually makes sense for the setting to be above 50%. While cancel control will work correctly when set to a low value, the bandwidth savings are reduced.

15.3 Cancel Control Scenario

In this section, we will examine the cancel control function using the following components:

Table 14. Cancel Control Scenario - Hardware, Software and Network Configuration

Host Name	IP Address	Operating System	Service
wtr05178.itso.ral.ibm.com	9.24.104.167	Windows NT Server 4.0	**Web Client** Netscape Navigator V4.5
rs600022.itso.ral.ibm.com	9.24.104.127	AIX 4.3.1	**Caching Proxy Server** Web Traffic Express V2.0
w3.itso.ibm.com	9.12.14.150	AIX 4.3.1	**Web Server** IBM HTTP Server

To enable cancel control on the WTE proxy server rs600022, we manually edit the WTE proxy configuration file, ibmproxy.conf. You will find the directive `ContinueCaching` commented out if it has never been configured before. We uncommented it and assigned to it a percent value of `20` as shown:

```
ContinueCaching 20
```

This means that if the WTE has started caching the file requested and has already cached more than 20%, it will continue to complete caching the file even if the user terminates the request.

As we said in Point 6 on page 238, a good value for cancel control should be equal to or greater than 50%. Here, we set it to 20% for testing purposes, in order to make it easier to see what the effects of cancel control are.

After setting the `ContinueCaching` directive, we have to restart the WTE server.

On the WTE proxy server, we reset WTE by removing its cache and all its logs by performing the following steps:

1. Remove the cache and all its logs with the following AIX commands:

   ```
   rm -r /wte/cache
   rm -r /wte/logs
   ```

2. Restart the WTE server.

On the Web client machine, we perform the following operations:

1. Configure the browser to use the proxy server `http://rs600022.itso.ral.ibm.com.`
2. Point the browser to the URL `http://w3.itso.ibm.com`.
3. Click the **RS/6000 Web Server - Powered by Apache, SG24-5132-00** URL link to download the mentioned redbook, as shown:

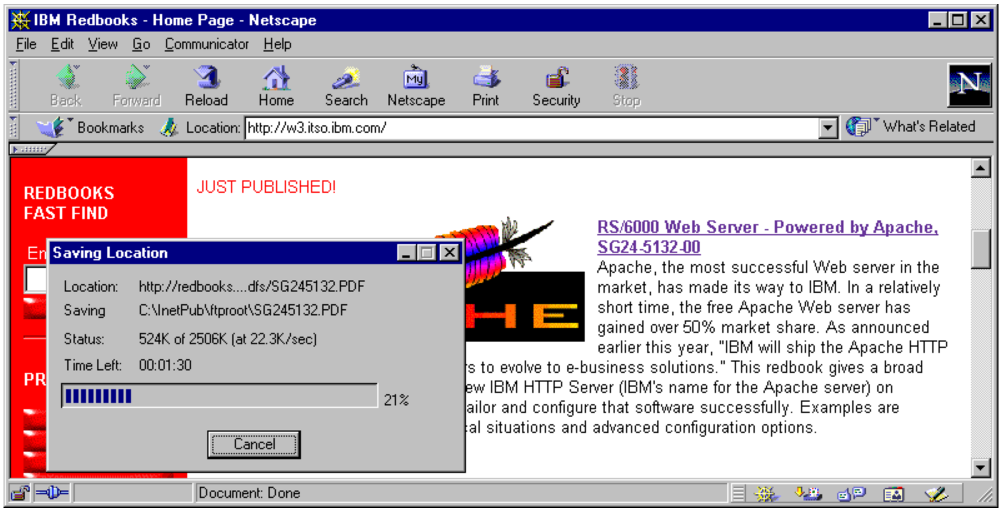

JUST PUBLISHED!

RS/6000 Web Server - Powered by Apache, SG24-5132-00
Apache, the most successful Web server in the market, has made its way to IBM. In a relatively short time, the free Apache Web server has gained over 50% market share. As announced earlier this year, "IBM will ship the Apache HTTP
s to evolve to e-business solutions." This redbook gives a broad
ew IBM HTTP Server (IBM's name for the Apache server) on
ailor and configure that software successfully. Examples are
al situations and advanced configuration options.

Saving Location

Location: http://redbooks....dfs/SG245132.PDF
Saving C:\InetPub\ftproot\SG245132.PDF
Status: 524K of 2506K (at 22.3K/sec)
Time Left: 00:01:30
21%

Figure 170. Cancel Control - Cancel at 21% while WTE Is Caching

4. When the saving was above 20%, we terminated the download by clicking the **Cancel** button, as shown in the figure above.

What happens to the WTE server? It continues its caching activity. This can be shown by listing the /wte/cache directory using the following AIX command:

```
ls -al -Ri /wte/cache
```

This is the output of the command immediately after stopping the downloading on the client machine:

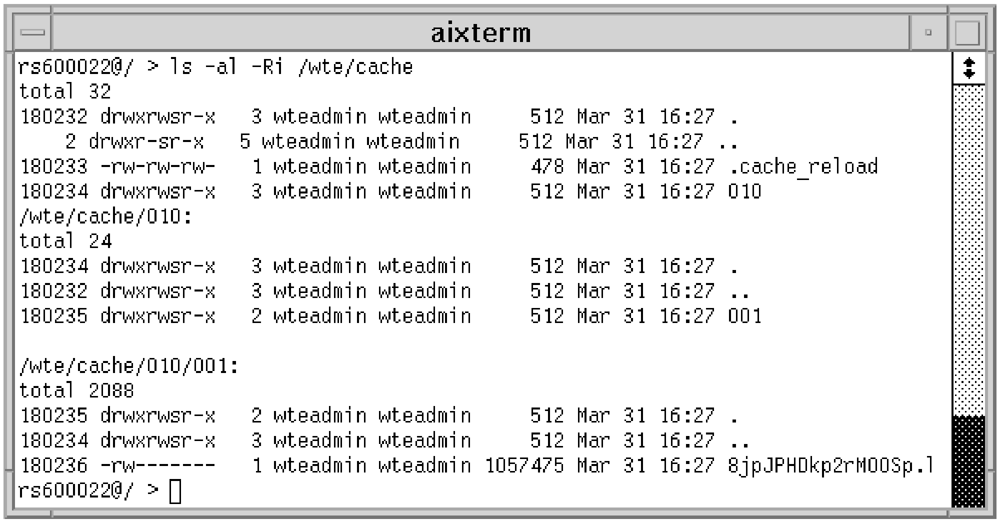

```
rs600022@/ > ls -al -Ri /wte/cache
total 32
180232 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 .
     2 drwxr-sr-x   5 wteadmin wteadmin     512 Mar 31 16:27 ..
180233 -rw-rw-rw-   1 wteadmin wteadmin     478 Mar 31 16:27 .cache_reload
180234 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 010
/wte/cache/010:
total 24
180234 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 .
180232 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 ..
180235 drwxrwsr-x   2 wteadmin wteadmin     512 Mar 31 16:27 001

/wte/cache/010/001:
total 2088
180235 drwxrwsr-x   2 wteadmin wteadmin     512 Mar 31 16:27 .
180234 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 ..
180236 -rw-------   1 wteadmin wteadmin 1057475 Mar 31 16:27 8jpJPHDkp2rMOOSp.l
rs600022@/ > 
```

Figure 171. Cancel Control - Listing of the Cached Files while WTE Is Caching 1

We executed the command again one minute later, and we noticed that the file 8jpJPHDkp2rMOOSp.1 had grown from 1057475 bytes to 2567366 bytes, as shown:

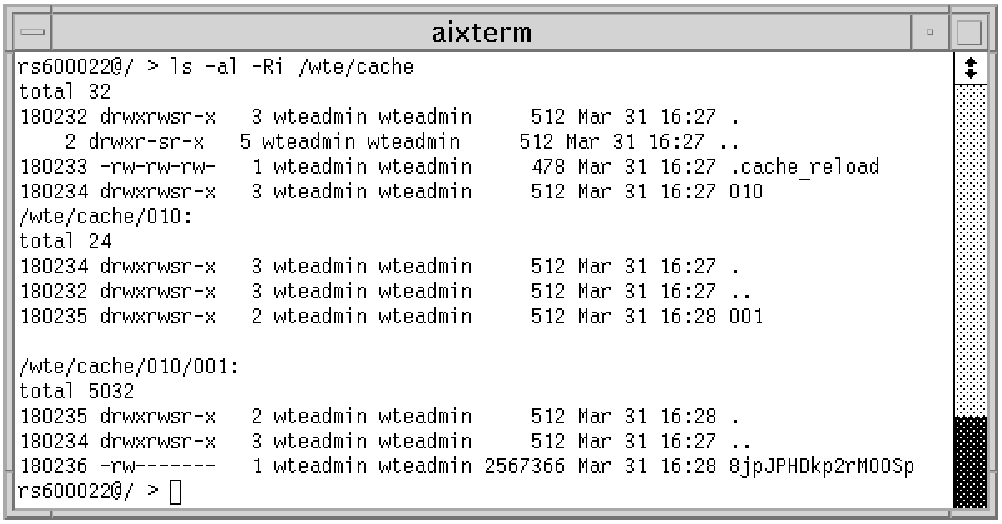

```
rs600022@/ > ls -al -Ri /wte/cache
total 32
180232 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 .
    2 drwxr-sr-x   5 wteadmin wteadmin     512 Mar 31 16:27 ..
180233 -rw-rw-rw-   1 wteadmin wteadmin     478 Mar 31 16:27 .cache_reload
180234 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 010
/wte/cache/010:
total 24
180234 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 .
180232 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 ..
180235 drwxrwsr-x   2 wteadmin wteadmin     512 Mar 31 16:28 001

/wte/cache/010/001:
total 5032
180235 drwxrwsr-x   2 wteadmin wteadmin     512 Mar 31 16:28 .
180234 drwxrwsr-x   3 wteadmin wteadmin     512 Mar 31 16:27 ..
180236 -rw-------   1 wteadmin wteadmin 2567366 Mar 31 16:28 8jpJPHDkp2rMOOSp
rs600022@/ > 
```

Figure 172. Cancel Control - Listing of the Cached Files while WTE Is Caching 2

From the browser, we again retrieved the same file, as we had already done above. Downloading now is very much faster. Looking at the WTE proxy log file, we noted that there had been only one access to the Web site where the file was located, even after the two requests were done:

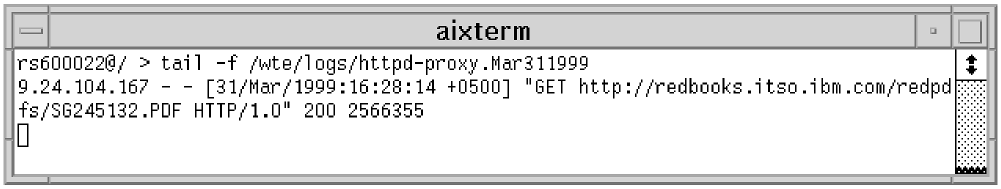

```
rs600022@/ > tail -f /wte/logs/httpd-proxy.Mar311999
9.24.104.167 - - [31/Mar/1999:16:28:14 +0500] "GET http://redbooks.itso.ibm.com/redpd
fs/SG245132.PDF HTTP/1.0" 200 2566355
```

Figure 173. Cancel Control - the Proxy Log File

Moreover, looking at the WTE cache log file, we noticed that the client, having IP address 9.24.104.167, had been served from the cache for the second request, as shown:

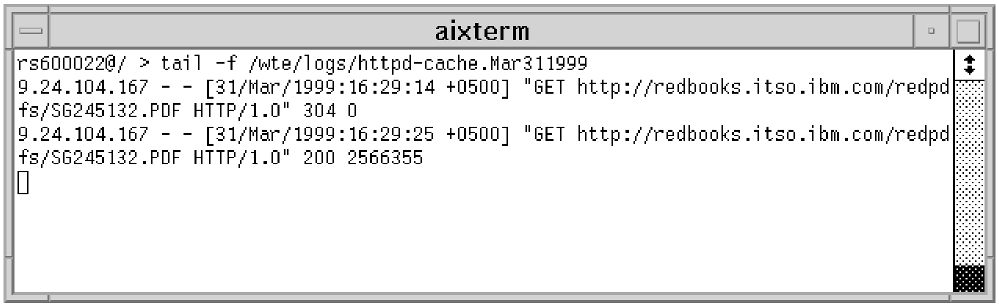

```
rs600022@/ > tail -f /wte/logs/httpd-cache.Mar311999
9.24.104.167 - - [31/Mar/1999:16:29:14 +0500] "GET http://redbooks.itso.ibm.com/redpd
fs/SG245132.PDF HTTP/1.0" 304 0
9.24.104.167 - - [31/Mar/1999:16:29:25 +0500] "GET http://redbooks.itso.ibm.com/redpd
fs/SG245132.PDF HTTP/1.0" 200 2566355
```

Figure 174. Cancel Control - the Cache Log File

The WTE caching proxy server had already cached the entire file. Hence, it was not necessary to access the Web server again to download the file, and the client could be served from the WTE cache directly.

This completes the demonstration of the cancel control function of WTE.

Chapter 16. WTE SSL Tunneling Support

Secure Sockets Layer (SSL) connections provide server authentication and privacy of data transmitted through the network. Optionally, SSL connections provide client authentication.

The SSL protocol involves encryption and decryption processes, and in a Web environment an SSL connection is established directly between the client browser and the destination content server. Client and server could use asymmetric key encryption to communicate securely. However, this way to proceed would be too slow and would be too CPU-intensive. For this reason, with SSL, client and server use asymmetric encryption only to agree across the network on a common encryption key and then use this key to encrypt the communication. This shared key is known as the *session key*. For more details on SSL, see the IBM redbook *Java 2 Network Security*, SG24-2109-01.

IBM Web Traffic Express (WTE) V2.0, the caching and filtering component of IBM WebSphere Performance Pack Version 2, offers SSL support. A WTE server will not make any attempt to cache or decrypt the information that the client and the server exchange during an SSL connection, but it will establish a connection to the destination content server, pass the requests to it, and serve the responses back to the client without looking at the data.

SSL tunneling is often used to encrypt the HTTP communication between a client and a Web server. However, it is often used also for other protocols. For example, Lotus Notes 4.6 can use SSL tunneling to send its communication through an HTTP proxy server (such as WTE).

16.1 How to Set Up SSL Tunneling in WTE

WTE uses the `SSLTunneling` directive to allow SSL tunneling to any port on the destination host. Setting this directive to `On` allows SSL tunneling to any port on the destination server. Setting this directive to `Off` allows SSL tunneling only to ports given in the `Proxy` rules. If there are no `Proxy` rules for SSL tunneling, and the `SSLTunneling` directive is set to `Off`, then SSL tunneling is not allowed. If the `SSLTunneling` directive is `On`, you must also enable the method `CONNECT`, using the `Enable` directive.

Turning SSL tunneling on and off can be done not only by manually editing the WTE configuration file, ibmproxy.conf, but also using the Configuration and Administration Forms, as shown:

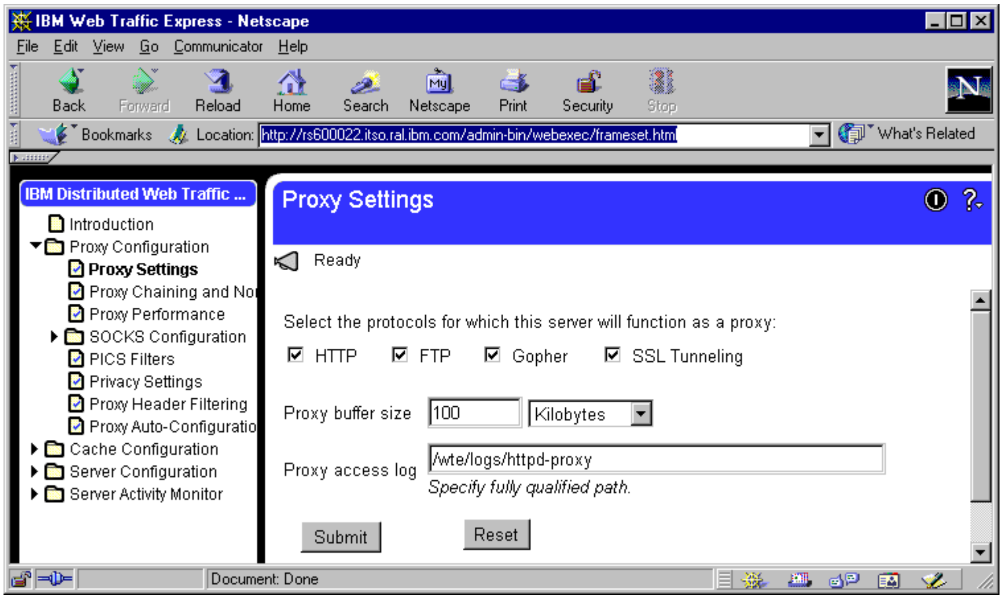

Proxy Settings

Ready

Select the protocols for which this server will function as a proxy:

☑ HTTP ☑ FTP ☑ Gopher ☑ SSL Tunneling

Proxy buffer size 100 Kilobytes

Proxy access log /wte/logs/httpd-proxy

Specify fully qualified path.

Submit Reset

Figure 175. Proxy Settings - SSL Tunneling

To turn SSL tunneling on, mark the **SSL Tunneling** button. To set it off, unmark it. Click the **Submit** button to make the change to the ibmproxy.conf configuration file and restart the WTE server to activate the changes.

16.2 SSL Tunneling Scenario

To experiment with SSL tunneling, the Web server must support SSL, and be configured to accept SSL traffic. We used Lotus Domino Go Webserver Version 4.6.2.5 as it supports SSL 3.0. The following figure shows a graphical representation of our scenario:

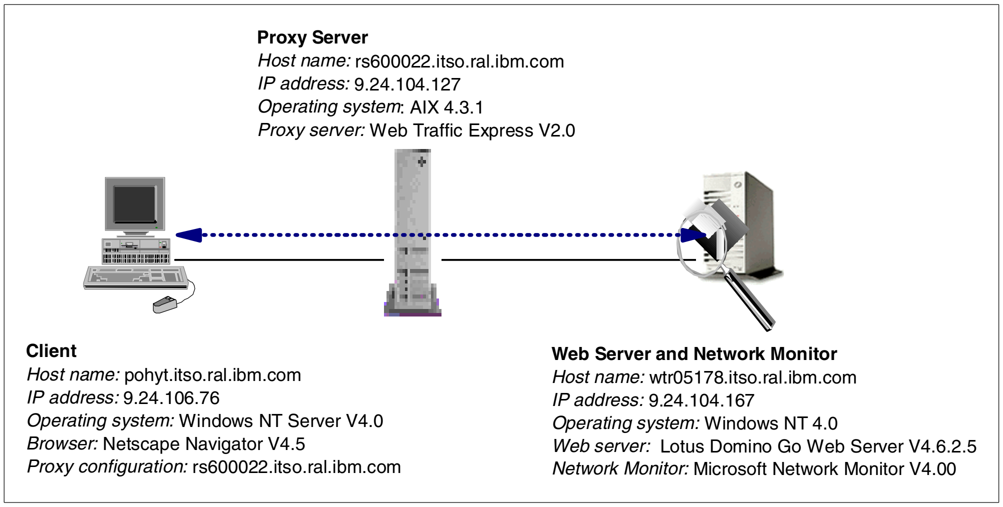

Proxy Server
Host name: rs600022.itso.ral.ibm.com
IP address: 9.24.104.127
Operating system: AIX 4.3.1
Proxy server: Web Traffic Express V2.0

Client
Host name: pohyt.itso.ral.ibm.com
IP address: 9.24.106.76
Operating system: Windows NT Server V4.0
Browser: Netscape Navigator V4.5
Proxy configuration: rs600022.itso.ral.ibm.com

Web Server and Network Monitor
Host name: wtr05178.itso.ral.ibm.com
IP address: 9.24.104.167
Operating system: Windows NT 4.0
Web server: Lotus Domino Go Web Server V4.6.2.5
Network Monitor: Microsoft Network Monitor V4.00

Figure 176. SSL Tunneling Scenario

16.2.1 Configuration of the Web Server Machine

As our purpose is mainly to demonstrate the configuration of WTE for SSL connections, we will describe only briefly the setup of Lotus Domino Go Webserver.

From the front page of Lotus Domino Go Webserver, click **CONFIGURATION AND ADMINISTRATION FORMS** and enter the Web server administrator's User Name and Password, as shown:

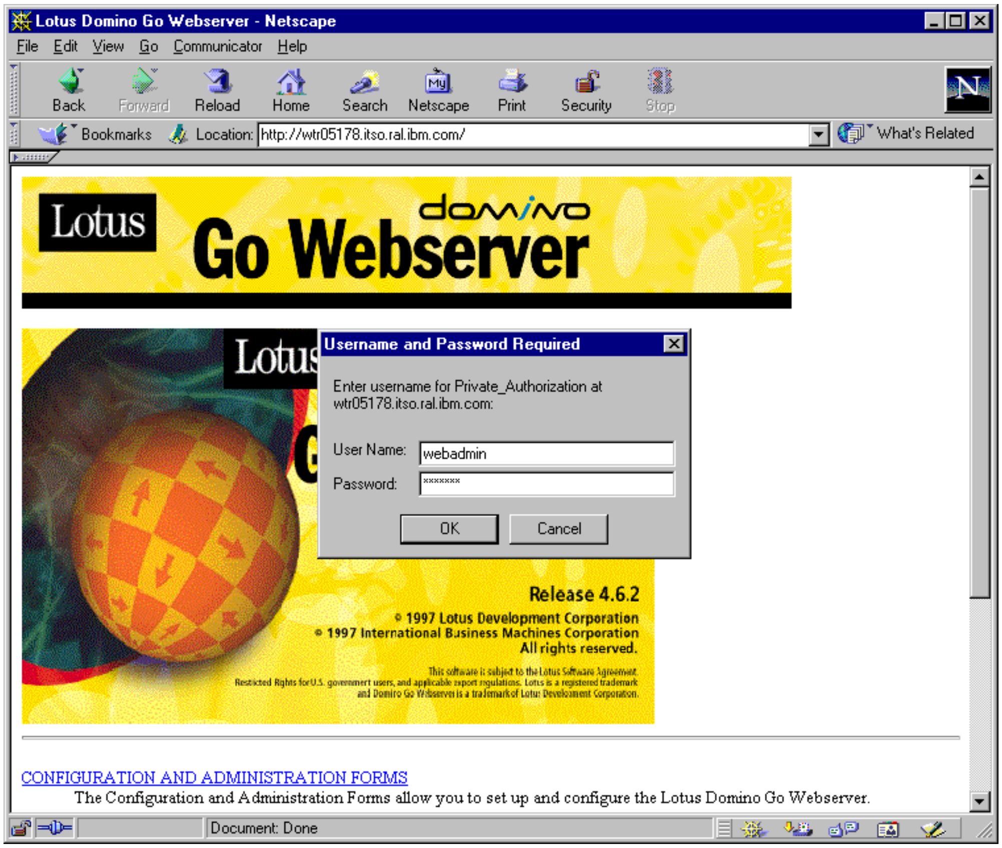
Lotus Domino Go Webserver - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks
What's Related
Lotus
domino
Go Webserver
OK
Cancel

Location: http://wtr05178.itso.ral.ibm.com/

Username and Password Required

Enter username for Private_Authorization at wtr05178.itso.ral.ibm.com:

User Name: webadmin

Password: xxxxxxx

Release 4.6.2
© 1997 Lotus Development Corporation
© 1997 International Business Machines Corporation
All rights reserved.

This software is subject to the Lotus Software Agreement. Restricted Rights for U.S. government users, and applicable export regulations. Lotus is a registered trademark and Domino Go Webserver is a trademark of Lotus Development Corporation.

CONFIGURATION AND ADMINISTRATION FORMS
The Configuration and Administration Forms allow you to set up and configure the Lotus Domino Go Webserver.

Document: Done

Figure 177. Lotus Domino Go Webserver

After logging in, you will be presented with the forms menu. Scroll down to the Security section, as shown in the following figure:

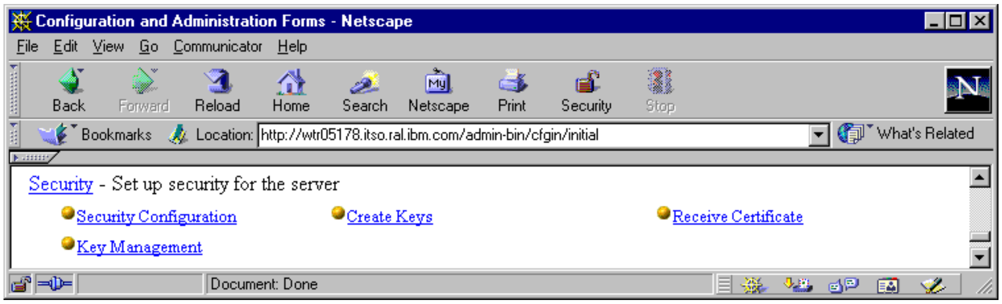
Configuration and Administration Forms - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks
What's Related

Location: http://wtr05178.itso.ral.ibm.com/admin-bin/cfgin/initial

Security - Set up security for the server

- Security Configuration
- Create Keys
- Receive Certificate
- Key Management

Document: Done

Figure 178. Lotus Domino Go Webserver - Configuration and Administration Forms

Click the **Create Keys** option, and the following window will be displayed:

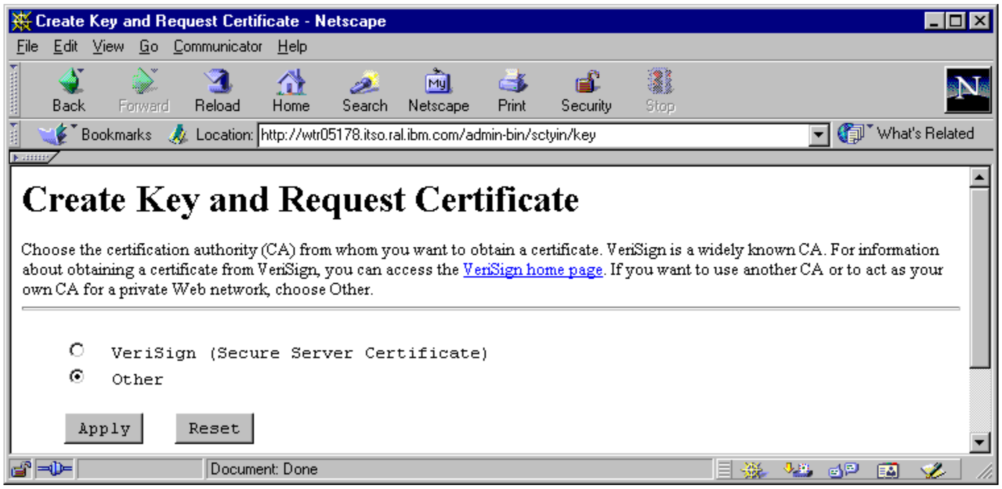
Create Key and Request Certificate - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/admin-bin/sctyin/key
What's Related
Document: Done

Create Key and Request Certificate

Choose the certification authority (CA) from whom you want to obtain a certificate. VeriSign is a widely known CA. For information about obtaining a certificate from VeriSign, you can access the VeriSign home page. If you want to use another CA or to act as your own CA for a private Web network, choose Other.

○ VeriSign (Secure Server Certificate)
● Other

Apply Reset

Figure 179. (Part 1 of 3). Lotus Domino Go Webserver - Create Keys

As you can see, we had to select the certification authority (CA) from which to obtain a certificate. Rather than requesting the certificate from VeriSign, we marked the **Other** box because we wanted to use our own CA for our scenario. Click the **Apply** button. The next screen is the Other Certificate form, where we create the public-private key pair. We type in the parameters as shown in the following two windows below:

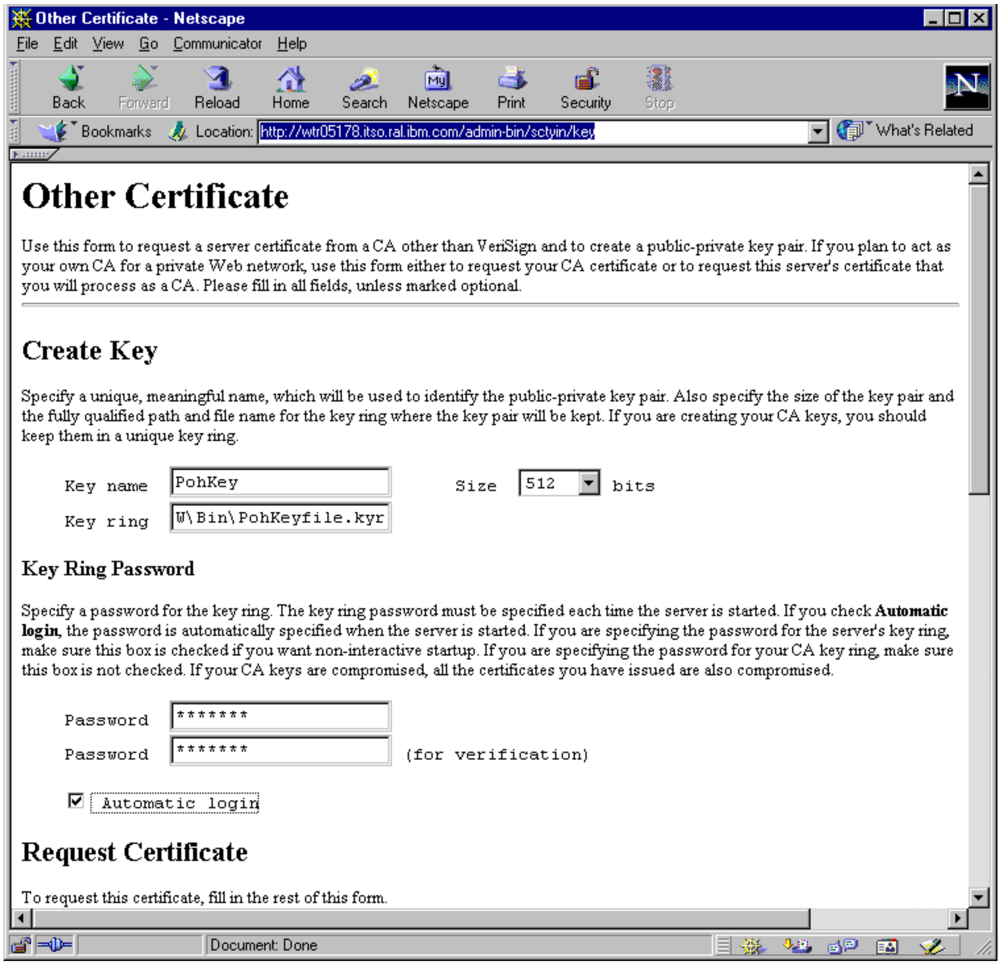
Other Certificate - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/admin-bin/sctyin/key
What's Related
Document: Done

Other Certificate

Use this form to request a server certificate from a CA other than VeriSign and to create a public-private key pair. If you plan to act as your own CA for a private Web network, use this form either to request your CA certificate or to request this server's certificate that you will process as a CA. Please fill in all fields, unless marked optional.

Create Key

Specify a unique, meaningful name, which will be used to identify the public-private key pair. Also specify the size of the key pair and the fully qualified path and file name for the key ring where the key pair will be kept. If you are creating your CA keys, you should keep them in a unique key ring.

Key name	PohKey	Size	512 bits
Key ring	W\Bin\PohKeyfile.kyr		

Key Ring Password

Specify a password for the key ring. The key ring password must be specified each time the server is started. If you check **Automatic login**, the password is automatically specified when the server is started. If you are specifying the password for the server's key ring, make sure this box is checked if you want non-interactive startup. If you are specifying the password for your CA key ring, make sure this box is not checked. If your CA keys are compromised, all the certificates you have issued are also compromised.

Password *******

Password ******* (for verification)

☑ Automatic login

Request Certificate

To request this certificate, fill in the rest of this form.

Figure 180. (Part 2 of 3). Lotus Domino Go Webserver - Create Keys

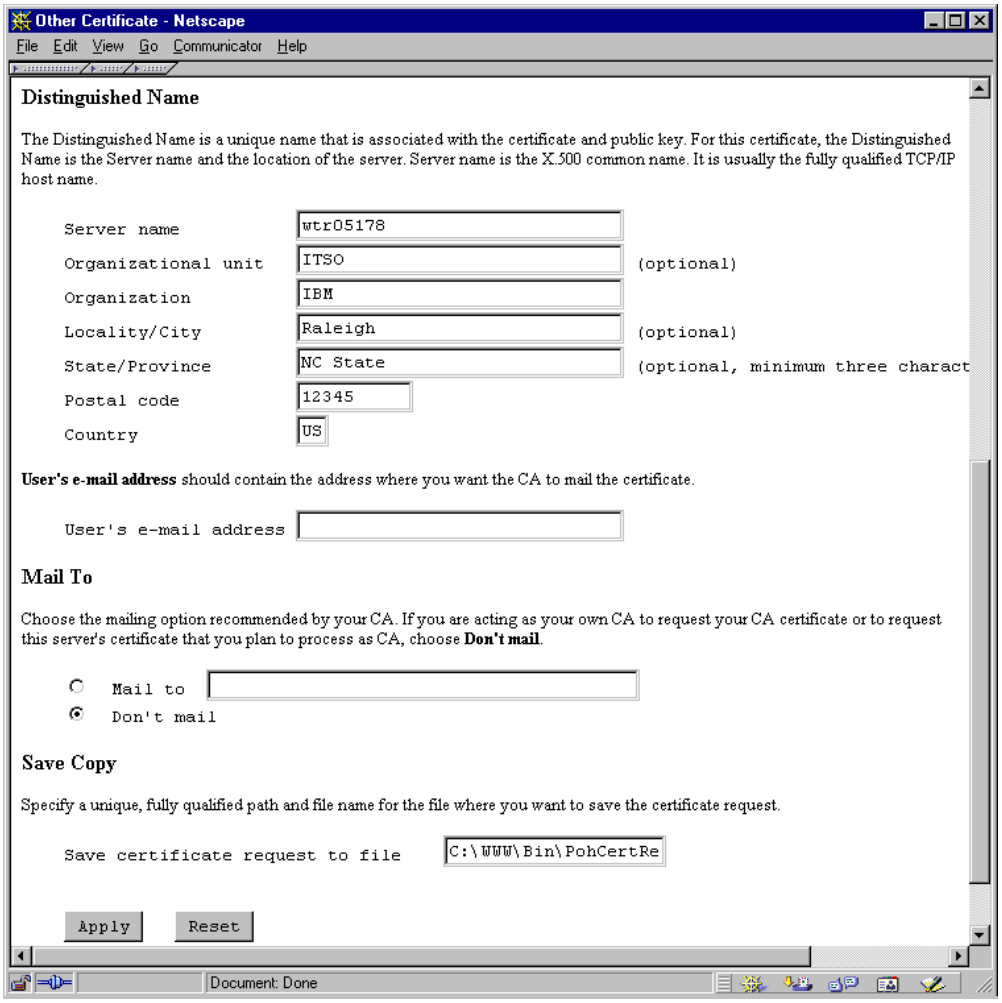

Other Certificate - Netscape
File Edit View Go Communicator Help

Distinguished Name

The Distinguished Name is a unique name that is associated with the certificate and public key. For this certificate, the Distinguished Name is the Server name and the location of the server. Server name is the X.500 common name. It is usually the fully qualified TCP/IP host name.

Server name	wtr05178	
Organizational unit	ITSO	(optional)
Organization	IBM	
Locality/City	Raleigh	(optional)
State/Province	NC State	(optional, minimum three charact
Postal code	12345	
Country	US	

User's e-mail address should contain the address where you want the CA to mail the certificate.

User's e-mail address

Mail To

Choose the mailing option recommended by your CA. If you are acting as your own CA to request your CA certificate or to request this server's certificate that you plan to process as CA, choose **Don't mail**.

○ Mail to

● Don't mail

Save Copy

Specify a unique, fully qualified path and file name for the file where you want to save the certificate request.

Save certificate request to file C:\WWW\Bin\PohCertRe

Figure 181. (Part 3 of 3). Lotus Domino Go Webserver - Create Keys

After entering the information, click the **Apply** button to make the changes and you will be presented with a Confirmation screen as follows:

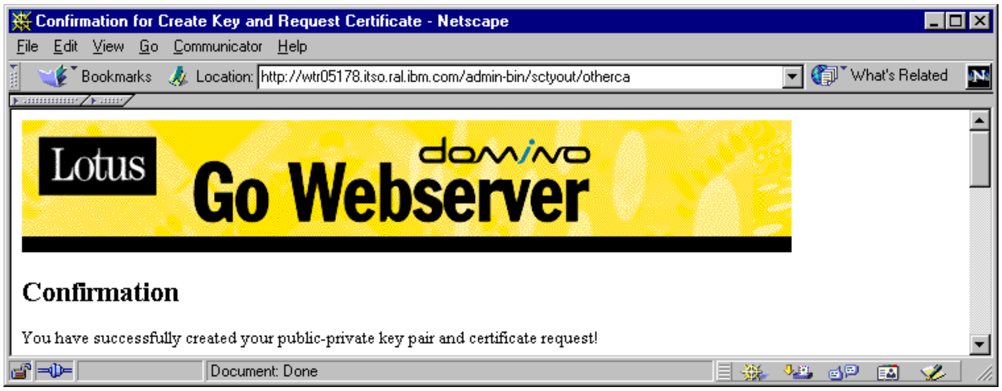

Confirmation for Create Key and Request Certificate - Netscape
File Edit View Go Communicator Help
Bookmarks
Location: http://wtr05178.itso.ral.ibm.com/admin-bin/sctyout/otherca
What's Related
Lotus
domino
Go Webserver

Confirmation

You have successfully created your public-private key pair and certificate request!

Figure 182. Lotus Domino Go Webserver - Confirmation

We now have a public-private key pair and a certificate request which we can send to a CA, such as Thawte or VeriSign. The CA would accept our certificate request and produce a certificate for us. However, Lotus Domino Go Webserver allows us to import the certificate request itself to generate a self-signed certificate. From the Configuration and Administration Forms, we clicked the **Receive Certificate** option for the following form:

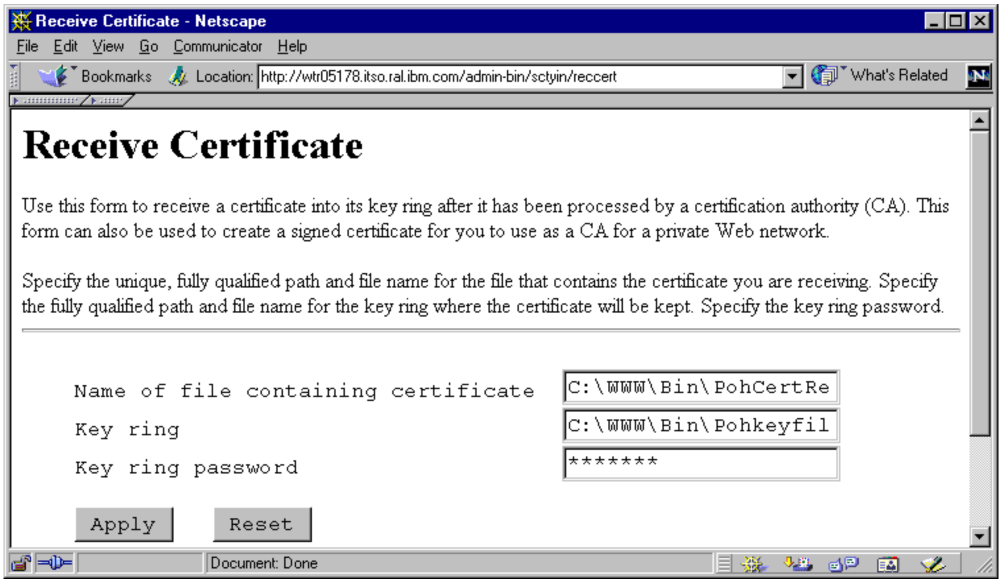
Receive Certificate - Netscape
File Edit View Go Communicator Help
Bookmarks
Location: http://wtr05178.itso.ral.ibm.com/admin-bin/sctyin/reccert
What's Related
Document: Done

Receive Certificate

Use this form to receive a certificate into its key ring after it has been processed by a certification authority (CA). This form can also be used to create a signed certificate for you to use as a CA for a private Web network.

Specify the unique, fully qualified path and file name for the file that contains the certificate you are receiving. Specify the fully qualified path and file name for the key ring where the certificate will be kept. Specify the key ring password.

Name of file containing certificate	C:\WWW\Bin\PohCertRe
Key ring	C:\WWW\Bin\Pohkeyfil
Key ring password	*******

Apply Reset

Figure 183. Lotus Domino Go Webserver - Receive Certificate

We enter the name of the text file containing the certificate request, `PohCertReq.txt`, the key ring file, `Pohkeyfile.kyr`, and the key ring password. Of course it is necessary to specify the full path for both the certificate request and key ring files. After clicking **Apply**, we got a confirmation window similar to Figure 182 on page 248.

Finally, from the Configuration and Administration Forms, we clicked the **Security Configuration** option to set the current key ring file. We obtained the following form:

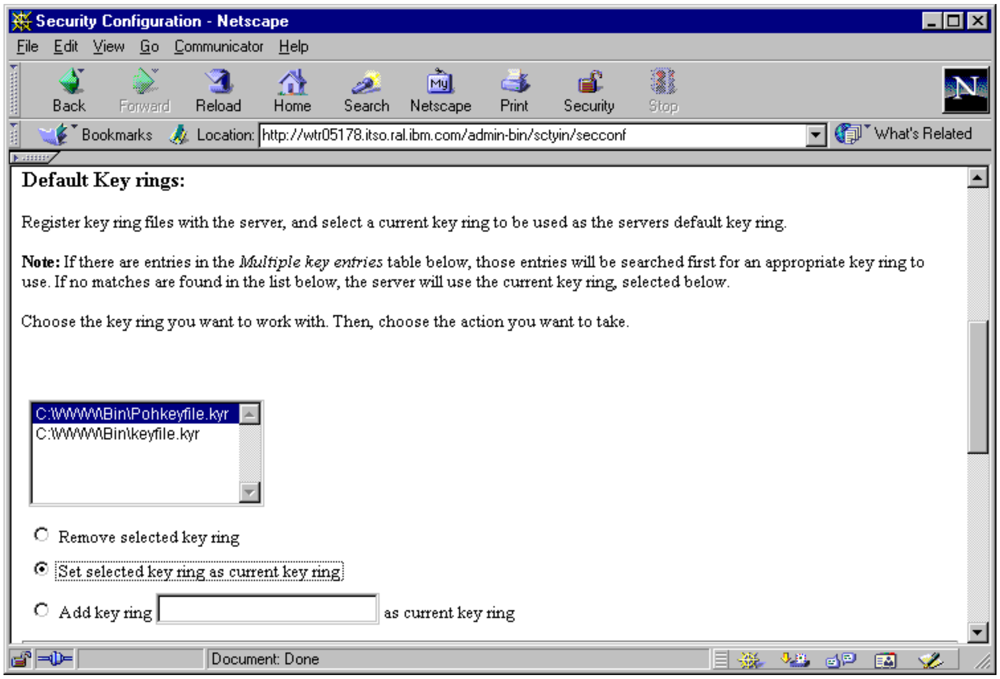

Security Configuration - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/admin-bin/sctyin/secconf
What's Related

Default Key rings:

Register key ring files with the server, and select a current key ring to be used as the servers default key ring.

Note: If there are entries in the *Multiple key entries* table below, those entries will be searched first for an appropriate key ring to use. If no matches are found in the list below, the server will use the current key ring, selected below.

Choose the key ring you want to work with. Then, choose the action you want to take.

C:\WWW\Bin\Pohkeyfile.kyr
C:\WWW\Bin\keyfile.kyr

○ Remove selected key ring

● Set selected key ring as current key ring

○ Add key ring ______ as current key ring

Figure 184. Lotus Domino Go Webserver - Security Configuration

We selected our key ring file and clicked on the option **Set selected key ring as current key ring**. Then, we click the **Apply** button to set the changes. The following confirmation screen appeared:

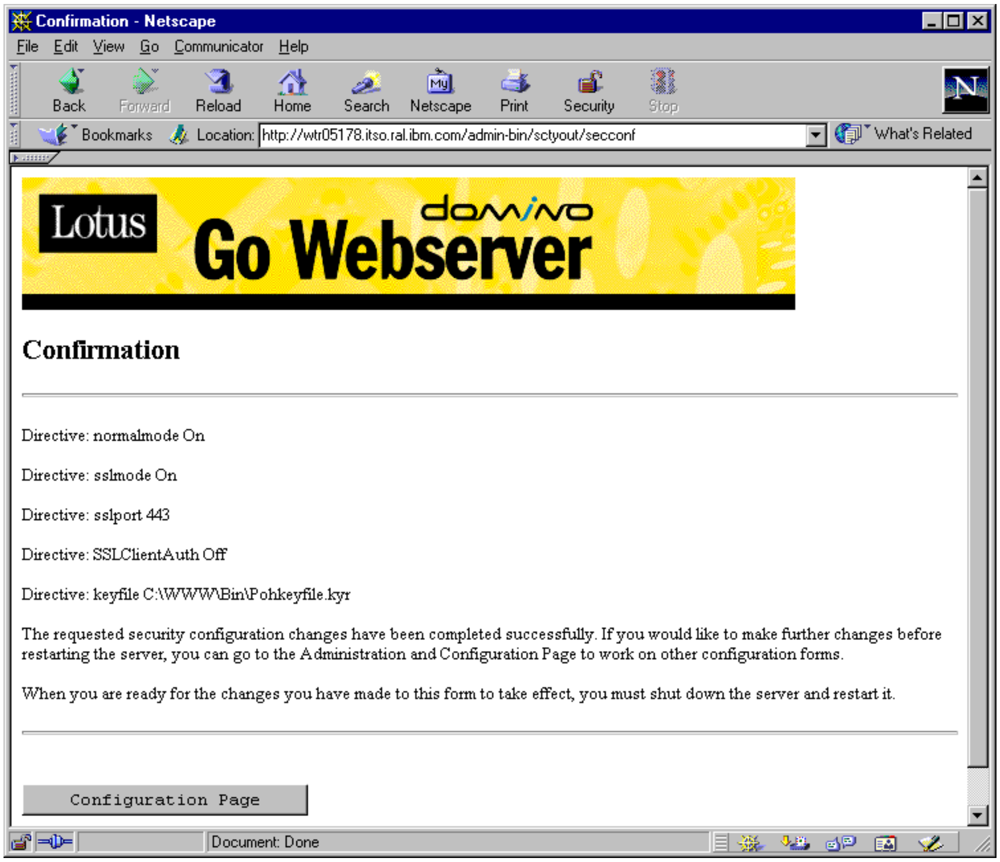
Confirmation - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://wtr05178.itso.ral.ibm.com/admin-bin/sctyout/secconf
What's Related
Lotus domino Go Webserver

Confirmation

Directive: normalmode On

Directive: sslmode On

Directive: sslport 443

Directive: SSLClientAuth Off

Directive: keyfile C:\WWW\Bin\Pohkeyfile.kyr

The requested security configuration changes have been completed successfully. If you would like to make further changes before restarting the server, you can go to the Administration and Configuration Page to work on other configuration forms.

When you are ready for the changes you have made to this form to take effect, you must shut down the server and restart it.

Figure 185. Lotus Domino Go Webserver - Confirmation

To activate these changes, we restarted Lotus Domino Go Webserver.

On the Web server machine, we also set up the Windows NT Network Monitor to monitor the traffic flowing through the Web server machine network interface (see Figure 176 on page 244).

16.2.2 Configuration of the Client Machine

On the client Web browser, we set the proxy to point to the WTE proxy server, rs600022.itso.ral.ibm.com, on port 80, as shown in the following figure:

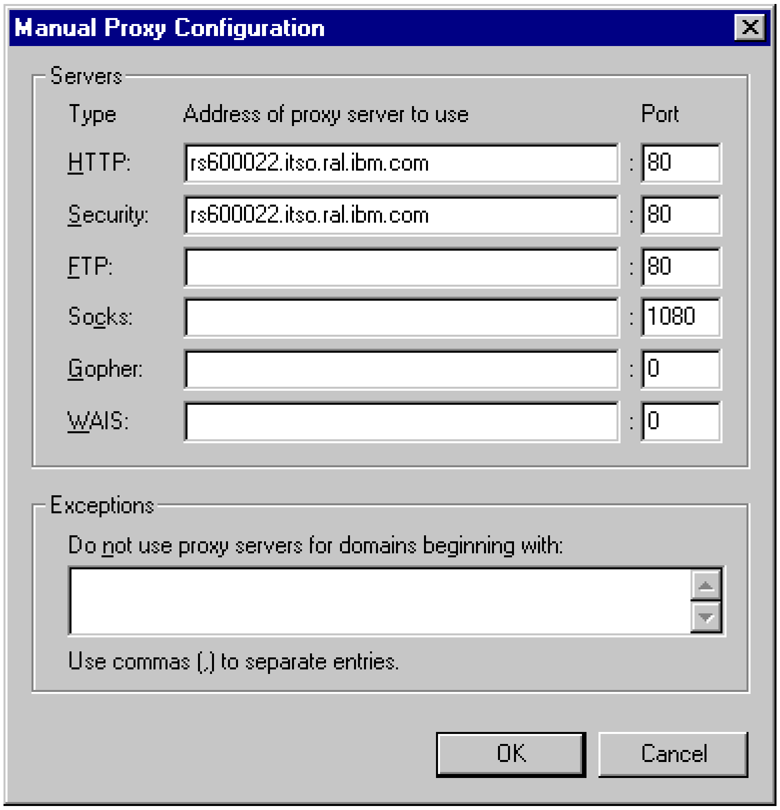

Manual Proxy Configuration

Servers

Type	Address of proxy server to use	Port
HTTP:	rs600022.itso.ral.ibm.com	80
Security:	rs600022.itso.ral.ibm.com	80
FTP:		80
Socks:		1080
Gopher:		0
WAIS:		0

Exceptions

Do not use proxy servers for domains beginning with:

Use commas (,) to separate entries.

Figure 186. Proxy Server Configuration on the Browser

16.2.3 SSL Tunneling Experience

Before experimenting with SSL tunneling, it was necessary to configure the WTE proxy server machine to work with SSL tunneling. The configuration of WTE is very simple, and is discussed in 16.1, “How to Set Up SSL Tunneling in WTE” on page 243.

For the purpose of comparison, we first pointed the browser to the Web server home page using the standard `http:` method. The URL we invoked was `http://wtr05178.itso.ral.ibm.com`, and the following is what we captured on the Windows NT Network Monitor. We can see from the captured screen that there was no encryption of the data, as shown:

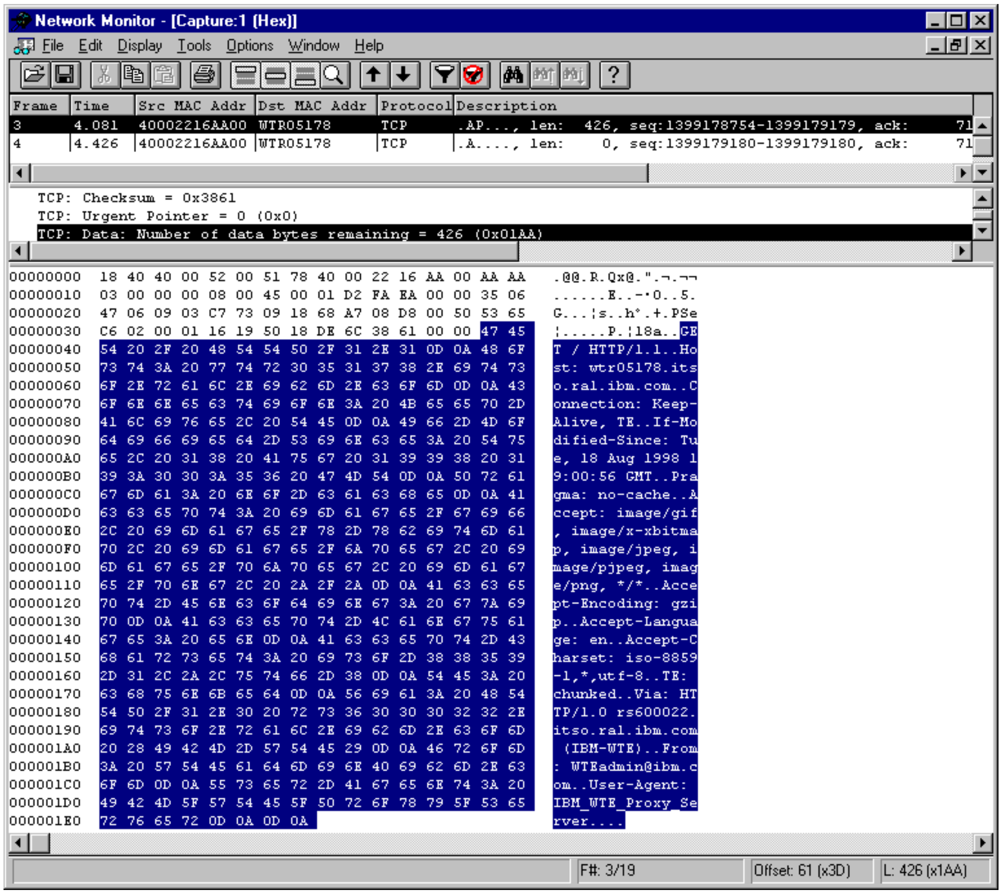
Network Monitor - [Capture:1 (Hex)]
File Edit Display Tools Options Window Help
?

```
Frame  Time   Src MAC Addr  Dst MAC Addr  Protocol Description
3      4.081  40002216AA00  WTR05178      TCP      .AP..., len:  426, seq:1399178754-1399179179, ack:    71
4      4.426  40002216AA00  WTR05178      TCP      .A...., len:    0, seq:1399179180-1399179180, ack:    71
```

```
   TCP: Checksum = 0x3861
   TCP: Urgent Pointer = 0 (0x0)
   TCP: Data: Number of data bytes remaining = 426 (0x01AA)
```

```
00000000  18 40 40 00 52 00 51 78 40 00 22 16 AA 00 AA AA   .@@.R.Qx@.".¬.¬¬
00000010  03 00 00 00 08 00 45 00 01 D2 FA EA 00 00 35 06   ......E..-·O..5.
00000020  47 06 09 03 C7 73 09 18 68 A7 08 D8 00 50 53 65   G...¦s..h°.+.PSe
00000030  C6 02 00 01 16 19 50 18 DE 6C 38 61 00 00 47 45   ¦.....P.¦l8a..GE
00000040  54 20 2F 20 48 54 54 50 2F 31 2E 31 0D 0A 48 6F   T / HTTP/1.1..Ho
00000050  73 74 3A 20 77 74 72 30 35 31 37 38 2E 69 74 73   st: wtr05178.its
00000060  6F 2E 72 61 6C 2E 69 62 6D 2E 63 6F 6D 0D 0A 43   o.ral.ibm.com..C
00000070  6F 6E 6E 65 63 74 69 6F 6E 3A 20 4B 65 65 70 2D   onnection: Keep-
00000080  41 6C 69 76 65 2C 20 54 45 0D 0A 49 66 2D 4D 6F   Alive, TE..If-Mo
00000090  64 69 66 69 65 64 2D 53 69 6E 63 65 3A 20 54 75   dified-Since: Tu
000000A0  65 2C 20 31 38 20 41 75 67 20 31 39 39 38 20 31   e, 18 Aug 1998 1
000000B0  39 3A 30 30 3A 35 36 20 47 4D 54 0D 0A 50 72 61   9:00:56 GMT..Pra
000000C0  67 6D 61 3A 20 6E 6F 2D 63 61 63 68 65 0D 0A 41   gma: no-cache..A
000000D0  63 63 65 70 74 3A 20 69 6D 61 67 65 2F 67 69 66   ccept: image/gif
000000E0  2C 20 69 6D 61 67 65 2F 78 2D 78 62 69 74 6D 61   , image/x-xbitma
000000F0  70 2C 20 69 6D 61 67 65 2F 6A 70 65 67 2C 20 69   p, image/jpeg, i
00000100  6D 61 67 65 2F 70 6A 70 65 67 2C 20 69 6D 61 67   mage/pjpeg, imag
00000110  65 2F 70 6E 67 2C 20 2A 2F 2A 0D 0A 41 63 63 65   e/png, */*..Acce
00000120  70 74 2D 45 6E 63 6F 64 69 6E 67 3A 20 67 7A 69   pt-Encoding: gzi
00000130  70 0D 0A 41 63 63 65 70 74 2D 4C 61 6E 67 75 61   p..Accept-Langua
00000140  67 65 3A 20 65 6E 0D 0A 41 63 63 65 70 74 2D 43   ge: en..Accept-C
00000150  68 61 72 73 65 74 3A 20 69 73 6F 2D 38 38 35 39   harset: iso-8859
00000160  2D 31 2C 2A 2C 75 74 66 2D 38 0D 0A 54 45 3A 20   -1,*,utf-8..TE: 
00000170  63 68 75 6E 6B 65 64 0D 0A 56 69 61 3A 20 48 54   chunked..Via: HT
00000180  54 50 2F 31 2E 30 20 72 73 36 30 30 30 32 32 2E   TP/1.0 rs600022.
00000190  69 74 73 6F 2E 72 61 6C 2E 69 62 6D 2E 63 6F 6D   itso.ral.ibm.com
000001A0  20 28 49 42 4D 2D 57 54 45 29 0D 0A 46 72 6F 6D    (IBM-WTE)..From
000001B0  3A 20 57 54 45 61 64 6D 69 6E 40 69 62 6D 2E 63   : WTEadmin@ibm.c
000001C0  6F 6D 0D 0A 55 73 65 72 2D 41 67 65 6E 74 3A 20   om..User-Agent: 
000001D0  49 42 4D 5F 57 54 45 5F 50 72 6F 78 79 5F 53 65   IBM_WTE_Proxy_Se
000001E0  72 76 65 72 0D 0A 0D 0A                           rver....
```

F#: 3/19 Offset: 61 (x3D) L: 426 (x1AA)

Figure 187. NT Network Monitor - HTTP Request

After this experience, we restarted the Windows NT Network Monitor on the Web server machine and again from the same browser, we accessed the Web server. However, this time we pointed the browser to the URL `https://wtr05178.itso.ral.ibm.com`, which by default uses port 443. The Netscape browser immediately displayed a warning message as it did not recognize the private CA that we set up. Hence, for the first SSL connection, there will be a few screens allowing you to accept the certificate. We accepted to trust the Web server's self-signed certificate, and this way we could successfully retrieve the Web server's front page. Then, we stopped the Windows NT Network Monitor. We can see from the captured screen, that the data was all encrypted as shown:

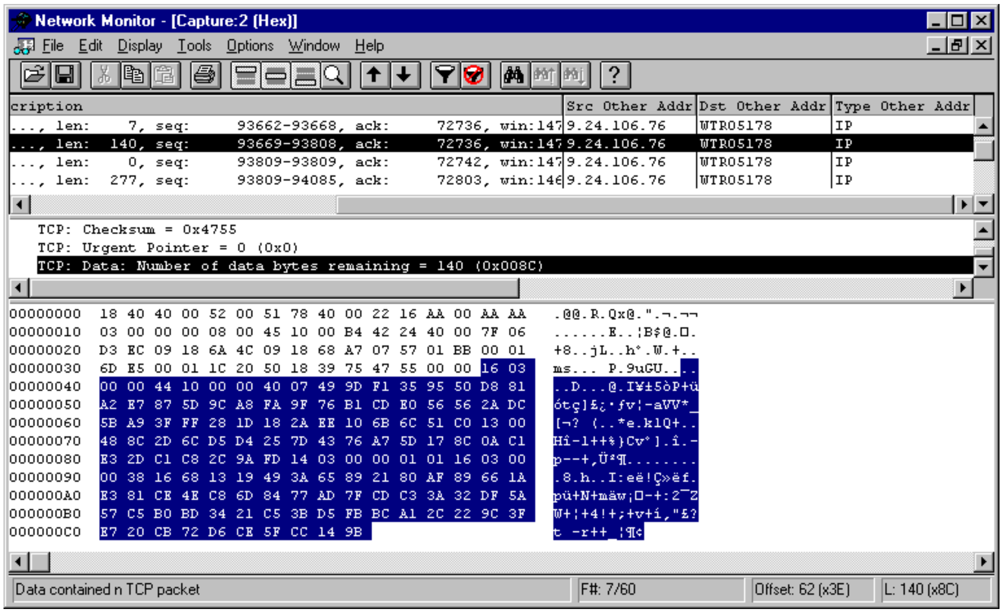
Network Monitor - [Capture:2 (Hex)]
File Edit Display Tools Options Window Help

```
cription                                                      Src Other Addr Dst Other Addr Type Other Addr
..., len:    7, seq:    93662-93668, ack:    72736, win:147 9.24.106.76    WTR05178       IP
..., len:  140, seq:    93669-93808, ack:    72736, win:147 9.24.106.76    WTR05178       IP
..., len:    0, seq:    93809-93809, ack:    72742, win:147 9.24.106.76    WTR05178       IP
..., len:  277, seq:    93809-94085, ack:    72803, win:146 9.24.106.76    WTR05178       IP

    TCP: Checksum = 0x4755
    TCP: Urgent Pointer = 0 (0x0)
    TCP: Data: Number of data bytes remaining = 140 (0x008C)

00000000  18 40 40 00 52 00 51 78 40 00 22 16 AA 00 AA AA   .@@.R.Qx@.".¬.¬¬
00000010  03 00 00 00 08 00 45 10 00 B4 42 24 40 00 7F 06   ......E..¦B$@.□.
00000020  D3 EC 09 18 6A 4C 09 18 68 A7 07 57 01 BB 00 01   +8..jL..h°.W.+..
00000030  6D E5 00 01 1C 20 50 18 39 75 47 55 00 00 16 03   ms... P.9uGU....
00000040  00 00 44 10 00 00 40 07 49 9D F1 35 95 50 D8 81   ..D...@.I¥±5òP+ü
00000050  A2 E7 87 5D 9C A8 FA 9F 76 B1 CD E0 56 56 2A DC   ótç]£¿·ƒv¦-aVV*_
00000060  5B A9 3F FF 28 1D 18 2A EE 10 6B 6C 51 C0 13 00   [¬? (..*e.klQ+..
00000070  48 8C 2D 6C D5 D4 25 7D 43 76 A7 5D 17 8C 0A C1   Hî-l++%}Cv°].î.-
00000080  E3 2D C1 C8 2C 9A FD 14 03 00 00 01 01 16 03 00   p--+,Ü²¶........
00000090  00 38 16 68 13 19 49 3A 65 89 21 80 AF 89 66 1A   .8.h..I:eë!Ç»ëf.
000000A0  E3 81 CE 4E C8 6D 84 77 AD 7F CD C3 3A 32 DF 5A   pü+N+mäw¡□-+:2¯Z
000000B0  57 C5 B0 BD 34 21 C5 3B D5 FB BC A1 2C 22 9C 3F   W+¦+4!+;+v+í,"£?
000000C0  E7 20 CB 72 D6 CE 5F CC 14 9B                     t -r++_¦¶¢
```

Data contained n TCP packet | F#: 7/60 | Offset: 62 (x3E) | L: 140 (x8C)

Figure 188. NT Network Monitor - HTTP Request

Finally, we take a look at the WTE proxy server cache access. We can only see the HTTP URL being accessed. No HTTPS URL was cached. To view the cache access log during the access, we use the `tail` AIX command on the proxy server as shown:

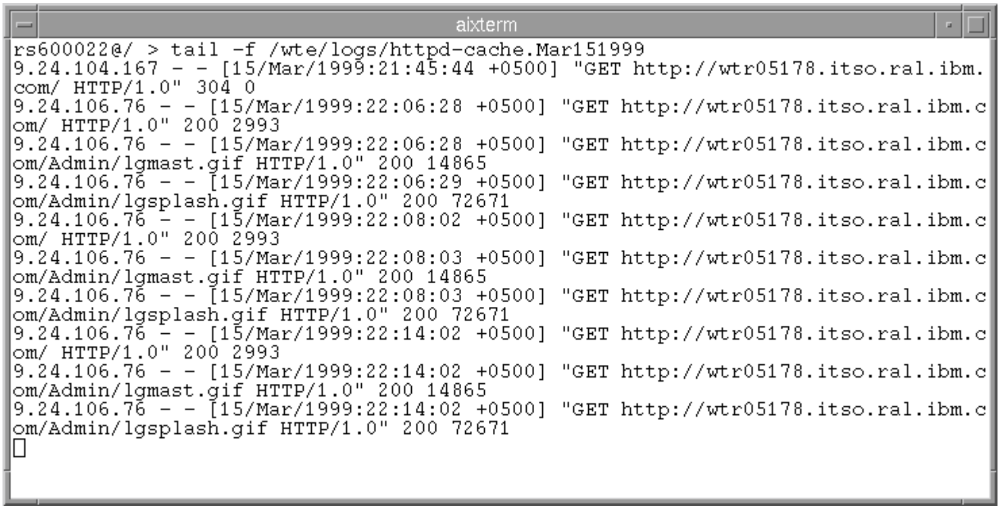

```
rs600022@/ > tail -f /wte/logs/httpd-cache.Mar151999
9.24.104.167 - - [15/Mar/1999:21:45:44 +0500] "GET http://wtr05178.itso.ral.ibm.
com/ HTTP/1.0" 304 0
9.24.106.76 - - [15/Mar/1999:22:06:28 +0500] "GET http://wtr05178.itso.ral.ibm.c
om/ HTTP/1.0" 200 2993
9.24.106.76 - - [15/Mar/1999:22:06:28 +0500] "GET http://wtr05178.itso.ral.ibm.c
om/Admin/lgmast.gif HTTP/1.0" 200 14865
9.24.106.76 - - [15/Mar/1999:22:06:29 +0500] "GET http://wtr05178.itso.ral.ibm.c
om/Admin/lgsplash.gif HTTP/1.0" 200 72671
9.24.106.76 - - [15/Mar/1999:22:08:02 +0500] "GET http://wtr05178.itso.ral.ibm.c
om/ HTTP/1.0" 200 2993
9.24.106.76 - - [15/Mar/1999:22:08:03 +0500] "GET http://wtr05178.itso.ral.ibm.c
om/Admin/lgmast.gif HTTP/1.0" 200 14865
9.24.106.76 - - [15/Mar/1999:22:08:03 +0500] "GET http://wtr05178.itso.ral.ibm.c
om/Admin/lgsplash.gif HTTP/1.0" 200 72671
9.24.106.76 - - [15/Mar/1999:22:14:02 +0500] "GET http://wtr05178.itso.ral.ibm.c
om/ HTTP/1.0" 200 2993
9.24.106.76 - - [15/Mar/1999:22:14:02 +0500] "GET http://wtr05178.itso.ral.ibm.c
om/Admin/lgmast.gif HTTP/1.0" 200 14865
9.24.106.76 - - [15/Mar/1999:22:14:02 +0500] "GET http://wtr05178.itso.ral.ibm.c
om/Admin/lgsplash.gif HTTP/1.0" 200 72671
```

Figure 189. SSL Scenario - Cache Access Log

Part 2. WebSphere Performance Pack Component Integration

Chapter 17. Content Based Routing

Content Based Routing (CBR) is an entirely new component within this version of SecureWay Network Dispatcher (ND). The purpose of CBR is to route client requests to specific servers based on the content of the URL request.

In this chapter we will show some scenarios using CBR. See *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858 for more information on how CBR works.

17.1 Installation of the CBR Function

The CBR component of ND Version 2.1 is supported on three operating systems: IBM AIX 4.2.1 or later, Microsoft Windows NT 4.0 and Sun Solaris 2.6 or later. Refer to the appropriate platform section of *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858 for details on how to use the Java InstallShield on your respective platform to install ND.

When you reach the point where you choose which ND component to install, select these three to install CBR on your machine:

- **Content Based Routing Runtime**
- **Content Based Routing Administration**
- **Content Based Routing License**

The following figure shows these three items as being selected:

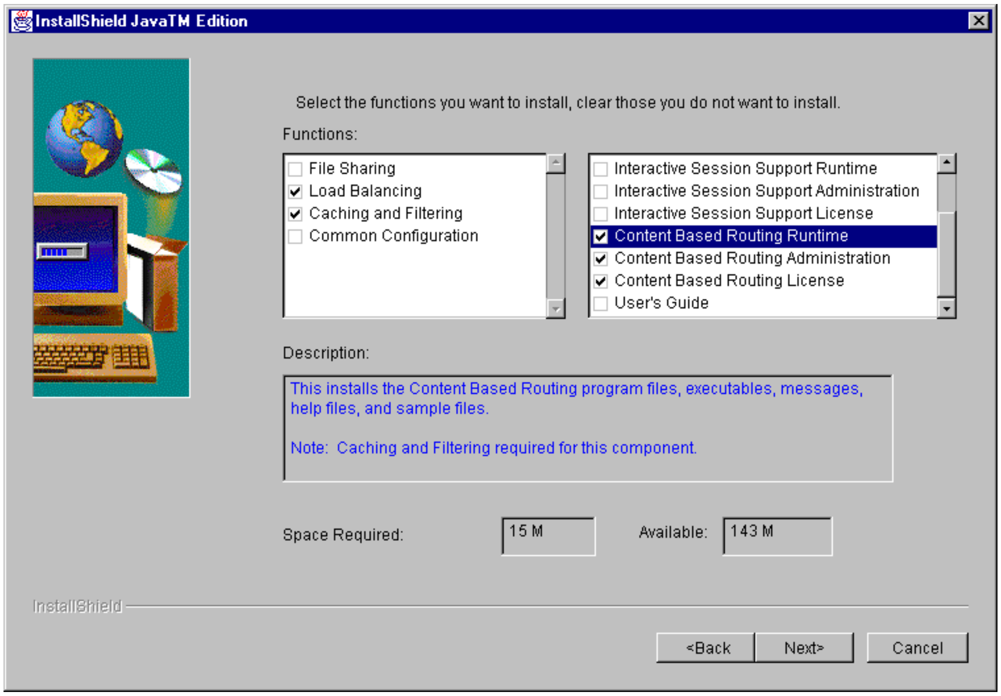
InstallShield JavaTM Edition

Select the functions you want to install, clear those you do not want to install.

Functions:

☐ File Sharing
☑ Load Balancing
☑ Caching and Filtering
☐ Common Configuration

☐ Interactive Session Support Runtime
☐ Interactive Session Support Administration
☐ Interactive Session Support License
☑ Content Based Routing Runtime
☑ Content Based Routing Administration
☑ Content Based Routing License
☐ User's Guide

Description:

This installs the Content Based Routing program files, executables, messages, help files, and sample files.

Note: Caching and Filtering required for this component.

Space Required: 15 M

Available: 143 M

Figure 190. Java InstallShield WebSphere Performance Pack Component Selection

When you select the **Content Based Routing Runtime** ND component, the **Caching and Filtering** WebSphere Performance Pack Version 2 component is automatically selected to be installed, as CBR cannot be installed without it.

The next screen asks if you would like to use the WTE proxy server to cache documents. If you click on the **Caching** checkbox, you will be asked to fill in three blanks:

- Cache Root Directory
- Cache Access Log
- Cache Size

The following figure shows this on the Windows NT platform:

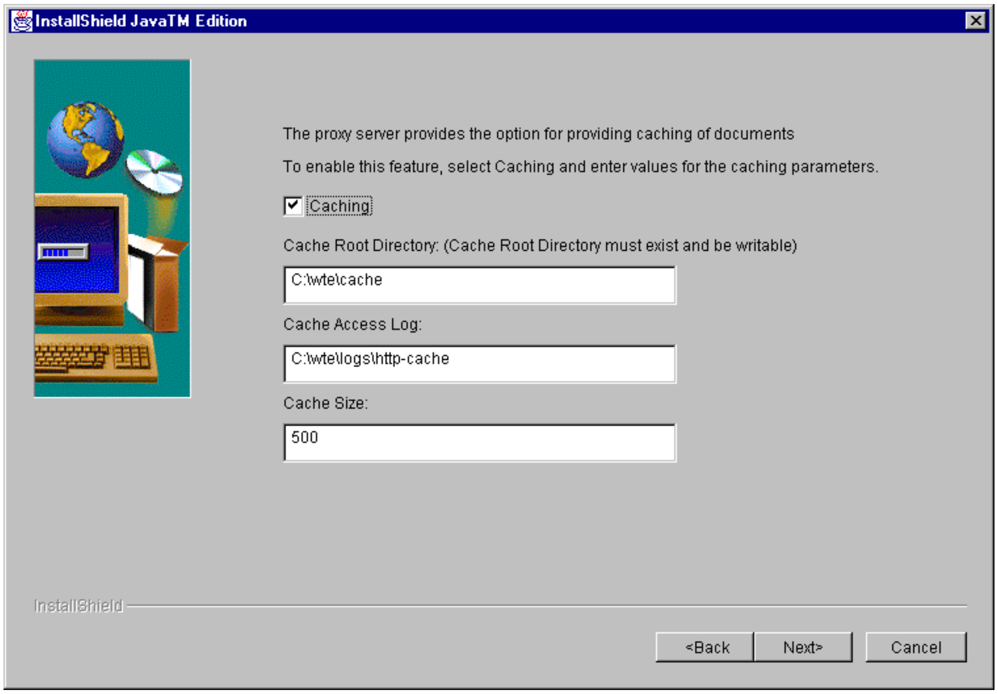

The proxy server provides the option for providing caching of documents

To enable this feature, select Caching and enter values for the caching parameters.

☑ Caching

Cache Root Directory: (Cache Root Directory must exist and be writable)

C:\wte\cache

Cache Access Log:

C:\wte\logs\http-cache

Cache Size:

500

Figure 191. CBR Install Process Asks for Caching Details

In our case, we wanted to use the caching function of the WTE proxy server, so we checked the **Caching** item in the dialog window, and entered the values shown in Figure 191 on page 258. 17.3.1, “WTE Configuration Overview” on page 260 explains how we selected these values. WTE and CBR can be installed with caching disabled for simplicity and then caching can be enabled and configured at a later time.

After we clicked on the **Next** button, CBR and WTE were installed.

17.1.1 Installation Locations

If CBR is installed from the WebSphere Performance Pack Version 2 InstallShield as demonstrated in 17.1, “Installation of the CBR Function” on page 257, the WTE component is also automatically installed. Alternatively, if CBR is installed from the SecureWay Network Dispatcher InstallShield, then WTE is not automatically installed and must be installed manually.

Note that in each case, the default base install directories will be different on Windows NT as shown in the following table:

Table 15. CBR and WTE Default Installation Base Directories

Platform	WebSphere Performance Pack components	
	WTE	**CBR**
AIX	/usr/lpp/internet/server_root	/usr/lpp/nd/cbr
Solaris	/opt/internet/server_root	/opt/nd/cbr
Windows NT WSPP InstallShield	C:\WSPP\WWW\	C:\WSPP\IBM\nd\cbr
Windows NT WTE or CBR InstallShield	C:\WWW	C:\Program Files\IBM\nd\cbr

17.2 Configuration of the CBR Function

In order to perform the configuration you must be the root user on AIX and Solaris or, if the ND server is installed on Windows NT, a member of the Administrators group.

Configuration of CBR cannot take place until WTE has been configured to start the CBR subprocess. We performed CBR configuration in three steps:

1. Because the CBR server is a subprocess of WTE, the first step we performed was to confirm that the WTE server was functional. Optionally, during this step, the WTE server can be configured with values different from the defaults set at installation time (for example, caching can be enabled).
2. Once this is done, four lines must be added to the WTE configuration file to enable CBR functionality. The WTE configuration file is:
   - C:\WINNT\ibmproxy.conf on Windows NT
   - /etc/ibmproxy.conf on AIX and Solaris

   Notice that the WTE server must be stopped and restarted for the changes to take effect.
3. Once the modifications have been made to the WTE configuration file, and the WTE server has been restarted using the modified configuration file, only then can CBR configuration be done. CBR configuration is very similar in style to Dispatcher configuration, in that one or more clusters are usually configured with their associated ports and servers. Similar to rule-based load balancing, rules are then defined, but in the case of CBR the type of the rule is Content. There are three methods of configuring CBR available:
   - `cbrcontrol` commands can be issued from the command line. As well, you can make use of the `cbrcontrol` *persistent session.* The persistent session is a limited shell like utility that can be used to enter `cbrcontrol` commands. Use the `cbrcontrol` command without any parameters to start the persistent session. You will receive the

     ```
     cbrcontrol >>
     ```

     prompt to which you can respond by entering `cbrcontrol` commands without the `cbrcontrol` keyword.

     See *SecureWay Network Dispatcher User's Guide Version 2.1 for Solaris, Windows NT and AIX*, GC31-8496 for `cbrcontrol` command usage information.

- The `cbrcontrol` configuration commands can be placed in a script or command file that can be executed as a batch job. When the file is run the commands are executed in sequence. A CBR configuration file can be loaded into the execution environment with the command:

  ```
  cbrcontrol file load configfilename
  ```

  See Figure 198 on page 272 for an example of a file containing `cbrcontrol` commands.
- The ND graphical user interface (GUI) can be used to perform all of the configuration steps that can be performed with `cbrcontrol` commands.

In addition, all three methods of configuration can be performed by a user on a machine other than the one that is being configured. This new feature is referred to as Remote Authenticated Administration and is explained further in *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858.

The three configuration steps above are described in the next section.

17.3 CBR Scenario

In this scenario, we demonstrate how to use CBR by configuring a basic environment. In 17.3.1, “WTE Configuration Overview” on page 260 we give an overview of how to configure WTE. In 17.3.2, “WTE Configuration File CBR Modifications” on page 261 we give details on how to change the WTE configuration file to enable the CBR subprocess. Then in 17.3.3, “CBR Configuration” on page 263 we use the ND GUI to perform CBR configuration of our cluster environment and add the Content type rules. Following this, we show how the rules influence which server is selected to service our request.

17.3.1 WTE Configuration Overview

The WTE component of WebSphere Performance Pack Version 2 contains functionality to work with the CBR component of ND Version 2.1. The steps required to start the WTE server are different on AIX and Solaris than on Windows NT. However, on all these platforms, WTE can be started with all of the default values without making any modifications to the configuration file.

We chose to modify the installed default configuration of WTE to make two changes. The first change was to modify the log used by the WTE proxy server (the proxy server is the basic function of the WTE server that is started if no user configuration is done after WTE installation) and the second change was to enable caching. *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858, contains information on WTE including details on how to configure its different functions.

Following a summary list of tasks that we performed to configure WTE on our server:

1. We created a WTE admin user ID and password for use in the following steps.
2. We created a separate file system or directory (/wte on UNIX or C:\wte in the case of Windows NT) to contain the WTE cache root directory and log files. On AIX, we found it necessary to change its ownership with the following command:

```
chown nobody.nobody /wte
```

3. We used the WTE Configuration and Administration Forms, accessible by directing a Web browser to the machine where WTE is installed, to change the location of the proxy log file. The only item we changed was the proxy server log file.
4. In the Configuration and Administration Forms, clicking **Submit** caused the ibmproxy.conf file to be updated with our new log file, but **Submit** or even **Restart Server** (the bar icon at the top of the window) did not result in these changes being used by the server. So we had to stop and restart the WTE server.
5. At this point, we were able to start the Netscape Navigator browser on our client machine. We selected **Preferences** from the Edit menu. Then, we clicked **Advanced** and **Proxies** and in the Proxies panel, we selected the **Manual Proxy Configuration** radio button. In the HTTP field of the Manual Proxy Configuration window, we entered the IP address of the machine where we had installed and configured the WTE server, with the corresponding port field set to `80`.
6. We made a request for a Web page that was available in another subnet that we had connectivity to. At this point we were presented with the Web page and verified from the Server Activity Monitor, accessible from the WTE Configuration and Administration Forms, that the request had been proxied by our server.
7. The process of configuring the caching component is very similar to this and was done by selecting the **Cache Configuration** item in the WTE Configuration and Administration Forms window.

17.3.2 WTE Configuration File CBR Modifications

When WTE is installed, the WTE configuration file, ibmproxy.conf, contains many lines of preconfigured WTE directives. It is necessary to add four lines to this configuration file to enable WTE to start the CBR server subprocess. In the configuration file, we searched for the word `ServerInit` and discovered that `ServerInit` was only contained in the default ibmproxy.conf file as a comment. At the end of the `ServerInit` comment section, we added the four lines.

It is important to note that the first of the four lines is very long and therefore wraps over the right side of the line to the next line. The parameters on the first line are as follows:

- Client keys directory
- Server key directory
- Install path
- Class path
- RMI port
- Log directory
- Save directory

The four lines are shown in the following screens. In each of the screens, it appears that the `CBR_CLIENT_KEYS_DIRECTORY` item is on a new line, but this is due to a formatting problem in this publication. In fact, the `CBR_CLIENT_KEYS_DIRECTORY` item is on the same line as the `ServerInit` directive. On the `ServerInit` line there is one space after the `ServerInit` keyword and one space before

CBR_CLIENT_KEYS_DIRECTORY. Do not include the new line character or any spaces within the parameter string on this line.

Some of the fields have variable-like names included on the line, and others do not. The fields are comma delimited. Following are the four lines that are necessary to configure CBR. The first figure contains the lines we used on our Windows NT system in C:\WINNT\ibmproxy.conf:

```
ServerInit C:\WSPP\IBM\nd\cbr\lib\libndcbr.dll:ndServerInit
CBR_CLIENT_KEYS_DIRECTORY=C:\WSPP\IBM\nd\admin\keys\cbr,CBR_SERVER_KEYS_DIRECTORY=C:\WSPP\IBM\nd\cbr\key,
END_INSTALL_PATH=C:\WSPP\IBM\nd,C:\WSPP\IBM\nd\cbr\lib;C:\WSPP\IBM\nd\cbr\lib\ibmcbr.jar;C:\WSPP\IBM\nd\a
dmin\lib\ChartRuntime.jar,11099,C:\WSPP\IBM\nd\cbr\logs\,C:\WSPP\IBM\nd\cbr\configurations\

PreExit C:\WSPP\IBM\nd\cbr\lib\libndcbr.dll:ndPreExit

PostExit C:\WSPP\IBM\nd\cbr\lib\libndcbr.dll:ndPostExit

ServerTerm C:\WSPP\IBM\nd\cbr\lib\libndcbr.dll:ndServerTerm
```

Note that if you installed CBR with the ND InstallShield rather than the WebSphere Performance Pack InstallShield, the path contained in these lines will have to be changed from `\WSPP\IBM` to `\Progra~1\IBM`. You are required to type the DOS format of the Program Files directory because there can be no spaces within the path elements on these lines.

The changes to the WTE configuration file /etc/ibmproxy.conf on AIX are shown in the following screen:

```
ServerInit /usr/lpp/nd/cbr/lib/libndcbr.so:ndServerInit
CBR_CLIENT_KEYS_DIRECTORY=/usr/lpp/nd/admin/keys/cbr,CBR_SERVER_KEYS_DIRECTORY=/usr/lpp/nd/cbr/key,END_I
NSTALL_PATH=/usr/lpp/nd,/usr/lpp/nd/cbr/lib:/usr/lpp/nd/cbr/lib/ibmcbr.jar:/usr/lpp/nd/admin/lib/ChartRun
time.jar,11099,/usr/lpp/nd/cbr/logs/,/usr/lpp/nd/cbr/configurations/

PreExit /usr/lpp/nd/cbr/lib/libndcbr.so:ndPreExit

Postexit /usr/lpp/nd/cbr/lib/libndcbr.so:ndPostExit

ServerTerm /usr/lpp/nd/cbr/lib/libndcbr.so:ndServerTerm
```

The changes to the WTE configuration file /etc/ibmproxy.conf on Solaris are shown in the following screen:

```
ServerInit /opt/nd/cbr/lib/libndcbr.so:ndServerInit
CBR_CLIENT_KEYS_DIRECTORY=/opt/nd/admin/keys/cbr,CBR_SERVER_KEYS_DIRECTORY=/opt/nd/cbr/key,END_INSTALL_PA
TH=/opt/nd,/opt/nd/cbr/lib:/opt/nd/cbr/lib/ibmcbr.jar:/opt/nd/admin/lib/ChartRuntime.jar,11099,/opt/nd/cb
r/logs/,/opt/nd/cbr/configurations/

PreExit /opt/nd/cbr/lib/libndcbr.so:ndPreExit

Postexit /opt/nd/cbr/lib/libndcbr.so:ndPostExit

ServerTerm /opt/nd/cbr/lib/libndcbr.so:ndServerTerm
```

After adding these lines to the WTE configuration file and saving the file, there is another configuration issue you should take care of:

- On Windows NT, add the following to the Path system environment variable:

  ```
  C:\jdk1.1.7\bin;C:\Program Files\nd\cbr\lib
  ```

  Notice here that we are assuming that Version 1.1.7 of the Java Development Kit (JDK) was installed in the default installation directory C:\jdk1.1.7.
- On AIX, add the following to the LIBPATH system environment variable:

  ```
  /usr/jdk_base/lib:/usr/jdk_base/lib/aix/native_threads:/usr/lpp/nd/cbr/lib
  ```
- On Solaris, add the following to the LD_LIBRARY_PATH system environment variable:

  ```
  /opt/jre1.1.7/lib/sparc/native_threads:/opt/nd/cbr/lib
  ```

Now, you can restart WTE.

Starting the WTE Server on AIX Systems

On AIX, it is advisable to use the `httpd` command to start the WTE server in a CBR environment. The System Resource Control (SRC) command:

```
startsrc -s httpd
```

does not start WTE correctly in the CBR case because WTE cannot find the library files necessary to configure and start CBR. This can be verified by looking at the WTE error log file.

17.3.3 CBR Configuration

Once the ibmproxy.conf file has been modified with the four lines to enable CBR, and the WTE proxy server has been restarted, then CBR configuration can begin.

In this scenario we defined two clusters, A and B, each with two servers. The two servers in cluster A were AIX systems and the two servers in cluster B were Solaris systems. Each cluster already had one rule defined by us, which load balanced requests for a particular page to the two servers. In this section we describe the steps necessary to define two additional CBR rules.

17.3.3.1 Network Environment

The following figure is a graphical representation of our scenario environment:

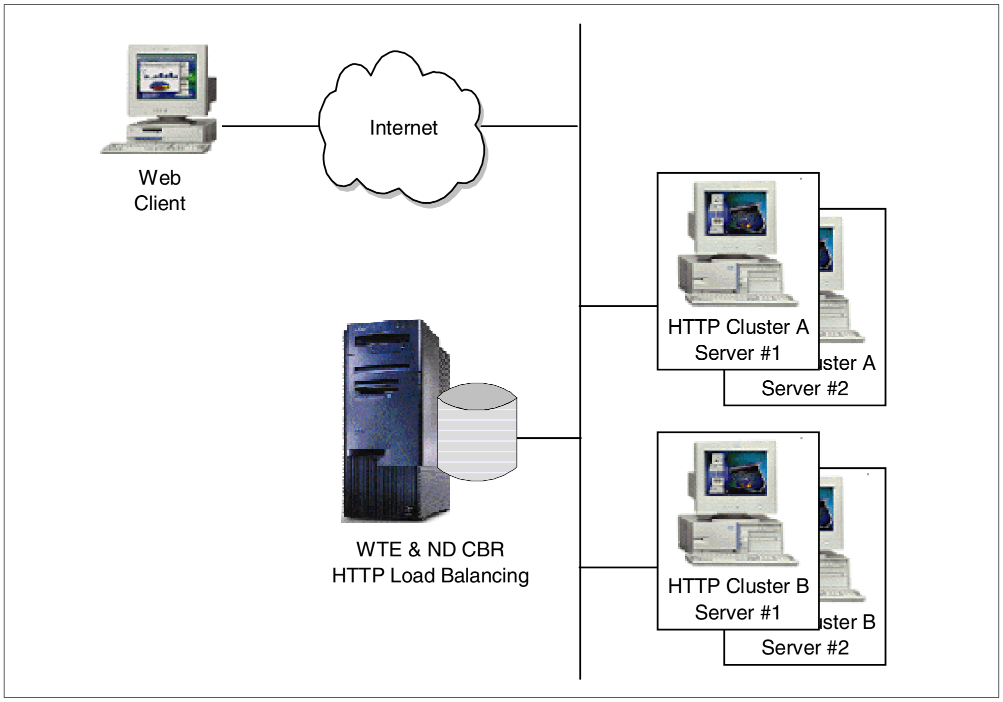

Figure 192. CBR Basic Scenario Environment

A summary of the software and network configuration of the environment where we performed our test is reported in the following table:

Table 16. CBR Basic Scenario - Hardware, Software and Network Configuration

Service		IP Address	Operating System
WTE & CBR server		9.67.133.67	Windows NT Server 4.0
Cluster A IP address 9.67.133.18	Web server 1	9.67.133.75	Solaris 2.6
	Web server 2	9.67.133.77	Solaris 2.6
Cluster B IP address 9.67.134.221	Web server 1	9.67.131.151	AIX 4.3.1
	Web server 2	9.67.131.153	AIX 4.3.1

17.3.3.2 Cluster, Port, Server and Rule Configuration

Once the CBR-enabled proxy server was running, we used the ND GUI to configure our CBR cluster. We used the `ndadmin` command to start the GUI and after right-clicking the **Content Based Routing** component, we selected **Connect to Host** from the pop-up menu. This is an effect of the new Remote Authenticated Administration feature which enables the ND component configuration to be done on a remote client machine. In this case, we were performing this configuration on the Dispatcher machine itself, however the connection still had to take place.

If, at the point where you try to connect to the host to do configuration either through the ND GUI or with the `cbrcontrol` commands, you receive a key related error message, then it is possible that there is an error in the four lines that were added to your ibmproxy.conf file.

We then started the Executor, added our clusters, and to each cluster a port and two servers. As the following figure demonstrates, for our testing purposes we had already added a rule named sunpage for cluster A and a rule named aixpage for cluster B. Following this configuration, the ND GUI appeared as follows:

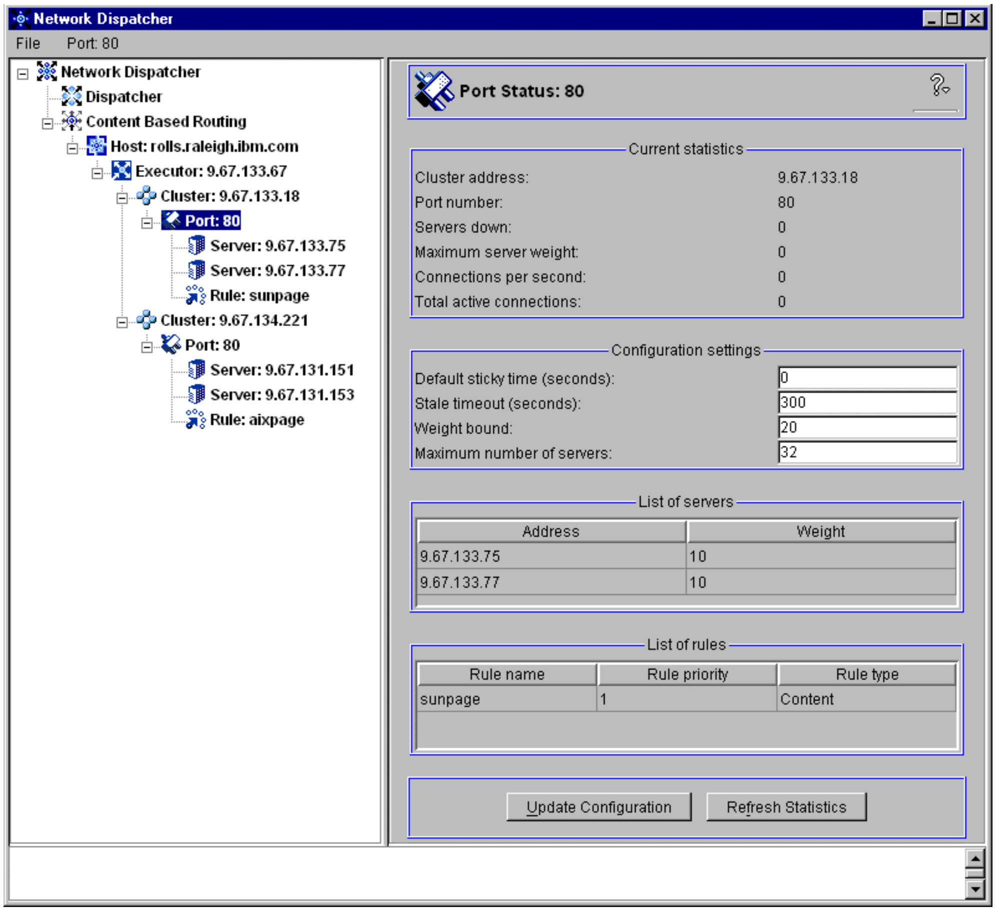

Current statistics

Cluster address:	9.67.133.18
Port number:	80
Servers down:	0
Maximum server weight:	0
Connections per second:	0
Total active connections:	0

Configuration settings

Default sticky time (seconds):	0
Stale timeout (seconds):	300
Weight bound:	20
Maximum number of servers:	32

List of servers

Address	Weight
9.67.133.75	10
9.67.133.77	10

List of rules

Rule name	Rule priority	Rule type
sunpage	1	Content

Figure 193. ND GUI Showing CBR Cluster, Port and Server Information

We will show you now details of how to add two more rules to the 9.67.133.18 cluster. To configure the content rules that CBR uses to determine which servers to load balance the request amongst, we right-clicked the **Port:80** item in the navigation portion of the GUI and selected **Add Rule...** from the pop-up menu, as shown in the following window:

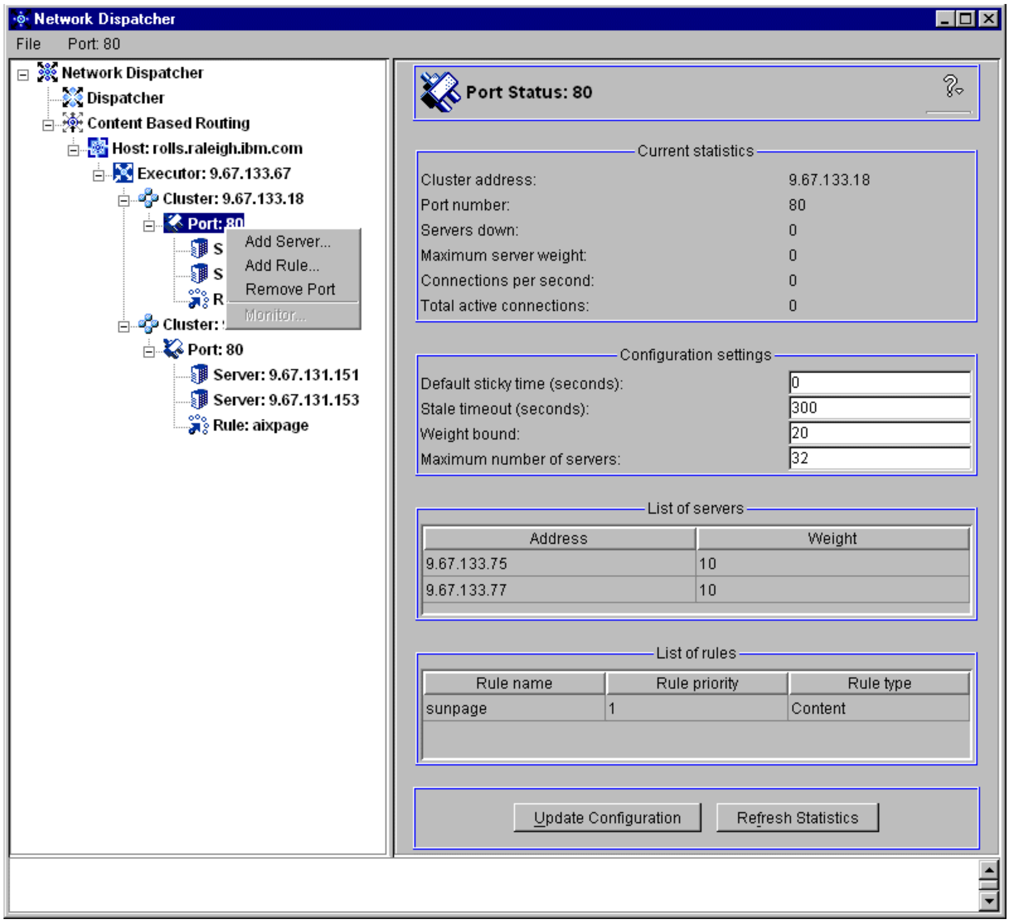

Current statistics

Cluster address:	9.67.133.18
Port number:	80
Servers down:	0
Maximum server weight:	0
Connections per second:	0
Total active connections:	0

Configuration settings

Default sticky time (seconds):	0
Stale timeout (seconds):	300
Weight bound:	20
Maximum number of servers:	32

List of servers

Address	Weight
9.67.133.75	10
9.67.133.77	10

List of rules

Rule name	Rule priority	Rule type
sunpage	1	Content

Figure 194. ND GUI Showing Add Rule Menu Item

Of course you should plan the logic that you want CBR to follow before you start adding rules to your configuration.

Notice that the network interface card on the CBR machine must be aliased to all the cluster addresses defined in the configuration above. This is the same thing you would do in a Dispatcher scenario. However, unlike a Dispatcher scenario, the loopback adapters on the TCP server machines must not be aliased to any cluster address.

Cluster Aliasing Note

The network interface card of the CBR machine must be aliased to all the cluster addresses used in the configuration. However, the loopback interface on the TCP server machines does not need to be aliased; for those who are familiar with this kind of aliasing operations in a Dispatcher scenario, it may seem unusual that similar aliasing is not done with CBR.

The reason for this is that when the packet is sent from the client machine, its destination IP address is the IP address of one of the clusters. When the packet arrives at the CBR machine, the packet is accepted because the network interface card on the CBR machine has been aliased to that cluster address. WTE receives the request and offers CBR the opportunity to examine its clusters and rules for a match. If a match is found, URL name translation is done.

If and when the proxy server sends the request to a clustered TCP server (the selection of which was done by CBR), WTE proxies a new request to the TCP server with the modified URL, and a destination IP address that is the IP address of the TCP server machine. When the TCP server replies, its response goes back to the CBR server, and is cached by WTE if caching is enabled and if WTE caching algorithms require caching of that particular page. For this reason, there is no requirement to alias the cluster address to the loopback interface on the TCP server machine.

On the contrary, Dispatcher keeps the cluster address as the destination address of the packet, and identifies the TCP servers through their Media Access Control (MAC) addresses. For this to work, the TCP servers need to have the cluster address aliased on their loopback interface. An advantage of this is that the server’s response can flow directly to the client, without any need for it to be proxied by the Dispatcher. This would not be a good solution with CBR though, because it would not allow caching.

Another positive effect of this is that a CBR cluster does not have the same restriction as a Dispatcher cluster with regard to the location of the TCP servers relative to the Dispatcher or CBR server. With CBR, because WTE proxies the request to the servers, there is no requirement for the servers to be located on the same LAN as the CBR machine.

Our next step was to create a content rule to direct all requests to cluster 9.67.133.18 that contained the string `productx.html` in the URL to server 9.67.133.75. To accomplish this, in the **Add Rule** dialog window, we filled in the fields as follows:

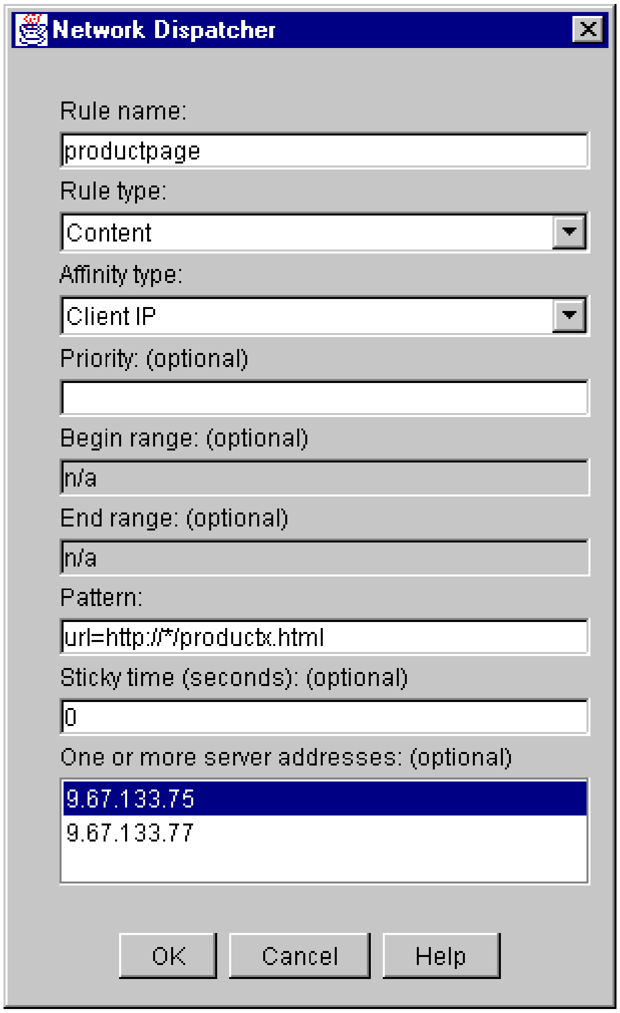
Network Dispatcher

Rule name: productpage

Rule type: Content

Affinity type: Client IP

Priority: (optional)

Begin range: (optional) n/a

End range: (optional) n/a

Pattern: url=http://*/productx.html

Sticky time (seconds): (optional) 0

One or more server addresses: (optional)
9.67.133.75
9.67.133.77

Figure 195. Add Rule Dialog Window for the productpage Rule

- We assigned the name `productpage` to this rule.
- We selected the type of the rule to be **Content**. This is the new rule type added for use by the CBR component.
- The Affinity type field contains one of two possible values: **Client IP** or **Cookie**. Client IP affinity as used in the Dispatcher component can also be used with CBR. The cookie affinity feature applies only to the CBR component and provides a new way to make clients *sticky* to a particular server. This function is enabled by setting the sticky time of a rule to a positive number, and setting the affinity to **Cookie**.

  We left the default value of **Client IP** in the Affinity type field. In the next rule we demonstrate setting cookie affinity and sticky time.
- We also did not put a value in the priority field for the rule. Priorities establish the order in which rules will be reviewed. This parameter accepts integer values. If you do not specify the priority of the first rule you add, CBR will set it by default to `1`. When a subsequent rule is added, by default its priority is calculated to be 10 plus the current lowest priority of any existing rule. For example, assume you have an existing rule whose priority is 30. You add a new rule and set its priority at `25` (which is a higher priority than 30). Then you

add a third rule without setting a priority. The priority of the third rule is calculated to be as 30 + 10 = 40, and so on.

- The Begin range parameter and the End range parameters are not used on Content type rules.
- The Pattern field is used to define the pattern of characters that CBR will match against each client request.

  The pattern must not contain any spaces and can make use of the special characters listed in the following table:

Table 17. Special Characters Allowed in the Pattern Field

Character	Function
*	Matches 0 to *x* of any character
(	Used for logic grouping
)	Used for logic grouping
&	Logical AND
|	Logical OR
!	Logical NOT

The reserved keywords shown in the following table must always be followed by the equal (=) sign:

Table 18. Keywords Followed by the Equal (=) Sign

Keyword	Value
client	Client IP address
url	URL in request
protocol	Protocol section of URL
path	Path section of URL
refer	Referred URL (quality of service)
user	User ID section of URL

In our case, we made use of the `url` reserved word and the wildcard (*) character. We specified the pattern as:

```
url=http://*/productx.html
```

This rule specifies that any URL that is a request for the page productx.html will be a match.

Other examples of valid rule patterns are:

- `url=http://*/*.gif`
- `client=9.32.*`
- `(path=index/*.gif&protocol=http)|(client=9.1.2.3)`
- `!(path=*.jpeg)`

- The Sticky time field is used along with the Affinity type field. We demonstrate the use of this in our next rule.
- The last item on the dialog window shown in Figure 195 on page 268 is a scrolled list of server addresses to optionally choose from. In this case we chose the server whose IP address was 9.67.133.75, meaning that we wanted

this server to serve the client requests containing the string productx.html in the URL.

We then clicked the **OK** button in the dialog window shown in Figure 195 on page 268.

The next rule was added to send requests to both servers in the 9.67.133.18 cluster for client requests for the page purchase.html. In this case, however, we specified an Affinity type of **Cookie** and a Sticky time of `30` seconds. This means that when a client request to the cluster is made and the URL matches the rule (that is, contains the string `purchase.html`), CBR would load balance the request to the best server of the two and then all subsequent HTTP requests from that client to this cluster address would also go to that server for a period of 30 seconds. We called this rule `purchase`. The dialog window in this case appeared as follows:

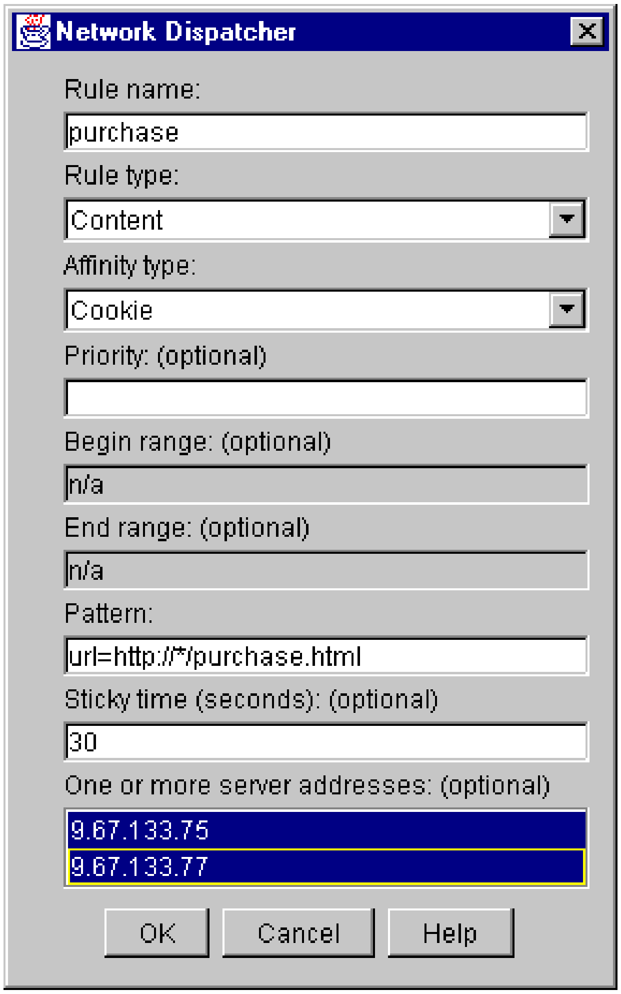

Network Dispatcher

Rule name: purchase

Rule type: Content

Affinity type: Cookie

Priority: (optional)

Begin range: (optional) n/a

End range: (optional) n/a

Pattern: url=http://*/purchase.html

Sticky time (seconds): (optional) 30

One or more server addresses: (optional)
9.67.133.75
9.67.133.77

Figure 196. Add Rule Dialog Window for the productpage Rule

We then clicked the **OK** button in the dialog window above. The GUI reported the configuration was updated with the added rules, as shown in the following figure:

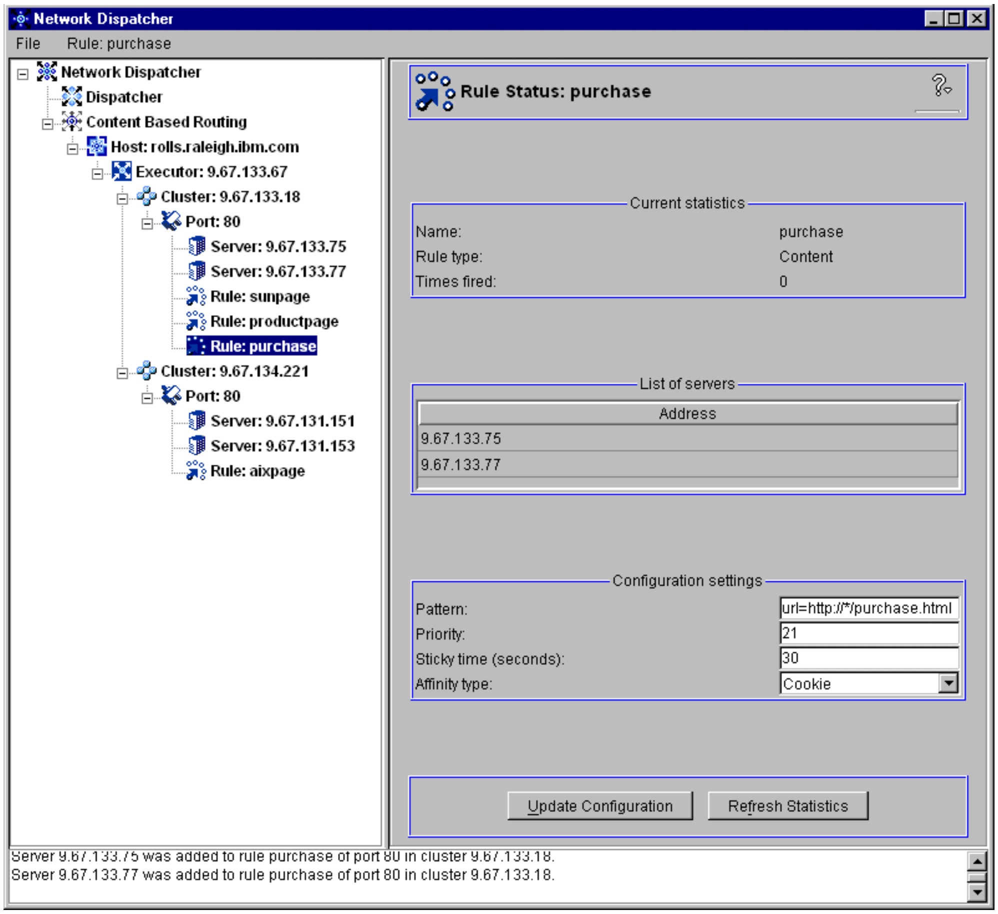

Name: purchase
Rule type: Content
Times fired: 0

Address
9.67.133.75
9.67.133.77

Pattern: url=http://*/purchase.html
Priority: 21
Sticky time (seconds): 30
Affinity type: Cookie

```
Server 9.67.133.75 was added to rule purchase of port 80 in cluster 9.67.133.18.
Server 9.67.133.77 was added to rule purchase of port 80 in cluster 9.67.133.18.
```

Figure 197. ND GUI Showing the Newly Added productpage Rule

17.3.4 CBR Manager and Advisors

As with the Dispatcher, starting the Manager and using the Advisors is optional. The Manager can be activated with the command:

```
cbrcontrol manager start
```

The typical advisor that is used with CBR is the HTTP advisor, which can be launched by entering the following command:

```
cbrcontrol advisor start http port
```

where *port* is the port configured for the cluster.

If you have configured any Advisors, you must change the Manager proportions to allow the Advisor information to be included in the load balancing decisions. To do this, you should use the command:

```
cbrcontrol manager proportions
```

For simplicity, in the scenarios described in this chapter, we did not make use of the CBR Manager or Advisors functions.

17.3.5 Saving the Configuration

Once we finished configuring CBR, we saved the configuration to a file. To do this from the ND GUI, we right-clicked **Host** and selected **Save Configuration File As...**.

In the resulting pop-up window, we were prompted to enter the name of the configuration file where we wanted to save this information. We entered the file name landon.cfg in the Save Configuration pop-up window. By default this file is placed in the directory configurations under *installbase*, where *installbase* is for the CBR component and varies by operating system. See Table 15 on page 259 for a list of the installbase locations.

The configuration file is saved in ASCII format and contains the list of commands that would be necessary to reconfigure your environment. Following is the content of the landon.cfg file:

```
cbrcontrol cluster add 9.67.133.18

cbrcontrol port add 9.67.133.18:80

cbrcontrol server add 9.67.133.18:80:9.67.133.77

cbrcontrol server add 9.67.133.18:80:9.67.133.75
cbrcontrol rule add 9.67.133.18:80:sunpage type content pattern url=http://*/sunpage.html priority 1 beginrange 0
endrange 0
cbrcontrol rule useserver 9.67.133.18:80:sunpage 9.67.133.75
cbrcontrol rule useserver 9.67.133.18:80:sunpage 9.67.133.77
cbrcontrol rule set 9.67.133.18:80:sunpage stickytime 30

cbrcontrol rule add 9.67.133.18:80:productpage type content pattern url=http://*/productx.html priority 11
beginrange 0 endrange 0
cbrcontrol rule useserver 9.67.133.18:80:productpage 9.67.133.75

cbrcontrol rule add 9.67.133.18:80:purchase type content pattern url=http://*/purchase.html priority 21 beginrange 0
endrange 0
cbrcontrol rule useserver 9.67.133.18:80:purchase 9.67.133.75
cbrcontrol rule useserver 9.67.133.18:80:purchase 9.67.133.77
cbrcontrol rule set 9.67.133.18:80:purchase stickytime 30
cbrcontrol rule set 9.67.133.18:80:purchase affinity cookie

cbrcontrol cluster add 9.67.134.221

cbrcontrol port add 9.67.134.221:80

cbrcontrol server add 9.67.134.221:80:9.67.131.153

cbrcontrol server add 9.67.134.221:80:9.67.131.151

cbrcontrol rule add 9.67.134.221:80:aixpage type content pattern url=http://*/aixpage.html priority 1 beginrange 0
endrange 0
cbrcontrol rule useserver 9.67.134.221:80:aixpage 9.67.131.151
cbrcontrol rule useserver 9.67.134.221:80:aixpage 9.67.131.153
```

Figure 198. Configuration File landon.cfg

Examination of the saved configuration file is interesting because it contains each of the `cbrcontrol` commands that could also be entered from the command line to configure the same environment. If you chose to configure CBR with commands, you can save your configured environment with the command:

cbrcontrol file save *configurationfilename*

You can then subsequently reload the configuration file with the command:

cbrcontrol file load *configurationfilename*

17.3.6 Scenario Results

We placed simple HTML files in the document root directories on each of our Web servers. On the Web server with IP address 9.67.133.75 we placed a file named productx.html and another called purchase.html. On the Web server with IP address 9.67.133.77 we placed a different file also called productx.html. Recall that we had defined a rule specifying that client requests containing the string `productx.html` in the URL be load balanced between both of these servers. Each of the files uniquely identified the server that it was located on and the name of the page.

Client Configuration

In a CBR scenario, the client machine does not need any special configuration. In particular, client browsers must not be set to redirect all the requests to the CBR machine. For example, in a real-life situation, it would not be appropriate to require end users to reconfigure their Web browsers before accessing a Web site that uses CBR. The use of CBR in a Web site is completely transparent to end users.

17.3.6.1 Client IP Affinity Demonstration

For our first request, with our browser we requested the productx.html page from our cluster address and received the following page:

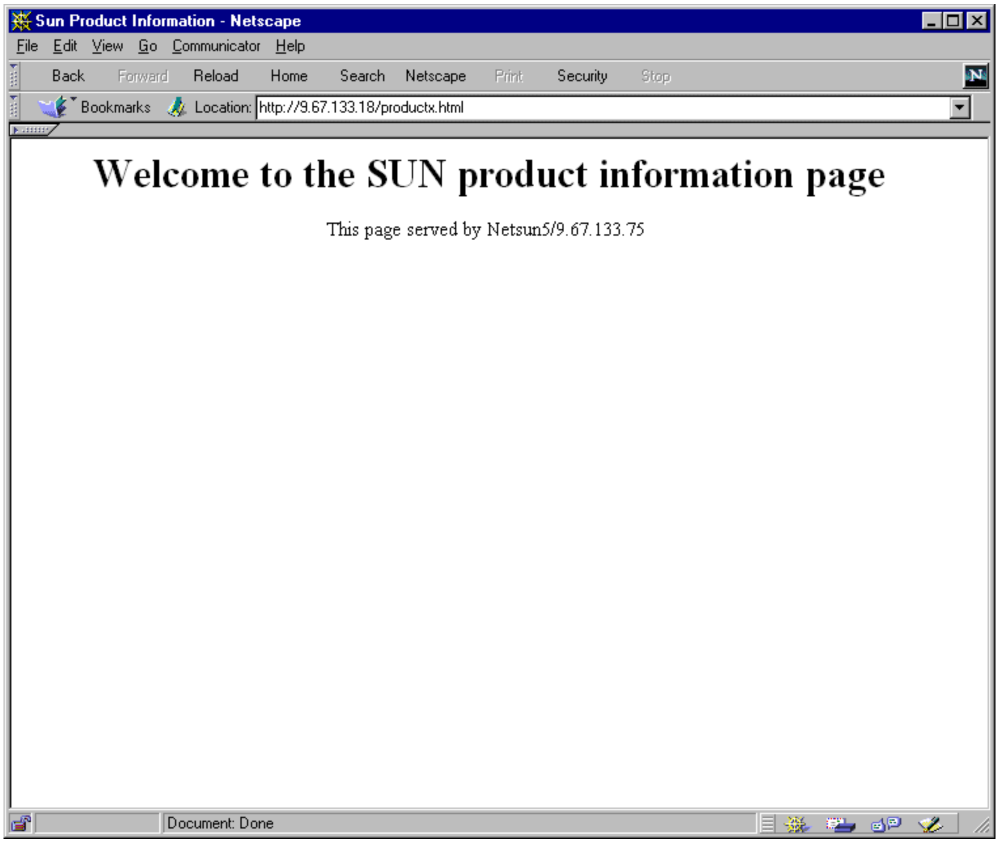
Sun Product Information - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks
Location: http://9.67.133.18/productx.html
Document: Done

Welcome to the SUN product information page

This page served by Netsun5/9.67.133.75

Figure 199. Sun Product Request

To verify that our productpage rule was matched by this request, we used the `cbrcontrol` persistent session command

```
rule rep ::
```

The command line version of this is:

```
cbrcontrol rule report cluster:port:rule
```

but as with other `cbrcontrol` commands, short forms of the keywords can be used. In order to specify that you would like the report to include all clusters, ports and rules, use only the two colon (`::`) delimiters on the command line. The command and its output are shown in the following figure:

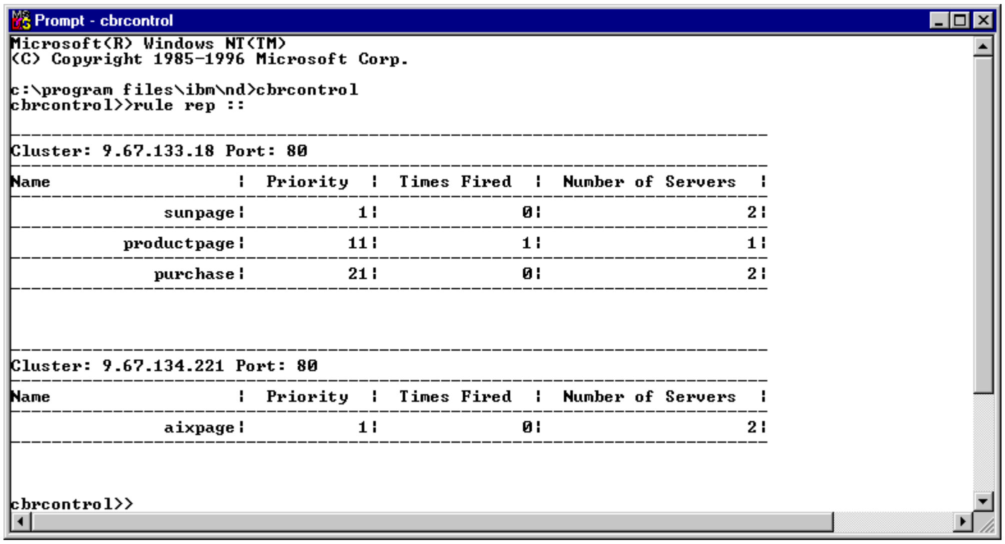

```
Microsoft(R) Windows NT(TM)
(C) Copyright 1985-1996 Microsoft Corp.

c:\program files\ibm\nd>cbrcontrol
cbrcontrol>>rule rep ::

---------------------------------------------------------------------------
Cluster: 9.67.133.18 Port: 80
---------------------------------------------------------------------------
Name                   |  Priority  |  Times Fired  |  Number of Servers  |
---------------------------------------------------------------------------
               sunpage|           1|              0|                    2|
---------------------------------------------------------------------------
           productpage|          11|              1|                    1|
---------------------------------------------------------------------------
              purchase|          21|              0|                    2|
---------------------------------------------------------------------------


---------------------------------------------------------------------------
Cluster: 9.67.134.221 Port: 80
---------------------------------------------------------------------------
Name                   |  Priority  |  Times Fired  |  Number of Servers  |
---------------------------------------------------------------------------
               aixpage|           1|              0|                    2|
---------------------------------------------------------------------------


cbrcontrol>>
```

Figure 200. Rule Report Showing productpage Rule Fired One Time

Since the sticky time is set to 0 seconds, Web requests from the same client should not stick to the same Web server, and subsequent requests should normally load balance between the two Web servers defined in cluster A. If the sticky time were set to a positive number of seconds, CBR would choose one of the two servers the first time a client request arrives that matches the productpage rule, and then would redirect all the client requests coming from the same IP address to the same server until the sticky time expires.

In this case, however, we forced the CBR server to redirect all the requests to the Web server having IP address 9.67.133.75, because this was the only server available in the productpage rule.

17.3.6.2 Cookie Affinity Demonstration

The next test we made was to demonstrate the use of Cookie affinity.

Web cookies are simple pieces of information passed between the client Web browser and the Web server during an HTTP transaction. Cookies do not contain any information about the client that the server does not already know and they cannot do anything on the client machine that the client itself cannot already do, provided the browser is within the specifications. Cookies were introduced as an answer to a fundamental problem of the Web's underlying HTTP 1.0 protocol: the lack of a state, or a persistent connection.

The first time a client accesses a Web site that serves cookies, the chosen server sends a cookie to the client browser, identifying the server and some information about which URLs the cookie is good for. The next time the client visits one of those URLs, the browser includes the cookie in its request. As long as the client's future requests contain the cookie, and each request arrives within the sticky time interval, the client will maintain affinity with the initial server.

Prior to making a request for the purchase page that will trigger the 30 second affinity rule, we examined the contents of the local cookie file (cookies.txt) in the browser directory on our client machine and it was empty:

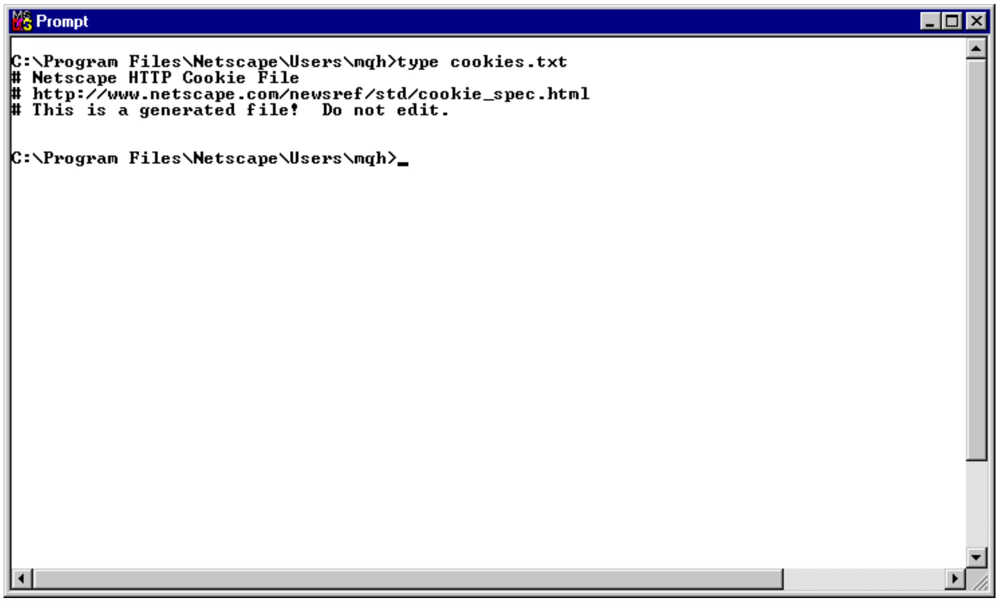

```
C:\Program Files\Netscape\Users\mqh>type cookies.txt
# Netscape HTTP Cookie File
# http://www.netscape.com/newsref/std/cookie_spec.html
# This is a generated file!  Do not edit.


C:\Program Files\Netscape\Users\mqh>_
```

Figure 201. Empty Cookie File Before Request Was Made

Recall that we had enabled a 30 second cookie affinity on our purchase rule by setting the sticky time to `30`, and setting the affinity to **Cookie** when we created the rule. This also could have been set with the command:

```
cbrcontrol rule set
```

Once a server is selected by CBR to respond to our request, subsequent requests were also responded to by the same server.

In the 30 second period we made three requests for `http://9.67.133.18/purchase.html`. Each request was responded to by the same server:

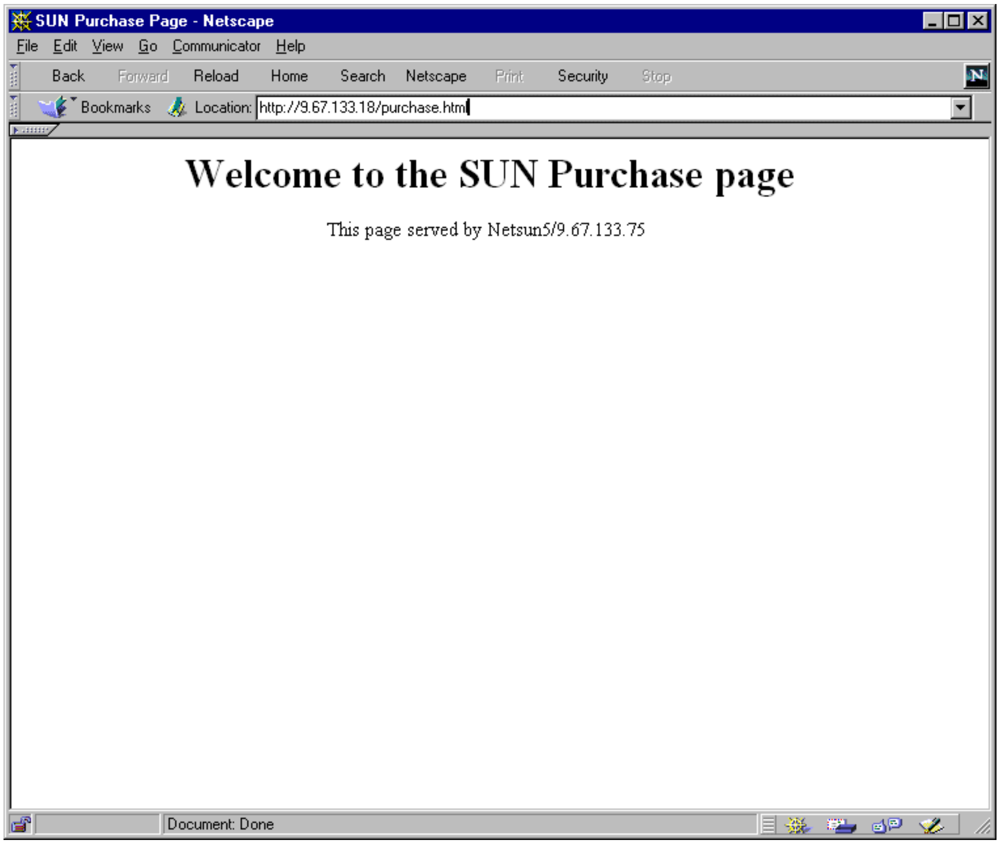
SUN Purchase Page - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://9.67.133.18/purchase.html
Document: Done

Welcome to the SUN Purchase page

This page served by Netsun5/9.67.133.75

Figure 202. Sun Purchase Using a Cookie

We verified that the cookie was set on our browser machine by examining the cookie.txt file:

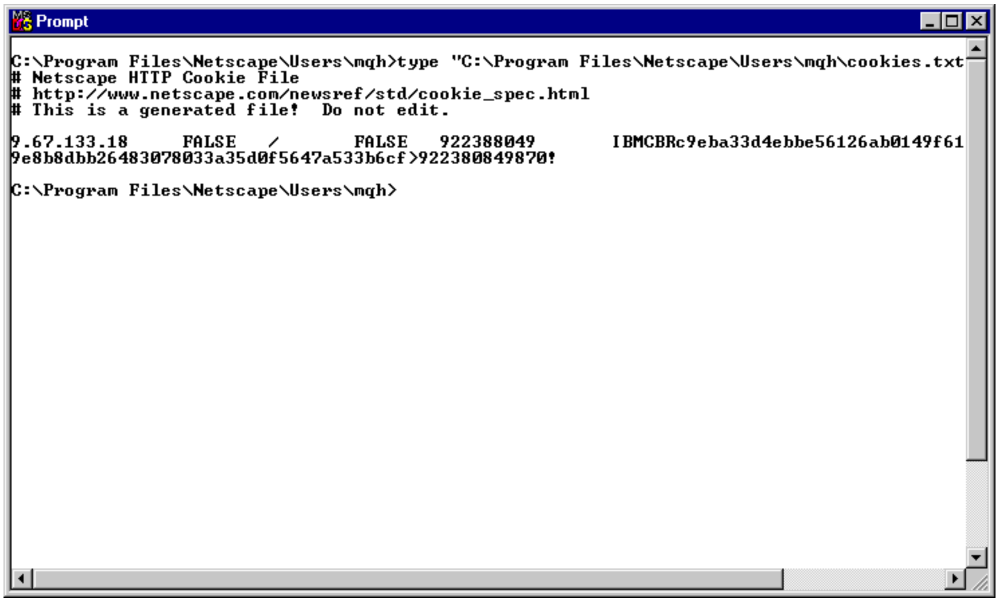

```
C:\Program Files\Netscape\Users\mqh>type "C:\Program Files\Netscape\Users\mqh\cookies.txt
# Netscape HTTP Cookie File
# http://www.netscape.com/newsref/std/cookie_spec.html
# This is a generated file!  Do not edit.

9.67.133.18     FALSE   /       FALSE   922388049       IBMCBRc9eba33d4ebbe56126ab0149f61
9e8b8dbb26483078033a35d0f5647a533b6cf>922380849870!

C:\Program Files\Netscape\Users\mqh>
```

Figure 203. Cookie File Containing one Cookie after Request Was Made

Even though the contents of the cookie are not understandable by us, they are meaningful to the Web server and browser.

After the 30 second sticky time expired, we saw that the CBR server was allowed to choose the other Web server to satisfy the request from the client.

17.4 WTE CacheByIncomingUrl Directive

A new WTE directive has been added for use with CBR in the ibmproxy.conf configuration file. This directive, `CacheByIncomingUrl`, specifies whether to use the incoming URL or the outgoing URL as the basis for generating cache file names. The values of this directive can be `on` or `off`:

- If `CacheByIncomingUrl` is set to `on`, the incoming URL will be used to generate the cache file name. In other words, when this directive is set to `on`, WTE keeps the original URL and uses it to cache the page that it gets back from the ND-changed URL.
- On the other hand, if `off` is specified, CBR rule matching and load balancing will be done on the incoming URL and the resulting URL will be used to generate the cache name. In this case, WTE uses the ND-changed URL to cache the page.

The `CacheByIncomingUrl` directive's default value is `off` in the ibmproxy.conf file.

Notice that `CacheByIncomingUrl` is a *hidden directive* of WTE. In other words, it is possible to alter its value only by manually editing the configuration file ibmproxy.conf; there is no way to change the value of this directive through the Configuration and Administration Forms.

Notice also that a Web page is cached only if WTE decides it should be. WTE does this by looking at the `Expires` tag. If there is not one, it estimates an expiration time by the `Last-Modified` parameter. You should keep this in mind if you are testing CBR caching with sample pages you have just created and which do not contain `Expires` header information. In this case, recently created pages would not be cached, since they would be interpreted by WTE as frequently changed pages, therefore resulting in an apparent failure of the CBR caching function.

Chapter 18. Remote Cache Access

It is not uncommon for a proxy server to receive more traffic than it can handle. One solution is to distribute the traffic between multiple, load balanced proxy servers. However, in this situation, the content of one cache is likely to overlap with the contents of the other caches. Besides unnecessary redundancy, this situation also requires additional bandwidth, because each copy in each server is fetched fresh from the origin. This problem can be minimized by chaining a hierarchy of proxies together, but this still results in additional traffic passing through a given server, and each additional link in the chain adds latency and increases the possibility of a failure. Though useful, proxy chaining does not solve the problem as it does not allow sharing the cache across multiple systems.

IBM Web Traffic Express (WTE) Version 2, the Caching and Filtering component of IBM WebSphere Performance Pack Version 2, has a feature called Remote Cache Access (RCA), which permits a number of WTE servers to share their cache. Eliminating duplicate objects results in bandwidth savings because:

- Objects are not fetched multiple times.
- A larger, combined logical cache yields a higher hit rate.

RCA, which was not available in the previous release of IBM WebSphere Performance Pack, allows multiple proxy servers to cooperate to form a cache array. RCA uses the Cache Array Routing Protocol (CARP) to determine which peer in the array should process the incoming request. If the requested file is not in the proxy's cache, it will query the other peers in the array to determine which proxy might have the object cached. Therefore, using RCA, multiple proxy servers can distribute the cache contents across their combined, logical cache to improve hit rates and reduce redundancy of cached content.

To run a RCA, WTE demands a shared file system. This can be IBM AFS Enterprise File System (AFS), DFS, NFS, the Windows NT file system, or any other, but the cache directories must have read and write permissions for all the cache array members.

18.1 How RCA Works

In the WTE cache array, all the machines know the existence of the others, and all run a special hash function that maps every URL to an address in the cache. An address in the cache is defined as the pair:

```
(array_member, cache_file_name)
```

This way, the cache space is the combination of all members (WTE servers) and all cache directories. Each URL has a unique location in the cache, and the hash function allows each member to compute the member that owns any given URL.

If a request is made to one proxy, this proxy will first check its local cache. If the object is not found, then it will compute the owner of the URL and ask the owner if it has the object already cached. If so, the first proxy will directly access the cache space of the other proxy, retrieve the requested object and serve it to the requesting client. If not, the requested object will be retrieved directly from the remote content server.

The following diagram shows a graphical representation of an RCA flow:

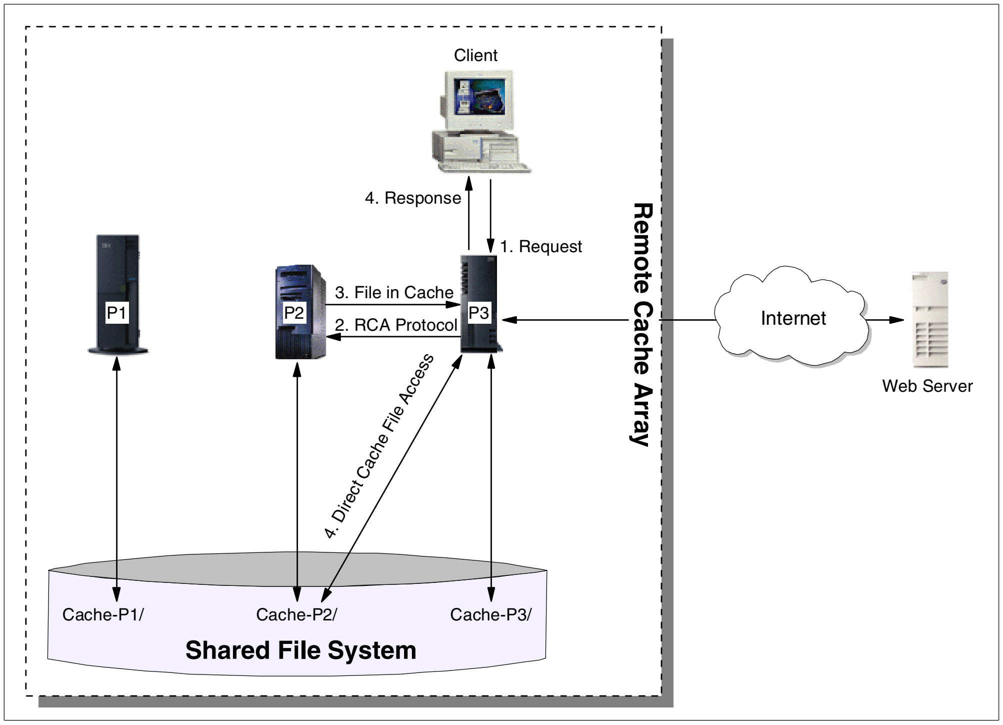

Figure 204. RCA Information Flow

The information flow on RCA is as follows:

1. The client makes a URL request to one of the proxies. In our example, this proxy is P3. Of course, the configuration of the figure above would make much more sense if a load balancing machine were placed between the client and the WTE proxy servers, to load balance the traffic between the proxy servers. Although this is not a mandatory requirement, it is a very good idea to distribute the workload between multiple WTE proxy servers using RCA by using IBM SecureWay Network Dispatcher (ND), the Load Balancing component of IBM WebSphere Performance Pack. In particular, ND offers the WTE Advisor to perform peak load management, as we will see in Chapter 19, “Peak Load Management” on page 285.
2. Using the hash function, P3 knows that this URL may be cached only in P2.
3. P3 makes an RCA request to P2 asking if it has already cached the requested object.
4. At this point, two situations can happen:
   - P2 responds that it has the URL in its cache:
     1. P3 directly accesses P2’s cache to pick up the file corresponding to the requested URL. P2 will not pass all the content of the cached object to P3. In fact, since P3 knows that P2 owns the URL, and since P2’s cache directory is accessible to P3 via the shared file system, P3 accesses P2’s cache directory directly.
     2. P3 returns the cached object to the requesting client
   - P2 responds that it does not have the requested object in its cache:

1. P3 will grab the URL in its original source, the Web or FTP server.
2. P3 will create the cache entry of the URL in P2's cache directory and use RCA to tell P2 about the modification.

Using this procedure, RCA is guaranteed to be an extremely fast protocol, and provides a way to build a scalable array.

18.2 Planning for RCA

When planning to use RCA, you should remember the following:

- The participating proxy servers should be as close as possible, with high-bandwidth connections. It is recommended that members are on the same LAN segment, or communicate over an SP switch, if at all possible. In addition, consider dedicating a network segment just to the inter-node communication (the RCA messages and the shared file system).
- The same file system must be used by all members of an array. You cannot use, for example, AFS on some nodes, NFS on others and the Windows NT file system on others. The recommended file system is AFS.
- Membership in the RCA array should be long-term. The configuration should be as stable as possible.
- Proxy servers should have similar capabilities (for example, CPU, memory, size, cache size).
- Network outages between members of the array must be infrequent.
- There should be less than 100 members in any array.
- All members of the array must run WTE Version 2. You cannot share a cache between WTE Version 2 and Distributed Web Traffic Express (DWTE) Version 1.1[1].
- The members of the array should be load-balanced using ND.

Conversely, RCA is not appropriate when:

- The participating proxy servers are not in close proximity.
- Frequent network outages are expected.
- Servers differ widely in capacity. For example, a PC server with a 500 MB cache and an RS/6000 server with a 20 GB cache should not be in the same array.
- Membership in the RCA array is short term.

18.3 RCA Scenario

Installing and configuring an RCA environment is a typical case where you can install, configure and integrate all of the WebSphere Performance Pack components. In fact, RCA nodes need to share the same file system, and by using AFS you can take advantage of the powerful file sharing features of WebSphere Performance Pack. Moreover, as we have already said, a scenario involving multiple proxy servers sharing the same cache makes more sense if the

[1] DWTE is another IBM caching and filtering proxy server. DWTE also implements RCA. Originally, IBM implemented RCA only in DWTE. Then, the RCA feature was ported also to WTE V2.0.

workload is distributed between the servers. For this purpose, you can use the ND component of IBM WebSphere Performance Pack.

In this section, we show you how we implemented an RCA scenario similar to the one represented below:

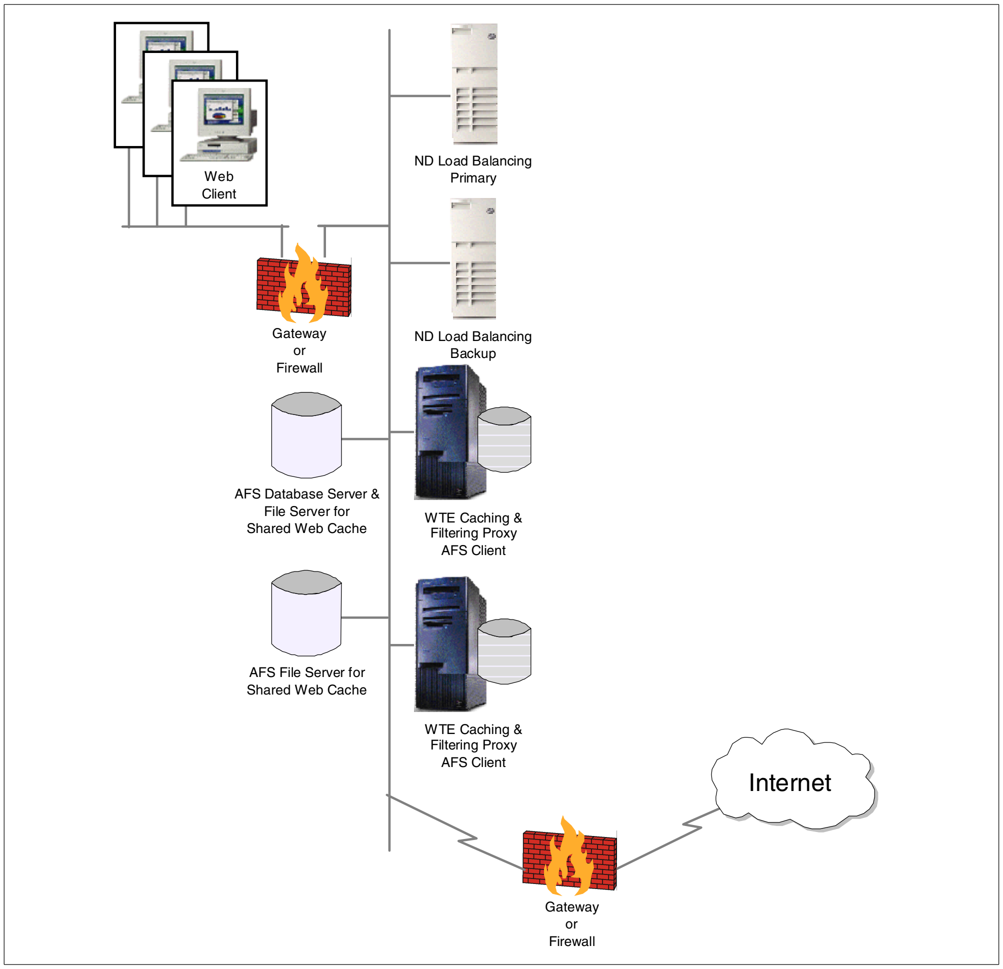
Web Client
ND Load Balancing Primary
ND Load Balancing Backup
Gateway or Firewall
AFS Database Server & File Server for Shared Web Cache
WTE Caching & Filtering Proxy AFS Client
AFS File Server for Shared Web Cache
WTE Caching & Filtering Proxy AFS Client
Internet
Gateway or Firewall

Figure 205. Graphical Representation of the RCA Scenario

18.3.1 Scenario Implementation

1. We constructed an AFS cell called cuzcuz.itso.ibm.com following the steps described in *IBM WebSphere Performance Pack: Web Content Management with IBM AFS Enterprise File System*, SG24-5857.

   One of the most useful characteristics of AFS is server ReadOnly volume replication, which means that we can have several copies of the same volume replicated across a number of file servers. This way, clients can access the closest server. Another advantage is that the workload will be balanced through all the servers. But in the case of RCA, the cache directory needs to be updated continuously, so the volumes must be accessed through the ReadWrite path. Even for a remote read access from another proxy, the ReadWrite volume must be available, because a ReadOnly volume is created and replicated manually and statically.

2. We created a separate volume for each proxy, but we didn't replicate those volumes.
3. We configured all the proxy servers as clients of the AFS cell.
4. We granted read and write permissions for all proxies when accessing any other proxy cache directory.

We set up two proxy AIX machines, pluto.itso.ral.ibm.com and rs600030.itso.ral.ibm.com. These machines had essentially the same configurations.

We also configured RCA. This can be done by either accessing the Configuration and Administration Forms or by including the following lines in the ibmproxy.conf configuration file of both the WTE machines:

```
Version    RCA/1.0
ArrayName Cake
Member rs600030 {
        RCAAddr         rs600030.itso.ral.ibm.com
        RCAPort         6969
        CacheSize       100M
        CacheRoot       /afs/.cuzcuz.itso.ibm.com/proxy-itso/cache-rs600030
        Timeout         1000 milliseconds
        BindSpecific    On
        ReuseAddr       Off
}

Member pluto {
        RCAAddr         pluto.itso.ral.ibm.com
        RCAPort         6969
        CacheSize       100M
        CacheRoot       /afs/.cuzcuz.itso.ibm.com/proxy-itso/cache-pluto
        Timeout         1000 milliseconds
        ReuseAddr       Off
}
```

Figure 206. RCA Configuration in ibmproxy.conf

Figure 206 on page 283 shows the configuration for an array called Cake. We can see that the cache directory for each proxy is inside the same AFS cell, and is under the ReadWrite path, .cuzcuz.itso.ibm.com. Through these lines in the WTE configuration file, both proxies know what the cache directories are, and can access them directly. The `RCAPort` directive tells each proxy which port it should connect to, and how to communicate with the other members.

It is very important that the RCA block of the configuration be identical in all the array members. Even the order in which members appear in the RCA block should not differ. One way to make this easier is to store the RCA configuration in a separate file, located somewhere in the shared file system (even on a replicated volume, if using AFS), and then including this configuration using the `RCAConfigFile` directive in the WTE ibmproxy.conf configuration file.

After restarting the WTE servers, we accessed the WTE Configuration and Administration Forms of both the WTE servers, and we verified that the forms reflected the configuration we had issued:

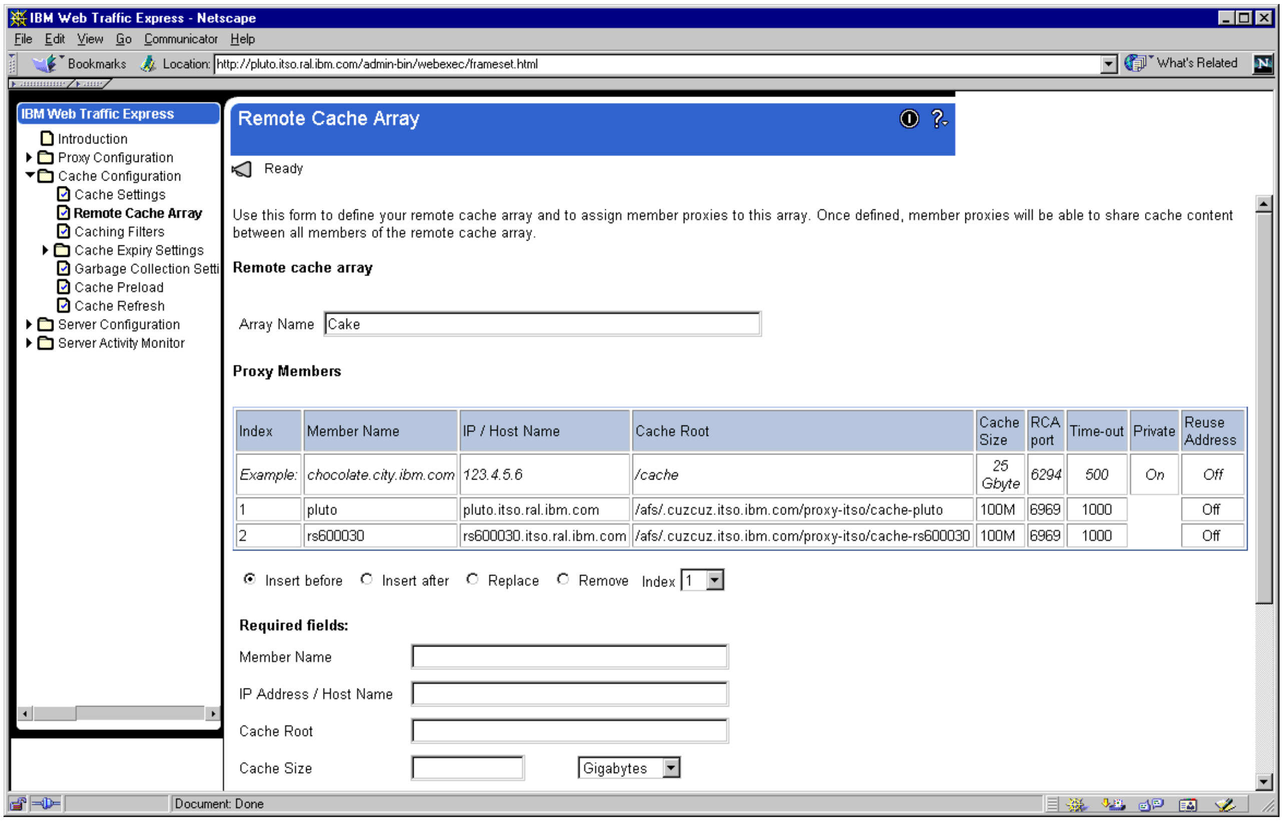

Remote Cache Array

Ready

Use this form to define your remote cache array and to assign member proxies to this array. Once defined, member proxies will be able to share cache content between all members of the remote cache array.

Remote cache array

Array Name: Cake

Proxy Members

Index	Member Name	IP / Host Name	Cache Root	Cache Size	RCA port	Time-out	Private	Reuse Address
Example:	*chocolate.city.ibm.com*	*123.4.5.6*	*/cache*	*25 Gbyte*	*6294*	*500*	*On*	*Off*
1	pluto	pluto.itso.ral.ibm.com	/afs/.cuzcuz.itso.ibm.com/proxy-itso/cache-pluto	100M	6969	1000		Off
2	rs600030	rs600030.itso.ral.ibm.com	/afs/.cuzcuz.itso.ibm.com/proxy-itso/cache-rs600030	100M	6969	1000		Off

● Insert before ○ Insert after ○ Replace ○ Remove Index 1

Required fields:

Member Name

IP Address / Host Name

Cache Root

Cache Size Gigabytes

Figure 207. RCA Configuration in the Web Interface

A simple test permitted us to demonstrate that RCA was functioning correctly.

Chapter 19. Peak Load Management

In WebSphere Performance Pack Version 2, a new Advisor, known as the WTE Advisor for SecureWay Network Dispatcher (ND), enhances load balancing management between multiple Web Traffic Express (WTE) nodes by preventing the ND server from sending new requests to a WTE node that is engaged in garbage collection or cache refresh. This is a new feature implemented for the first time in WebSphere Performance Pack Version 2.

Using the WTE Advisor, an ND server has the ability to detect and react to sudden increases in activity on a WTE node belonging to a cluster. The new WTE Advisor on ND opens a connection to the WTE node, sends a *WTE-specific* HTTP GET request, and interprets the response as a WTE load.

Notice that the WTE Advisor integrates the load balancing component with the caching and filtering component of WebSphere Performance Pack. This new feature is called *peak load management* and is part of the quality-of-service enhancements offered by IBM WebSphere Performance Pack Version 2.

19.1 How Peak Load Management Works

Peak load management is a new feature implemented for the first time in IBM WebSphere Performance Pack Version 2. You can perform peak load management by activating the WTE Advisor on a Dispatcher machine that is load balancing the workload between two or more WTE proxy servers.

The WTE Advisor is a process running on the Dispatcher machine, but is specific for WTE. This process will not work if the load balanced proxy server is not WTE. When the WTE Advisor runs on a Dispatcher machine, this machine sends a specific HTTP request to the load balanced WTE servers, and these reply back with four pieces of information:

- A boolean saying if the cache is reloading
- A boolean saying if garbage collection is running
- The number of active threads
- The maximum number of threads

The WTE Advisor is able to understand these pieces of information passed to it by the WTE machines, and feeds them back to the Dispatcher Manager. The Manager uses the data retrieved by the WTE Advisor to modify the weights associated with the WTE proxy servers, if necessary, in proportion to the importance given by the Dispatcher administrator to the Advisor’s input.

This way, the Dispatcher can detect when a WTE proxy server is performing some CPU-intense activities better than it could if the WTE Advisor were not running. The Dispatcher takes into account all these resource-intensive activities and makes its load balancing decision appropriately, as we demonstrate in the following section.

19.2 Peak Load Management Scenario

In this section, we show you a basic configuration for an ND machine load balancing multiple WTE proxy server nodes. We built the network with five workstations: one client, one Web server, two WTE proxy servers and one ND machine load balancing the traffic between the two WTE servers. We installed the components as shown in the following diagram:

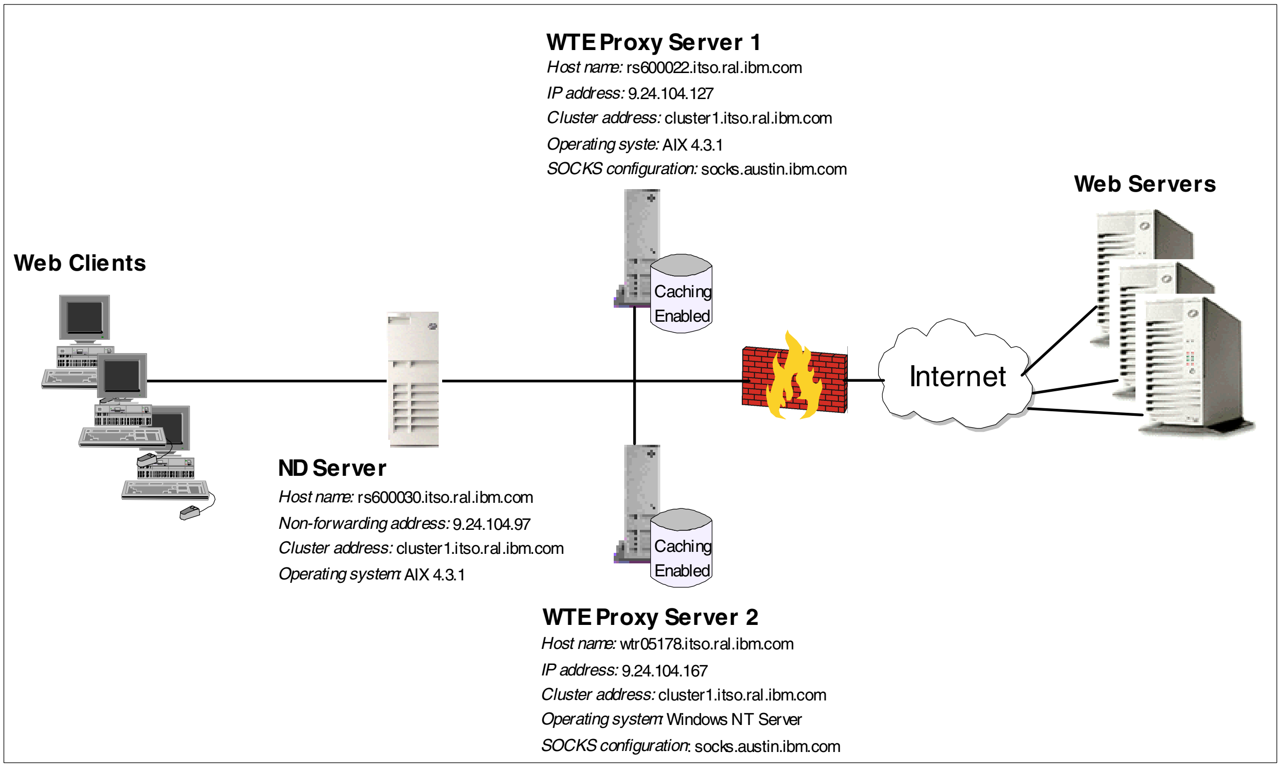
WTE Proxy Server 1
Host name: rs600022.itso.ral.ibm.com
IP address: 9.24.104.127
Cluster address: cluster1.itso.ral.ibm.com
Operating syste: AIX 4.3.1
SOCKS configuration: socks.austin.ibm.com
Web Servers
Web Clients
Caching Enabled
Internet
ND Server
Host name: rs600030.itso.ral.ibm.com
Non-forwarding address: 9.24.104.97
Cluster address: cluster1.itso.ral.ibm.com
Operating system: AIX 4.3.1
Caching Enabled
WTE Proxy Server 2
Host name: wtr05178.itso.ral.ibm.com
IP address: 9.24.104.167
Cluster address: cluster1.itso.ral.ibm.com
Operating system: Windows NT Server
SOCKS configuration: socks.austin.ibm.com

Figure 208. Peak Load Management Scenario

Notice that the diagram above tries to reflect a real-life situation, where multiple clients and multiple Web servers come into the scenario.

The following table provides more details on the WTE cluster configuration:

Table 19. WTE Cluster Configuration

Machine Role	Host Name	IP Address	Operating System
Dispatcher	rs600030	9.24.104.97	AIX 4.3.1
	cluster1	9.24.104.156	
WTE proxy server 1	rs600022	9.24.104.127	AIX 4.3.1
WTE proxy server 2	wtr05178	9.24.104.167	Windows NT Server 4.0

All the machines above belonged to the domain itso.ral.ibm.com. Notice that 9.24.104.156 was the cluster address in this configuration. The host name associated with this cluster address was cluster1.itso.ral.ibm.com.

All the machines shown in Figure 208 on page 286 were provided with a token-ring interface and connected to the same LAN.

The software we used in our scenario is described in the following list:

- Netscape Navigator 4.5 was the Web browser running on the Web client machine.

- The load balancing function was provided by the Dispatcher component of ND Version 2.1.
- The proxy server function was provided by WTE Version 2.0.
- The Web server function was provided by Lotus Domino Go Webserver V4.6.2.5.

19.2.1 Dispatcher Configuration

For the installation and configuration of the ND component of IBM WebSphere Performance Pack, refer to *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858. We configured the Dispatcher using the ND configuration graphical user interface (GUI).

The following diagram shows the Dispatcher GUI after the configuration was completed:

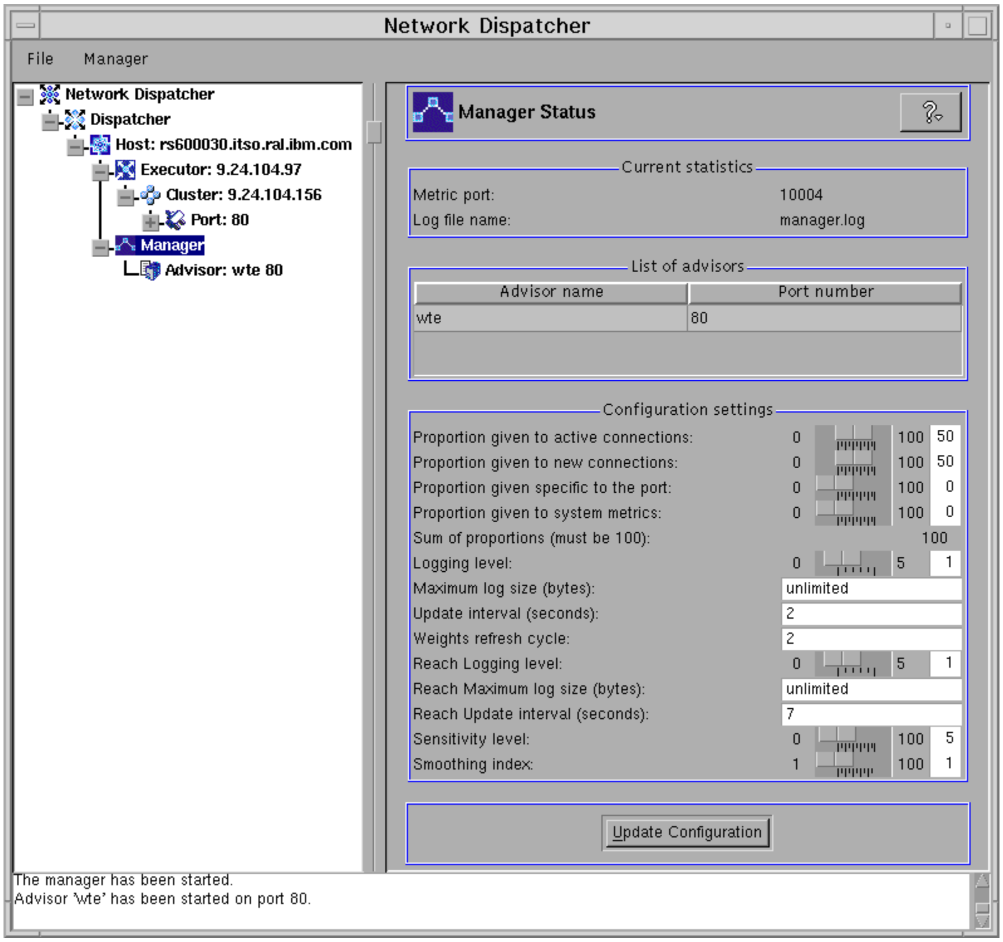

Network Dispatcher
Dispatcher
Host: rs600030.itso.ral.ibm.com
Executor: 9.24.104.97
Cluster: 9.24.104.156
Port: 80
Manager
Advisor: wte 80

Manager Status

Current statistics

Metric port:	10004
Log file name:	manager.log

List of advisors

Advisor name	Port number
wte	80

Configuration settings

Proportion given to active connections:	0	100	50
Proportion given to new connections:	0	100	50
Proportion given specific to the port:	0	100	0
Proportion given to system metrics:	0	100	0
Sum of proportions (must be 100):			100
Logging level:	0	5	1
Maximum log size (bytes):			unlimited
Update interval (seconds):			2
Weights refresh cycle:			2
Reach Logging level:	0	5	1
Reach Maximum log size (bytes):			unlimited
Reach Update interval (seconds):			7
Sensitivity level:	0	100	5
Smoothing index:	1	100	1

```
The manager has been started.
Advisor 'wte' has been started on port 80.
```

Figure 209. Network Dispatcher GUI - WTE Cluster with WTE Advisors

As you can see from the figure above, we defined a cluster of two WTE proxy servers, load balanced by an ND machine. In this cluster, we also activated the WTE Advisor on port 80. Advisors can be activated as explained in *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858. In our configuration, you can see that the Dispatcher was configured to load balance the traffic to the two WTE servers giving 40% of importance to the active connections, 40% of importance to the new connections, and 20% of importance to the feedback coming from Advisors (in this case, from the WTE Advisor). In this scenario, no importance was given to the input coming from system monitoring tools, such as ISS, because no system monitoring tool was part of the configuration.

It would not be appropriate to give a high importance to the Advisors’ inputs, but in this case we preferred to set the importance to 20% as we wanted to see the real effects of the WTE Advisors in the load balancing decisions made by the Dispatcher.

19.2.2 Configuration of the WTE Proxy Servers

No special configurations other than a basic one are required on the WTE servers when the WTE Advisor is running on the load balancing Dispatcher machine.

19.2.3 Testing the Peak Load Management Scenario

With the configuration described above, we performed some tests to verify how the Dispatcher reacts to a sudden activity increase on one of the WTE machines.

19.2.3.1 Load Balancing Basic WTE Activities

On the Web client machine, we configured the Web browser to access the Internet through the proxy server 9.24.104.156. This was the cluster address of our WTE proxy server cluster (see Table 19 on page 286). Then, we requested multiple URLs from the clients simultaneously. From the Dispatcher Server monitor GUI, we could see that the two WTE proxy servers were in use simultaneously as shown in the following figure:

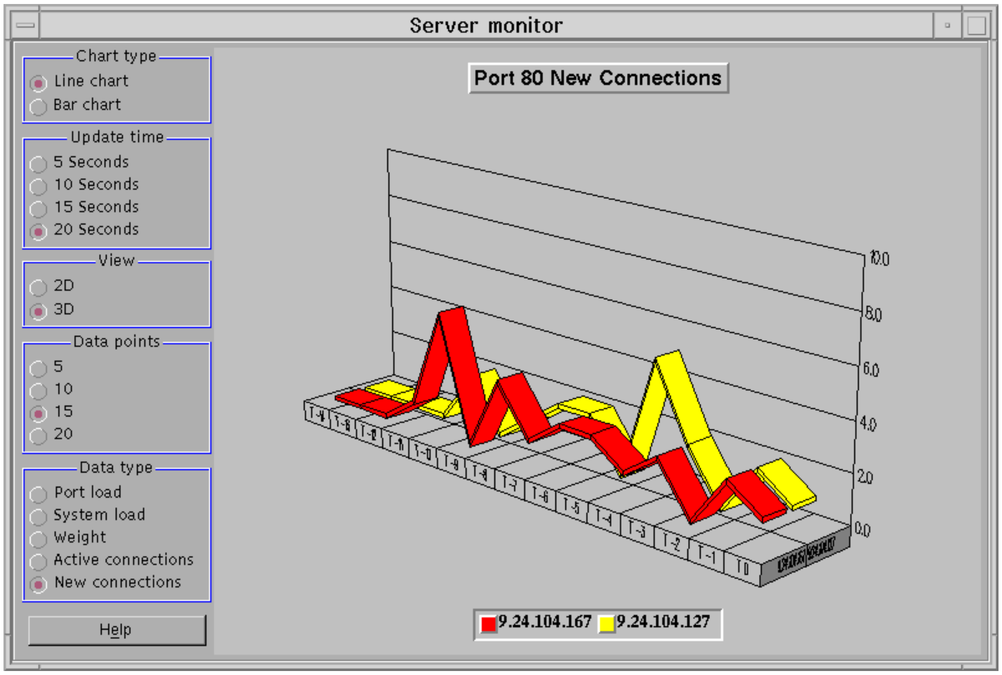

Server monitor

Chart type
● Line chart
○ Bar chart

Update time
○ 5 Seconds
○ 10 Seconds
○ 15 Seconds
● 20 Seconds

View
○ 2D
● 3D

Data points
○ 5
○ 10
● 15
○ 20

Data type
○ Port load
○ System load
○ Weight
○ Active connections
● New connections

Figure 210. Network Dispatcher GUI - WTE Proxy Servers Running

19.2.3.2 Consequences of the WTE Advisor Activation

As mentioned earlier, the WTE Advisor for the ND enhances load management on multiple WTE servers by preventing the Dispatcher from sending new requests to a WTE node that is currently engaged in garbage collection or cache refresh. This function is very useful, since garbage collection and cache refresh are very CPU-intensive activities for a WTE server, and it is very important that the Dispatcher have a way to detect when a WTE server is using most of its resources for these resource-consuming activities.

The following figure shows that, at the beginning of our tests, when there was no particular activity on either of the WTE proxy servers, both WTE servers had the same weights, as shown:

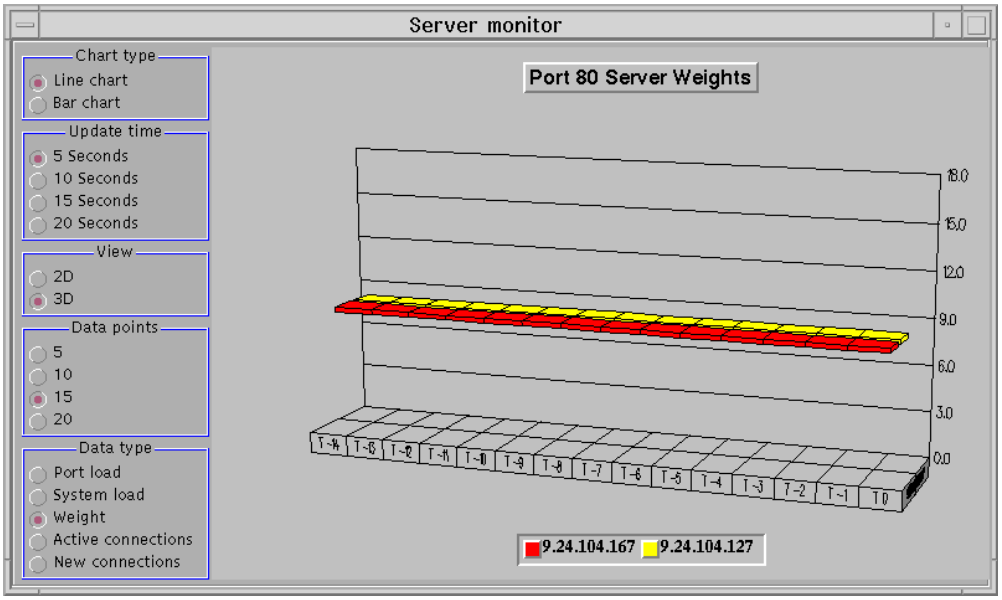

Server monitor

Chart type
● Line chart
○ Bar chart

Update time
● 5 Seconds
○ 10 Seconds
○ 15 Seconds
○ 20 Seconds

View
○ 2D
● 3D

Data points
○ 5
○ 10
● 15
○ 20

Data type
○ Port load
○ System load
● Weight
○ Active connections
○ New connections

Figure 211. Network Dispatcher GUI - WTE Proxy Servers - No Load

We then activated the cache refresh process on the WTE proxy server 1 (see Table 19 on page 286), whose hostname was rs600022. Although cache refresh can start automatically at specific times, for this test we manually forced the WTE proxy server 1 to run the cache refresh process. We did this with the command:

```
cacheagt -r /etc/ibmproxy.conf
```

Through the WTE Advisor, the Dispatcher immediately registered that cache refresh was running on the WTE proxy server 1, and the Dispatcher Manager decreased the weight of this WTE server to reduce the number of connections that were going to be forwarded to it. The Dispatcher Server monitor immediately registered this weight variation:

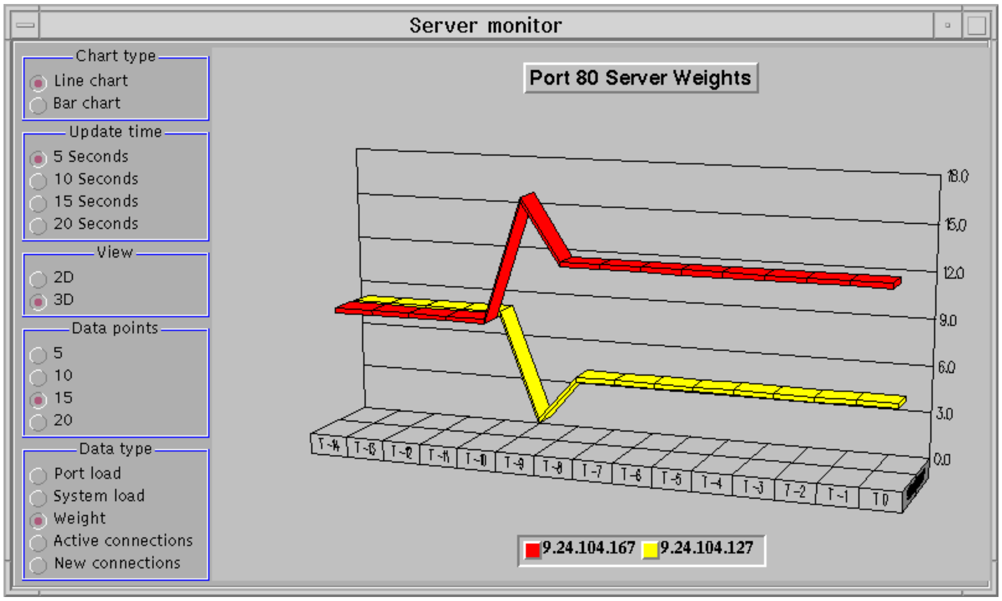

Server monitor

Chart type
● Line chart
○ Bar chart

Update time
● 5 Seconds
○ 10 Seconds
○ 15 Seconds
○ 20 Seconds

View
○ 2D
● 3D

Data points
○ 5
○ 10
● 15
○ 20

Data type
○ Port load
○ System load
● Weight
○ Active connections
○ New connections

Figure 212. Network Dispatcher GUI - WTE Proxy Servers - Cache Refresh in Progress

We can see that the weight of the WTE proxy server 1, having IP address 9.24.104.127, is now much lower than the weight of the WTE proxy server 2, having IP address 9.24.104.167. The following diagram shows the Dispatcher configuration GUI and the weight of each of the WTE proxy servers when the cache refresh process was running:

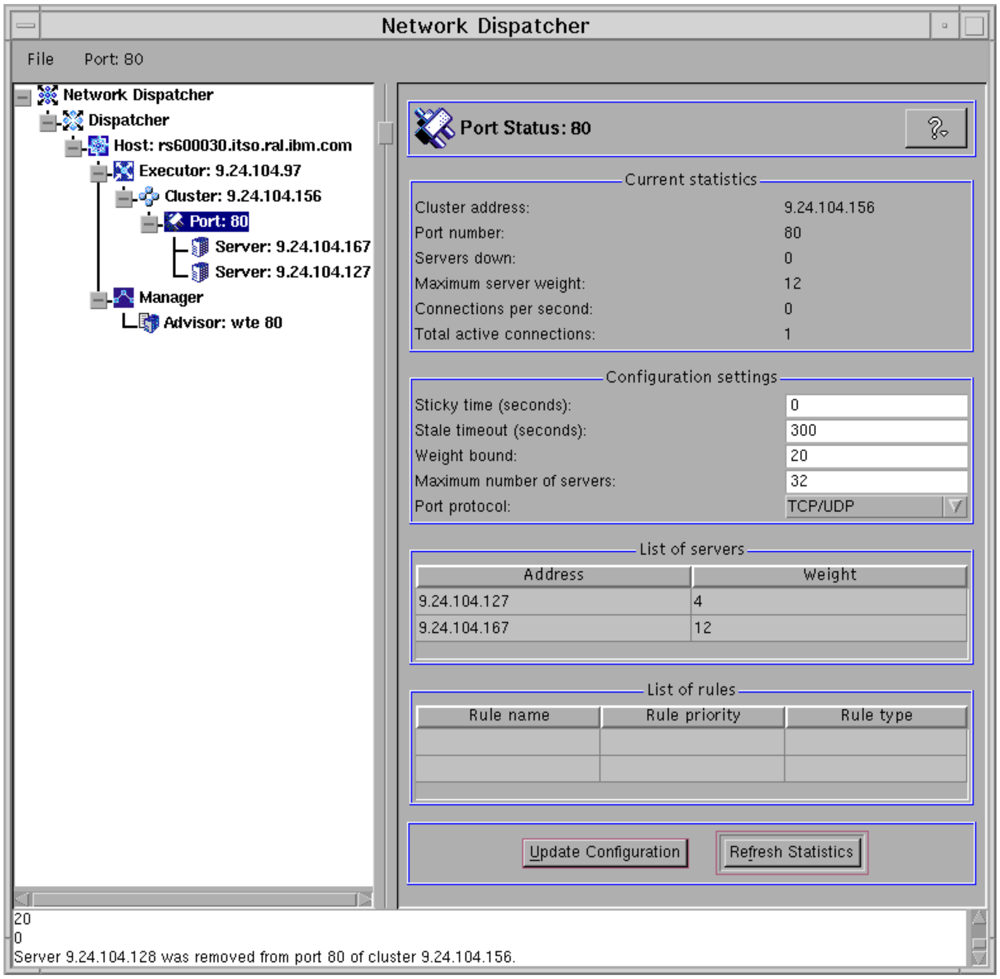

Current statistics

Cluster address:	9.24.104.156
Port number:	80
Servers down:	0
Maximum server weight:	12
Connections per second:	0
Total active connections:	1

Configuration settings

Sticky time (seconds):	0
Stale timeout (seconds):	300
Weight bound:	20
Maximum number of servers:	32
Port protocol:	TCP/UDP

List of servers

Address	Weight
9.24.104.127	4
9.24.104.167	12

List of rules

Rule name	Rule priority	Rule type

```
20
0
Server 9.24.104.128 was removed from port 80 of cluster 9.24.104.156.
```

Figure 213. Network Dispatcher GUI - Cache Refresh in Progress

This demonstration completes our scenario on the peak load management feature of IBM WebSphere Performance Pack Version 2.

Chapter 20. Common Configuration

The Common Configuration utility is a new feature of IBM WebSphere Performance Pack Version 2. In Version 2, it provides access to seven browser-based wizards that can be used to configure specific aspects of the three components of the product:

- IBM AFS Enterprise File System (AFS):
  - Getting Started with AFS
  - Add an AFS Web volume
  - Replicate an AFS Web volume
  - Move an AFS Web volume
  - Delete an AFS Web volume
- IBM Web Traffic Express (WTE):
  - Platform for Internet Content Selection (PICS) filtering wizard
- IBM SecureWay Network Dispatcher (ND):
  - Add a Network Dispatcher cluster

The wizards do not provide access to all the functions of the individual component configuration programs. Use the GUIs provided with each of the components to perform additional product-specific configuration.

20.1 Installation

The Common Configuration utility must be installed and configured as a separate component. The installation and configuration are supported on the three platforms where IBM WebSphere Performance Pack itself is supported.

20.1.1 Planning for the Installation

The first step is to choose a machine that will perform the function of *configuration server* for your environment. The Common Configuration utility allows for remote configuration of the WebSphere Performance Pack components. In other words, it is not necessary to install the Common Configuration utility on the same ND, AFS, or WTE machine you will use the Common Configuration utility for.

Access to the Common Configuration utility machine is Web based.

The following diagram shows how the Common Configuration utility works in relation to all of the other WebSphere Performance Pack components:

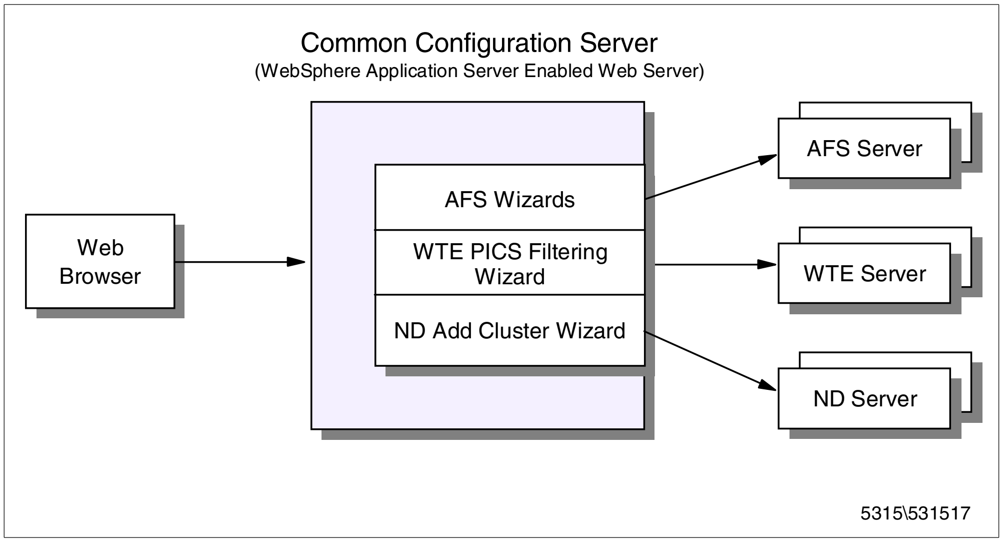
Common Configuration Server
(WebSphere Application Server Enabled Web Server)
Web Browser
AFS Wizards
WTE PICS Filtering Wizard
ND Add Cluster Wizard
AFS Server
WTE Server
ND Server
5315\531517

Figure 214. Common Configuration Server Environment

Because the Common Configuration utility is implemented by using IBM WebSphere Application Server and associated Java servlets (see 1.7.1, “IBM WebSphere Application Server” on page 23), it is necessary that your chosen configuration server machine meet these two prerequisites:

1. A Web server must be installed on the Common Configuration machine. This Web server must be supported by WebSphere Application Server on your platform. See 20.1.2, “Installing the Common Configuration Utility” on page 294 for a list of the supported Web servers on your platform.

   You need to know the document directory and the configuration file directory for whichever Web server you have installed.

2. The Web server must have the ability to serve files that are contained on one of the machine’s local file systems. It cannot be set up to serve files only from AFS.

Notice that it is not necessary to install WebSphere Application Server on the Common Configuration server machine. If WebSphere Application Server is not installed yet, the Common Configuration installation process installs the WebSphere Application Server filesets transparently during the installation routine.

Once these prerequisites are met, you can proceed to install the Common Configuration on your Common Configuration server machine.

20.1.2 Installing the Common Configuration Utility

The process of installing the Common Configuration utility with the Java InstallShield differs only slightly on the three platforms that it can be installed on. We will demonstrate the installation and note where the differences are.

We performed our installation of the Common Configuration utility on a uniprocessor IBM RS/6000 43P having 192 MB of RAM, 2.2 GB of hard disk and one token-ring interface. This RS/6000 had AIX Version 4.3.1, Java Runtime Environment (JRE) 1.1.6 and IBM HTTP Server Version 1.3.3 installed on it.

We also installed the Common Configuration utility on an IBM PC 365 running Windows NT Server4.0. This machine had Java Development Kit (JDK) 1.1.6 and IBM HTTP Server Version 1.3.3 installed on it.

One of the Web servers listed below that is supported by the WebSphere Application Server must be installed on your machine:

- On the AIX platform:
  - IBM HTTP Server V1.3.3
  - Lotus Domino Go Webserver V4.6.2
  - Netscape Enterprise V3.01
  - Netscape Enterprise V3.51
  - Apache HTTP Server V1.3.2
- On the Windows NT platform:
  - IBM HTTP Server V1.3.3
  - Lotus Domino Go Webserver V4.6.2
  - Netscape Enterprise v3.01
  - Netscape Enterprise v3.5
  - Apache v1.3.2
  - IIS 2/3.0
  - IIS 4.0
- On the Solaris platform:
  - IBM HTTP Server V1.3.3
  - Apache Server V1.3.2
  - Lotus Domino Go Webserver V4.6.2.5
  - Netscape Enterprise Server V3.01 and V3.51 (recommend V3.5.1)
  - Netscape FastTrack Server V3.01

The installation program for the Common Configuration utility makes use of Java InstallShield's setup class. For this reason you are required to pre-install the Java Virtual Machine (JVM). On AIX, only the Java Runtime Environment (JRE) is sufficient as it contains the JVM, the Java platform core classes, and supporting files. On Windows NT we found it necessary to install the entire JDK to perform the installation. We used JDK Version 1.1.6.

If Java is not installed on your machine or installed at a level lower than 1.1.6, you can find the install image for JDK 1.1.6:

- For AIX at `http://www.ibm.com/java/jdk/download/index.html` or on the IBM WebSphere Performance Pack Version 2 CD-ROM
- For Windows NT and Sun Solaris at `http://www.javasoft.com` or on the IBM WebSphere Performance Pack Version 2 CD-ROM

The installation of the JDK is described in the IBM redbook *Network Computing Framework Component Guide,* SG24-2119.

To prepare for the installation of the Common Configuration utility, follow the steps listed below:

1. Insert the IBM WebSphere Performance Pack Version 2 CD-ROM in the CD-ROM drive.
2. This step will differ on each platform:
   - On AIX, from a command line, enter the following commands:

```
mkdir /cdrom
mount -rv cdrfs /dev/cd0 /cdrom
cd /cdrom/aix
```

- On Windows NT, enter:

  ```
  E:
  cd nt
  ```

  where E is assigned to the CD-ROM drive.

- On Solaris, insert the IBM Websphere Performance Pack CD-ROM in the CD-ROM drive. The system will automatically mount the WebSphere Performance Pack CD-ROM as /cdrom/websphere. Then, from a command prompt, enter the following command:

  ```
  cd /cdrom/websphere/sun
  ```

3. To start the Java InstallShield installation program, enter:

   ```
   java setup
   ```

The first screen you will see is the Welcome window. After clicking the **Next** button, you are prompted to enter the destination location. Note that this is a working directory for Websphere Performance Pack. The default location on Windows NT is C:\WSPP; on Solaris /opt/WSPP; and on AIX /usr/lpp/WSPP, as shown in the following window:

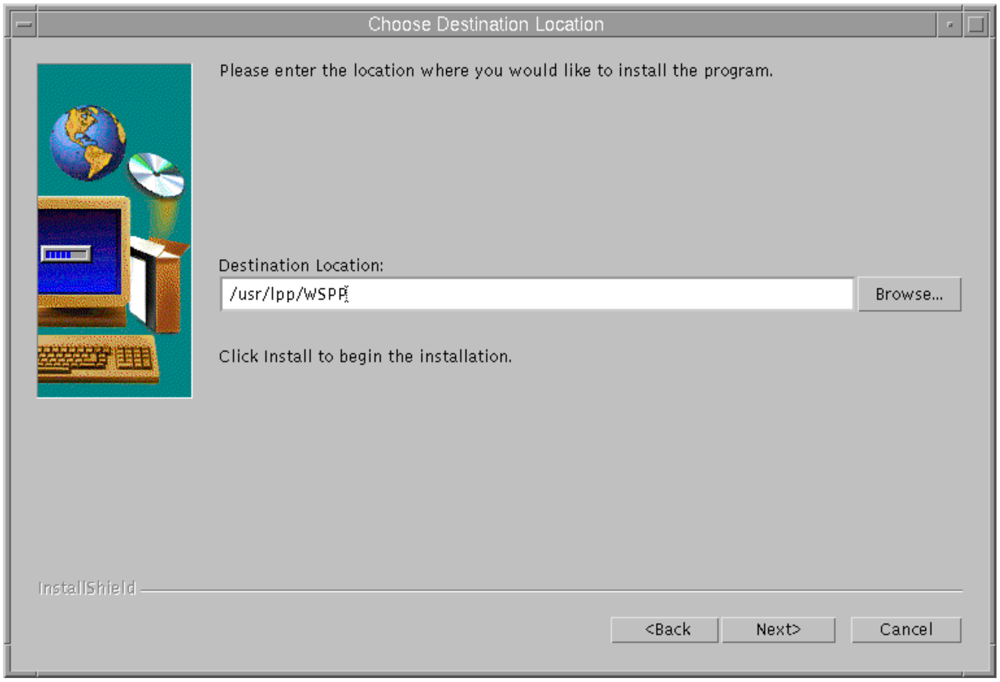

Please enter the location where you would like to install the program.

Destination Location: /usr/lpp/WSPP

Click Install to begin the installation.

Figure 215. Choose Destination Location

In the above window, you are prompted to click **Install** to proceed. The button does not exist, due to a known InstallShield for Java problem. Click the button labeled **Next** to continue instead. If the destination directory does not exist on your system, you will be presented with a question dialog asking you if you would like the location to be created. Then, you will be presented with a window allowing you to choose which IBM WebSphere Performance Pack V2 components you want to install:

- File Sharing
- Load Balancing
- Caching and Filtering
- Common Configuration

This is shown in the following figure:

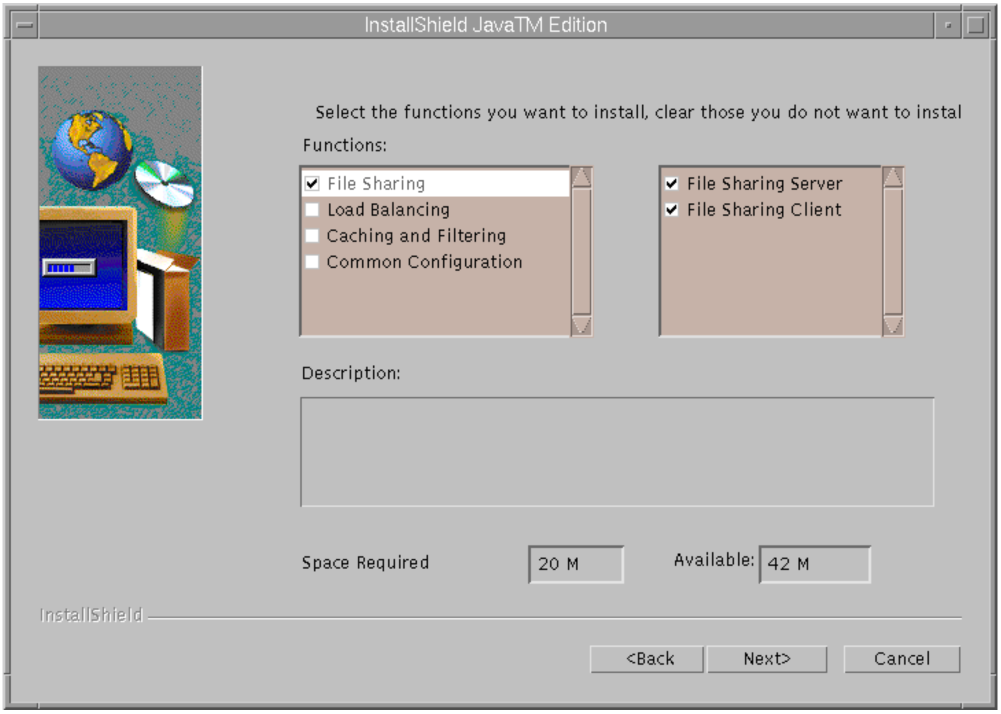

Select the functions you want to install, clear those you do not want to instal

Functions:

☑ File Sharing
□ Load Balancing
□ Caching and Filtering
□ Common Configuration

☑ File Sharing Server
☑ File Sharing Client

Description:

Space Required 20 M

Available: 42 M

Figure 216. Choose the Component to Install (accinst01)

When we selected **Common Configuration**, a description of that component was displayed as well as the space required to install it. This window also shows you how much space is available in the destination location file system.

If you do not have enough free space in your destination location file system and you click **Next**, you will see a warning window letting you know that there is not enough space to perform the installation. You will then have to use the operating system utilities to increase the size of the file system so that there is enough space for the installation to be performed. After we had successfully increased the size of the /usr file system, we were not able to continue with the Java InstallShield process. In order for the available space field (as seen in Figure 216 on page 297) to be updated to reflect the new amount of available space in the target directory, it was necessary to backup one screen by clicking the **Back** button followed by the **Next** button. Then, with the Common Configuration component selected, we were able to click the **Next** button.

In the next step, we had to provide some information to the install process. The type of information that the user is prompted for next depends on whether or not IBM WebSphere Application Server is already installed on the machine. Furthermore, in both cases the information that the user is prompted for is different for each of the supported platforms and so in the next sections we will show all of the possibilities.

20.1.2.1 Providing Information on AIX and Solaris

The information you provide during the installation depends on whether or not WebSphere Application Server is already installed on your system.

1. **WebSphere Application Server is not already installed**

   When WebSphere Application Server was not preinstalled on our AIX machine, we saw the following window:

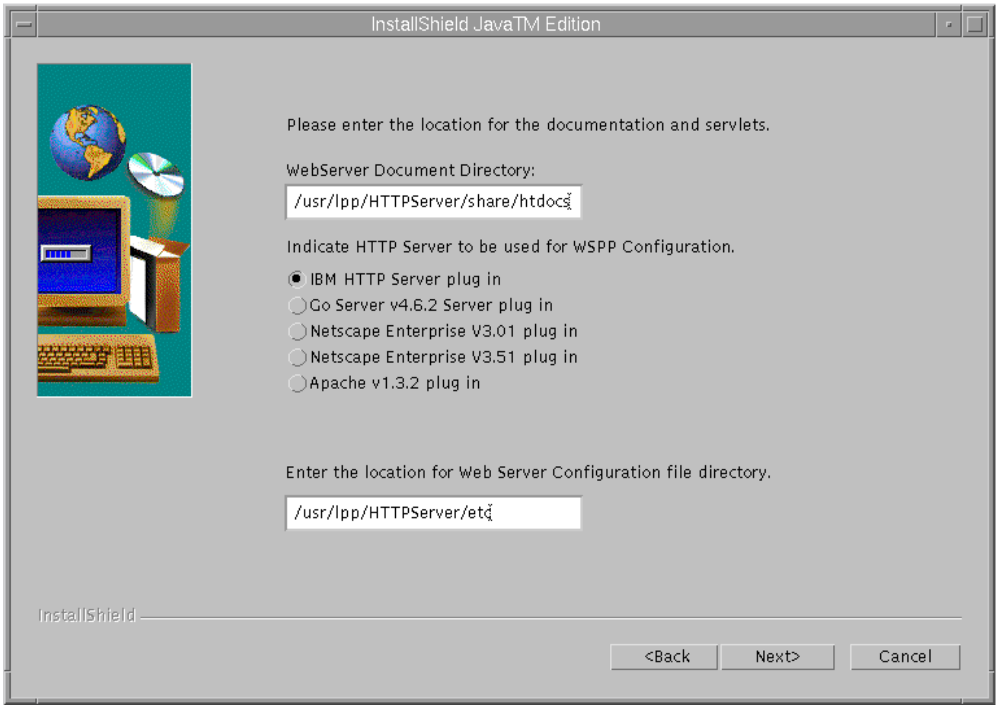

Please enter the location for the documentation and servlets.

WebServer Document Directory:

/usr/lpp/HTTPServer/share/htdocs

Indicate HTTP Server to be used for WSPP Configuration.

● IBM HTTP Server plug in
○ Go Server v4.6.2 Server plug in
○ Netscape Enterprise V3.01 plug in
○ Netscape Enterprise V3.51 plug in
○ Apache v1.3.2 plug in

Enter the location for Web Server Configuration file directory.

/usr/lpp/HTTPServer/etc

Figure 217. Request to Supply Web Server Configuration Details

   The above window will be displayed if you do not have WebSphere Application Server already installed. In order for the correct Web server plug-in to be installed, you will need to select one of the radio buttons to inform the installation process which Web server is installed on the machine.

   In addition, you are asked to provide two directories:

   1. The first piece of information the installation process asks you for is the Web server document directory. The installation process requires this information as it places some files in this directory that are used when the Common Configuration utility is launched.
   2. The second piece of information the installation process asks you for is the Web server configuration file directory. Because the installation in this case is going to install WebSphere Application Server, it must make modifications to the Web server configuration file so that the Web server will be able to work together with WebSphere Application Server.

   You will receive warning messages if you try to continue without filling in either of the two fields or if you set the location for the Web server configuration file incorrectly. As we were using the IBM HTTP Server, we selected the radio button for it and filled in the associated directory information.

2. **WebSphere Application Server is already installed**

When WebSphere Application Server was already installed on our AIX machine, we saw the following window:

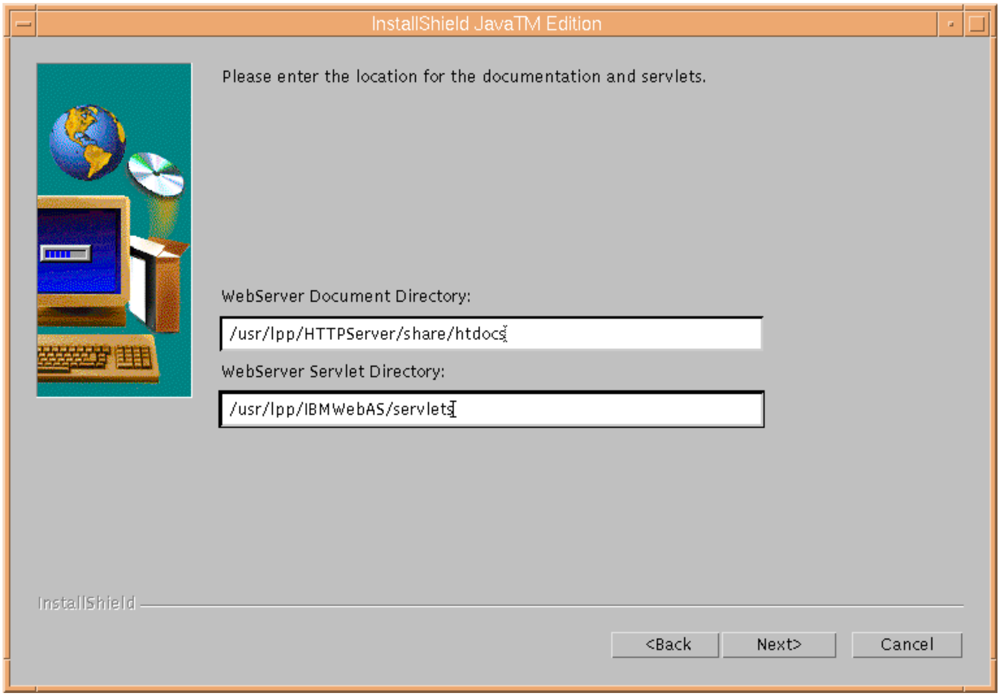

Please enter the location for the documentation and servlets.

WebServer Document Directory: /usr/lpp/HTTPServer/share/htdocs

WebServer Servlet Directory: /usr/lpp/IBMWebAS/servlets

Figure 218. Request to Supply Web Server and Servlet Configuration Details

In this case, you are again asked to supply two directory locations to the install process:

1. The first piece of information it asks you for is the same as in Point 1 on page 298: that is, the Web server document directory. The installation process requires this information as it places some files in this directory that are used when the Common Configuration utility is launched. We supplied the default document directory for IBM HTTP Server as that was the Web server we had preinstalled.
2. The second piece of information is different from Point 2 on page 298. In this case, it asks where the Web server servlet directory is. This directory is where the Common Configuration servlets will be placed during the installation process. We supplied the default servlet directory for WebSphere Application Server on AIX in the second field.

What we have just seen for AIX is very similar in the Solaris platform.

20.1.2.2 Providing Information on Windows NT

The information you provide during the installation depends on whether or not WebSphere Application Server is already installed on your system.

1. **WebSphere Application Server is not already installed**

   When WebSphere Application Server was not preinstalled on our Windows NT machine, we saw the following window:

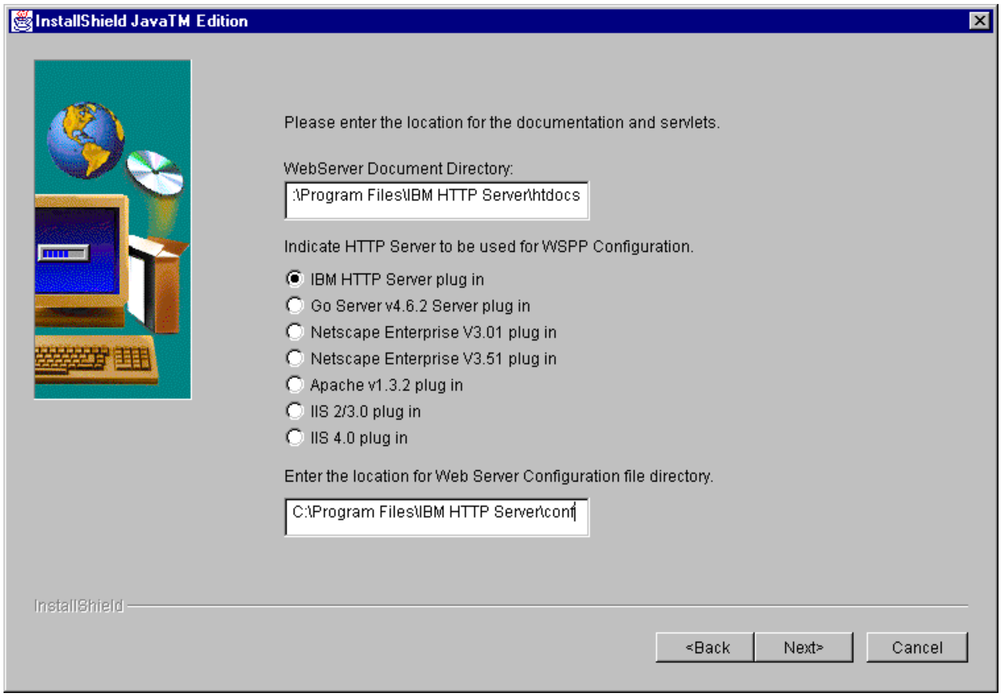

Please enter the location for the documentation and servlets.

WebServer Document Directory:

:\Program Files\IBM HTTP Server\htdocs

Indicate HTTP Server to be used for WSPP Configuration.

- ● IBM HTTP Server plug in
- ○ Go Server v4.6.2 Server plug in
- ○ Netscape Enterprise V3.01 plug in
- ○ Netscape Enterprise V3.51 plug in
- ○ Apache v1.3.2 plug in
- ○ IIS 2/3.0 plug in
- ○ IIS 4.0 plug in

Enter the location for Web Server Configuration file directory.

C:\Program Files\IBM HTTP Server\conf

Figure 219. Request to Supply Web Server Configuration Details

The above window will be displayed if you do not have WebSphere Application Server already installed. In order for the correct Web server plug-in to be installed, you will need to select one of the radio buttons to inform the installation process which Web server is installed on the machine.

In addition, you are asked to provide two directories:

1. The first piece of information the installation process asks you for is the Web server document directory. The installation process requires this information as it places some files in this directory that are used when the Common Configuration utility is launched.
2. The second piece of information the installation process asks you for is the Web server configuration file directory. Because the installation in this case is going to install WebSphere Application Server, it must make modifications to the Web server configuration file so that the Web server will work together with WebSphere Application Server.

You will receive warning messages if you try to continue without filling in either of the two fields or if you set the location for the Web server configuration file incorrectly. As we were using the IBM HTTP Server, we selected the radio button for it and filled in the associated directory information.

2. WebSphere Application Server is already installed

When WebSphere Application Server was already installed on our Windows NT machine, we saw the following window:

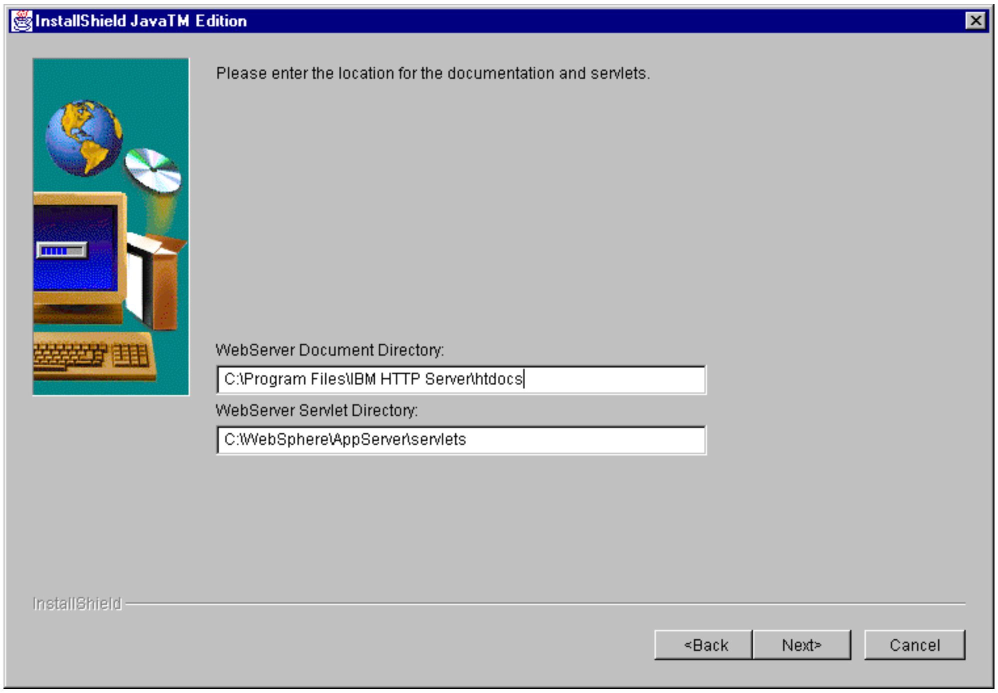

InstallShield JavaTM Edition

Please enter the location for the documentation and servlets.

WebServer Document Directory:

C:\Program Files\IBM HTTP Server\htdocs

WebServer Servlet Directory:

C:\WebSphere\AppServer\servlets

Figure 220. Request to Supply Web Server and Servlet Configuration Details

In this case, you are again asked to supply two directory locations to the install process:

1. The first piece of information the installation process asks you for is the same as in Point 1 on page 299: the Web server document directory. The installation process requires this information as it places in this directory some files that are used when the Common Configuration utility is launched. We supplied the default document directory for IBM HTTP Server as that was the Web server we had preinstalled.
2. The second piece of information is different from Point 2 on page 300. Now, in fact, the installation process asks you where the Web server servlet directory is. This directory is where the Common Configuration servlets will be placed during the installation process. We supplied the default servlet directory for WebSphere Application Server on Windows NT in the second field.

20.1.3 Remaining Installation Steps

After clicking **Next** we saw a window asking if we wanted the installation program to replace other programs already installed. We selected **No** in this case, since we had not installed any programs on our system yet:

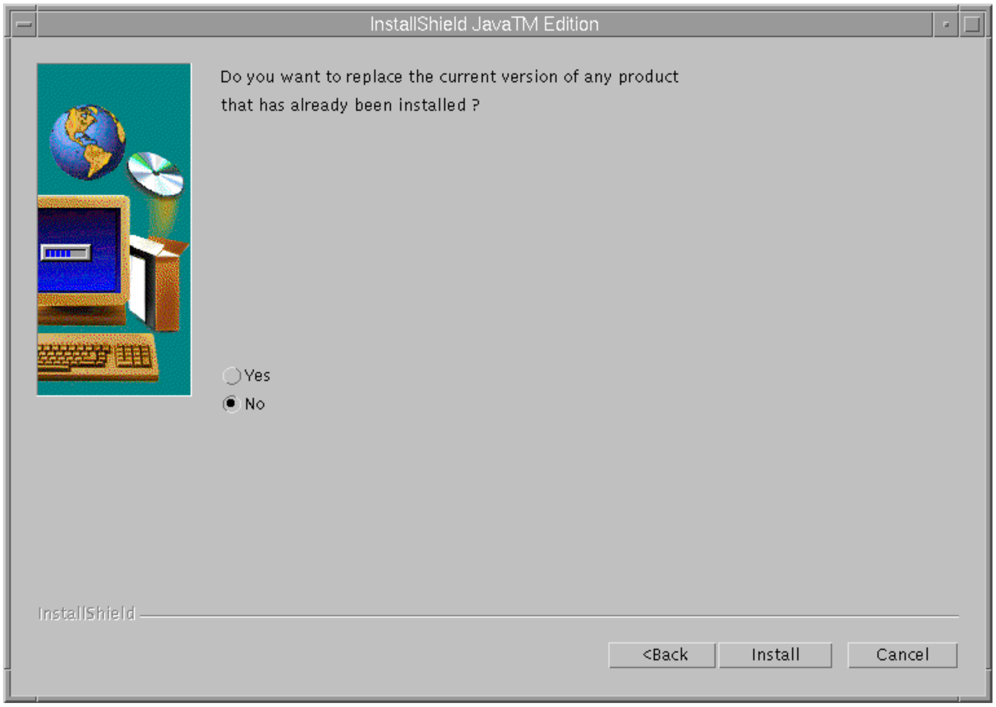

Do you want to replace the current version of any product that has already been installed ?

○ Yes

● No

Figure 221. Choose to Replace Version

After clicking **Install**, the Common Configuration servlets were installed. When WebSphere Application Server was not already installed on our system, we first saw the following screen displayed, showing the installation progress for WebSphere Application Server:

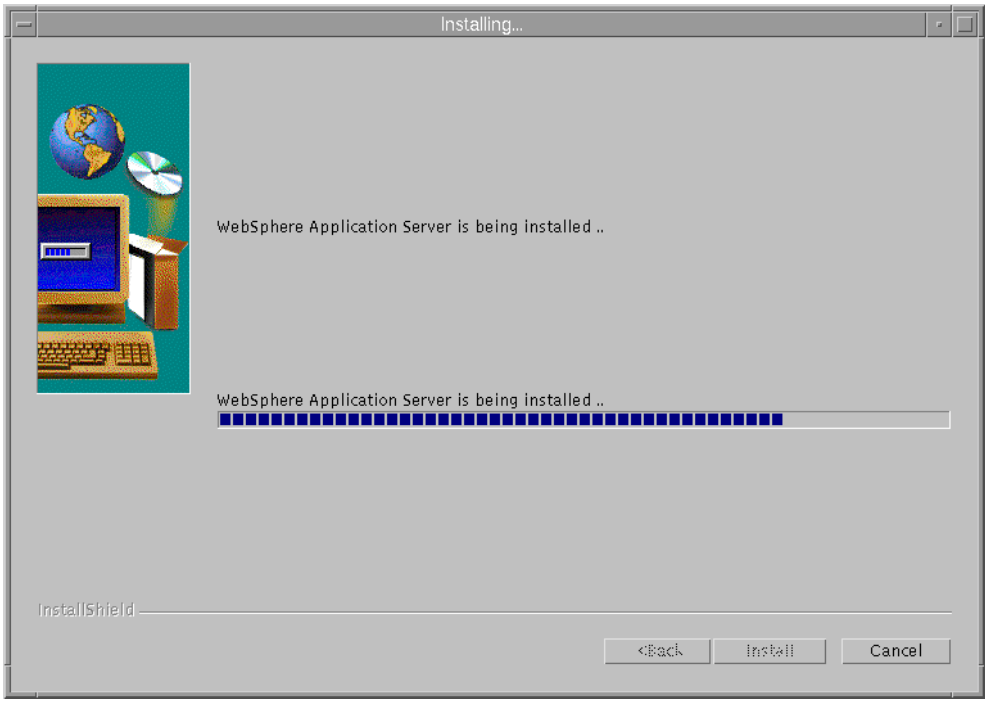

WebSphere Application Server is being installed ..

WebSphere Application Server is being installed ..

Figure 222. WebSphere Application Server Installation Progress Indicator

The following screen is always displayed during the installation of the Common Configuration utility, showing the installation progress for the Common Configuration utility itself:

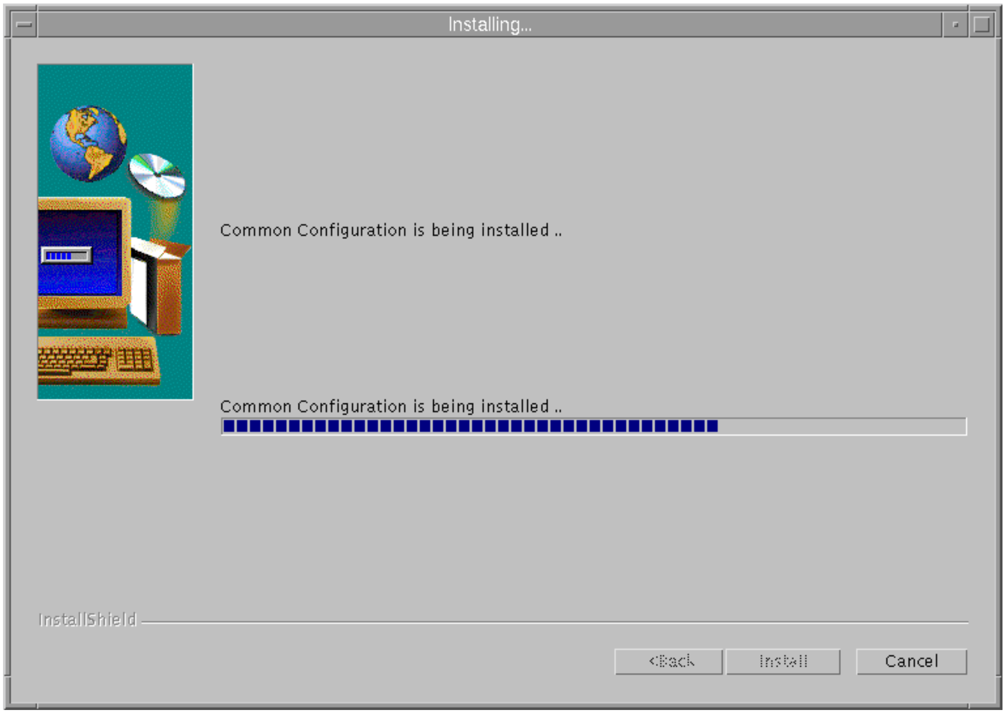

Common Configuration is being installed ..

Common Configuration is being installed ..

Figure 223. Common Configuration Installation Progress Indicator

Then, you will be informed that the installation completed successfully:

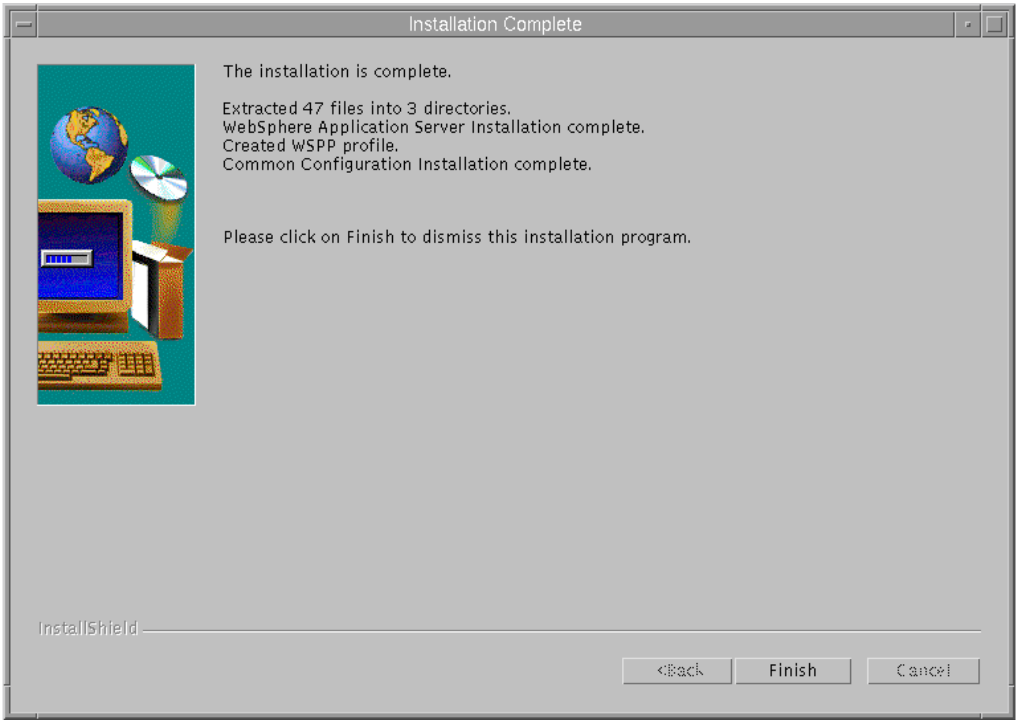

The installation is complete.

Extracted 47 files into 3 directories.
WebSphere Application Server Installation complete.
Created WSPP profile.
Common Configuration Installation complete.

Please click on Finish to dismiss this installation program.

Figure 224. Successful Completion Indication

We clicked on **Finish** to dismiss the installation process.

At this point, we could verify that the installation of the Common Configuration utility transparently installs WebSphere Application Server on the Common Configuration server. For example, on our AIX system, we saw that the filesets shown in the following figure were the only ones that were installed on our system as a result of the Common Configuration installation:

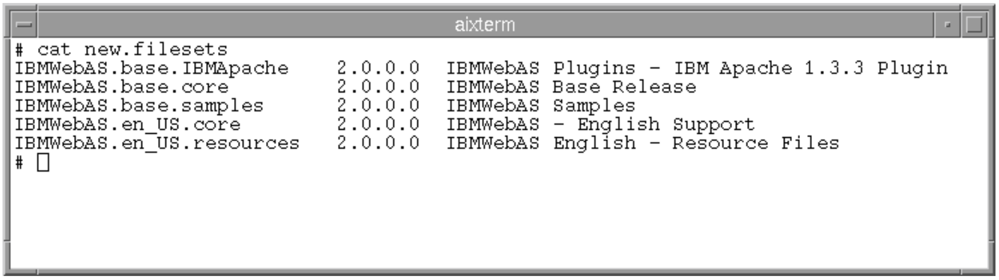

```
# cat new.filesets
IBMWebAS.base.IBMApache    2.0.0.0  IBMWebAS Plugins - IBM Apache 1.3.3 Plugin
IBMWebAS.base.core         2.0.0.0  IBMWebAS Base Release
IBMWebAS.base.samples      2.0.0.0  IBMWebAS Samples
IBMWebAS.en_US.core        2.0.0.0  IBMWebAS - English Support
IBMWebAS.en_US.resources   2.0.0.0  IBMWebAS English - Resource Files
# 
```

Figure 225. Installing the Common Configuration Utility Installs WebSphere Application Server

20.2 Preparing for the Common Configuration Utility

Communication between the Common Configuration server machine and the WebSphere Performance Pack Version 2 machine that the Common Configuration server is configuring is done via Remote Method Invocation (RMI).

A machine where WebSphere Performance Pack has been installed, and that needs to be configured through the Common Configuration utility, is also known as *WebSphere Performance Pack server machine*. On a WebSphere Performance Pack server machine, the installation program creates a file that must be executed before you can use the Common Configuration to configure WebSphere Performance Pack on this machine. The file is called wsppserver on AIX and Solaris, and wsppserver.cmd on Windows NT. It is located in the rmi subdirectory of the installation directory that you instructed the installation program to use at installation time.

Notice that the wsppserver file must not be executed on the Common Configuration server machine, but on the WebSphere Performance Pack server machine that is going to be the target of the Common Configuration utility. In other words, the wsppserver script must be executed on the machine where one of the three components (AFS, ND, and WTE) of WebSphere Performance Pack has been installed, and needs to be configured (typically remotely) through the Common Configuration utility.

On AIX, we followed these steps to execute the wsppserver file:

```
cd /usr/lpp/WSPP/rmi
chmod +x wsppserver
./wsppserver
```

We confirmed the success of the wsppserver execution by looking at the wsppcc.log file that is also present in the same directory as wsppserver:

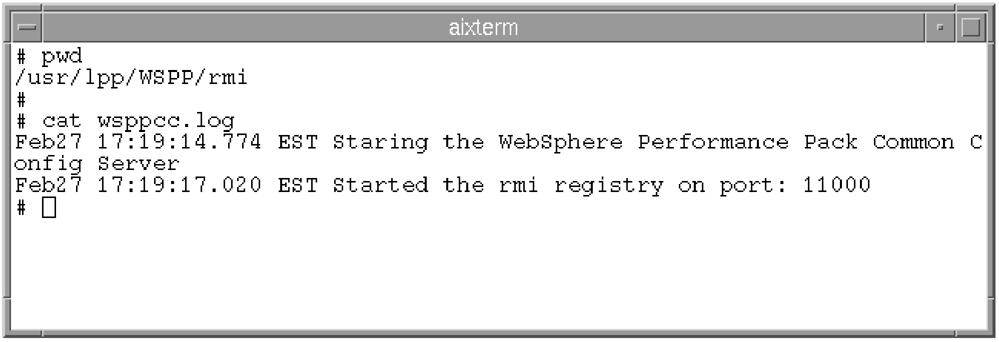

```
# pwd
/usr/lpp/WSPP/rmi
#
# cat wsppcc.log
Feb27 17:19:14.774 EST Staring the WebSphere Performance Pack Common C
onfig Server
Feb27 17:19:17.020 EST Started the rmi registry on port: 11000
# 
```

Figure 226. Confirming the Success of the wsppserver Execution

If you would like wsppserver to be run each time the WebSphere Performance Pack machine is rebooted, you could add it to an rc file on AIX or Solaris; on a Windows NT machine, add wsppserver to the Startup folder or registry.

To summarize, the steps that must be done before using the Common Configuration utility are:

1. Identify a Websphere Performance Pack server machine to be the target of the Common Configuration utility and on this machine execute the wsppserver script file. This file carries a .cmd extension on Windows NT.
2. Identify a configuration server machine and on this machine:

- Install the Common Configuration utility.
- Start the WebSphere Application Server-compliant Web server if it is not already running.

The next step that you must perform is dependent upon which WebSphere Performance Pack component you plan to configure with the Common Configuration utility.

20.3 Launching the Common Configuration Utility

To launch the Common Configuration utility from a Web browser, go to `http://common.configuration.server.name/wspp/startServer.html`, where common.configuration.server.name is the host name of the Common Configuration server machine. Following is the Common Configuration home page:

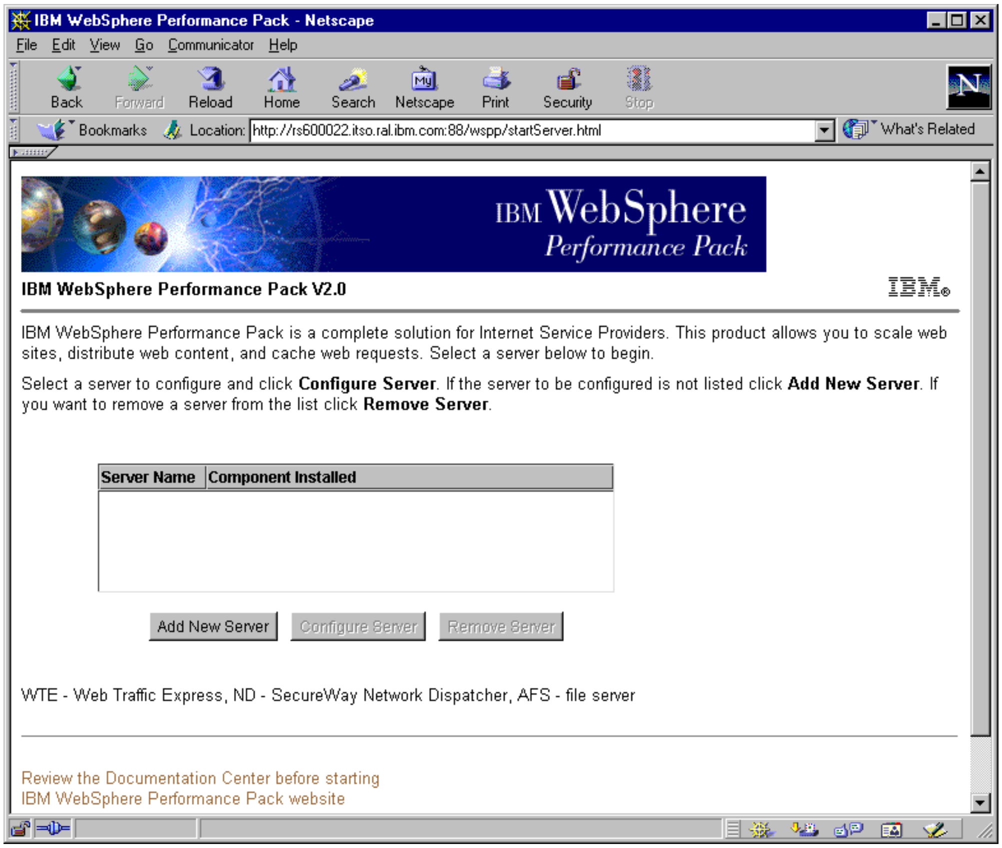

IBM WebSphere Performance Pack - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com:88/wspp/startServer.html
What's Related
IBM WebSphere Performance Pack

IBM WebSphere Performance Pack V2.0

IBM WebSphere Performance Pack is a complete solution for Internet Service Providers. This product allows you to scale web sites, distribute web content, and cache web requests. Select a server below to begin.

Select a server to configure and click **Configure Server**. If the server to be configured is not listed click **Add New Server**. If you want to remove a server from the list click **Remove Server**.

Server Name	Component Installed

WTE - Web Traffic Express, ND - SecureWay Network Dispatcher, AFS - file server

Review the Documentation Center before starting
IBM WebSphere Performance Pack website

Figure 227. Common Configuration Main Page

From the server selection list on the page, you can either click **Add New Server** or select a server from the list (if the list is not empty) and then click **Configure Server** or **Remove Server**.

The first time we accessed the Common Configuration utility, we selected **Add New Server** and in the subsequent Server Addition Request flield, entered the name of our WebSphere Performance Pack server machine:

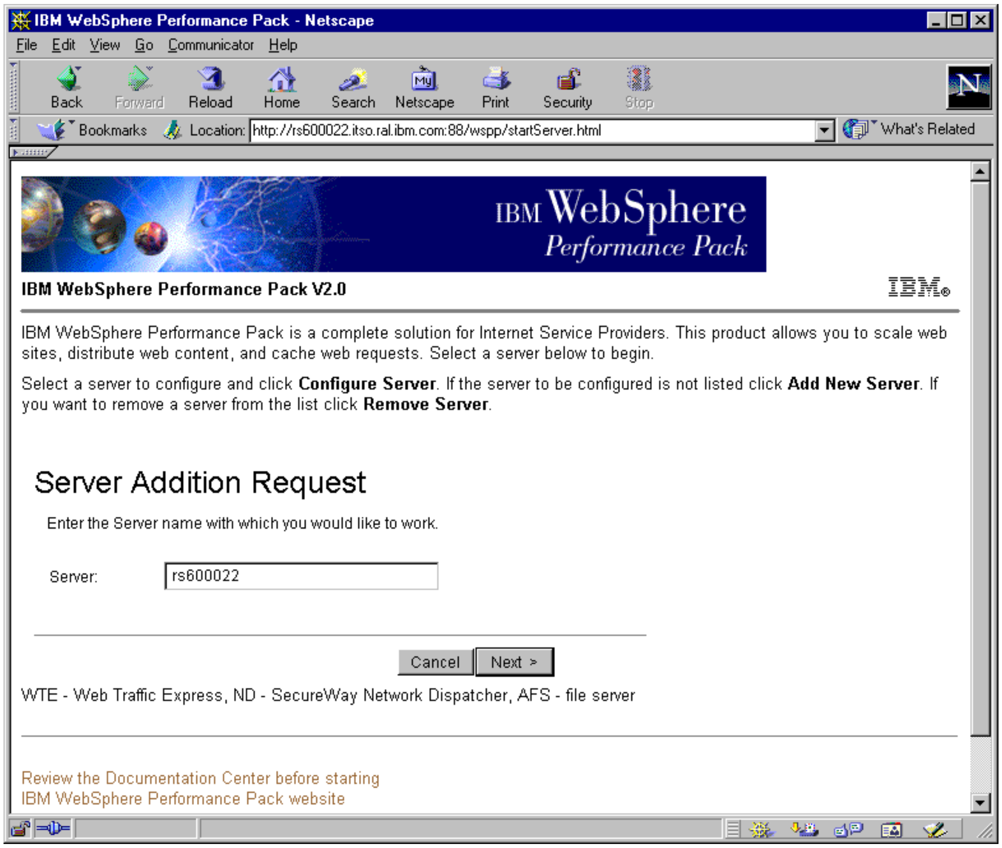
IBM WebSphere Performance Pack - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com:88/wspp/startServer.html
What's Related
IBM WebSphere Performance Pack

IBM WebSphere Performance Pack V2.0

IBM WebSphere Performance Pack is a complete solution for Internet Service Providers. This product allows you to scale web sites, distribute web content, and cache web requests. Select a server below to begin.

Select a server to configure and click **Configure Server**. If the server to be configured is not listed click **Add New Server**. If you want to remove a server from the list click **Remove Server**.

Server Addition Request

Enter the Server name with which you would like to work.

Server: rs600022

WTE - Web Traffic Express, ND - SecureWay Network Dispatcher, AFS - file server

Review the Documentation Center before starting
IBM WebSphere Performance Pack website

Figure 228. Server Addition Request

As you can see from the figure above, in this example, the Common Configuration server machine (see the host name in the URL) is the same as the WebSphere Performance Pack server machine (see the host name in the Server Addition Request field). However, there is no need for this to happen. The WebSphere Performance Pack server machines that a Common Configuration server can administer can be either remote or local.

After clicking the **Next** button, we saw a refreshed version of the Common Configuration main page, showing the name of the WebSphere Performance Pack server machine. The Common Configuration utility can immediately detect the WebSphere Performance Pack components that are installed on a WebSphere Performance Pack server and that, therefore, are subject to be configured through the Common Configuration utility. The following figure shows how the Common Configuration utility immediately detected the presence of WTE in the WebSphere Performance Pack server machine whose host name has been added in the Server Addition Request field:

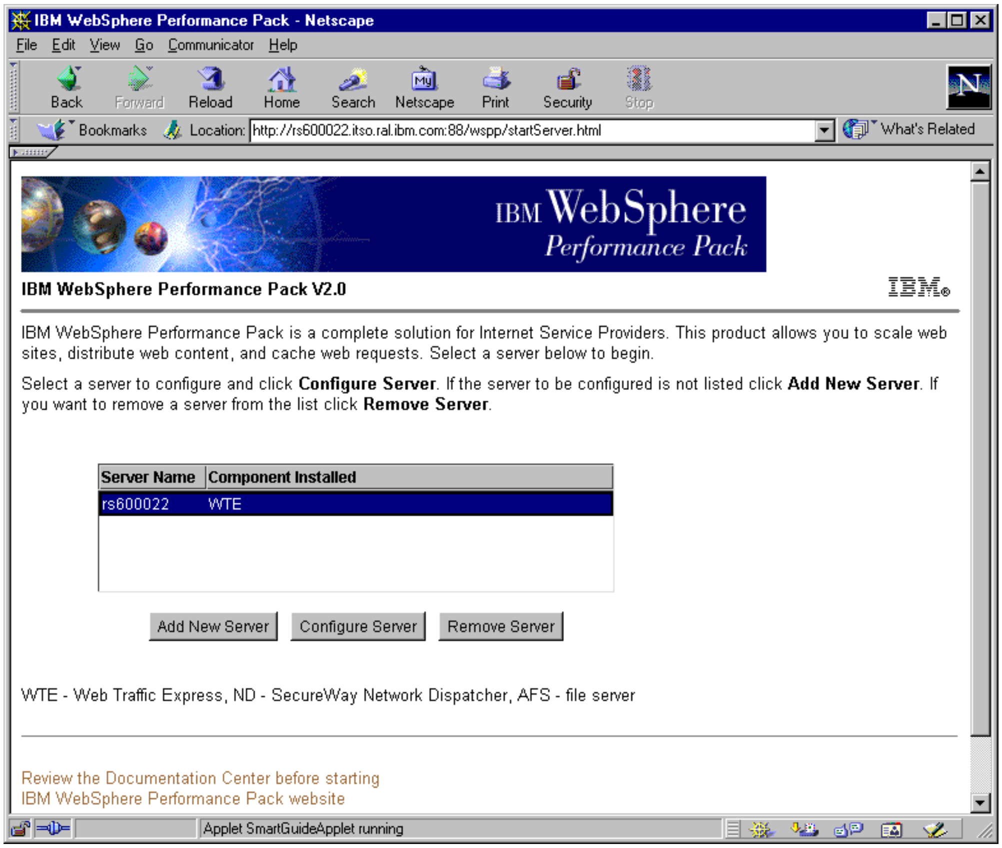

IBM WebSphere Performance Pack - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com:88/wspp/startServer.html
What's Related
IBM WebSphere Performance Pack
IBM

IBM WebSphere Performance Pack V2.0

IBM WebSphere Performance Pack is a complete solution for Internet Service Providers. This product allows you to scale web sites, distribute web content, and cache web requests. Select a server below to begin.

Select a server to configure and click **Configure Server**. If the server to be configured is not listed click **Add New Server**. If you want to remove a server from the list click **Remove Server**.

Server Name	Component Installed
rs600022	WTE

Add New Server | Configure Server | Remove Server

WTE - Web Traffic Express, ND - SecureWay Network Dispatcher, AFS - file server

Review the Documentation Center before starting
IBM WebSphere Performance Pack website

Applet SmartGuideApplet running

Figure 229. One Performance Pack Server Added to the Common Configuration Page

Multiple WebSphere Performance Pack server machines (the local machine as well as remote machines) can be added to the list. The Common Configuration server can automatically detect the WebSphere Performance Pack components that are installed on all the WebSphere Performance Pack server machines, so that these can be configured through the Common Configuration utility. The following figure shows an example of this:

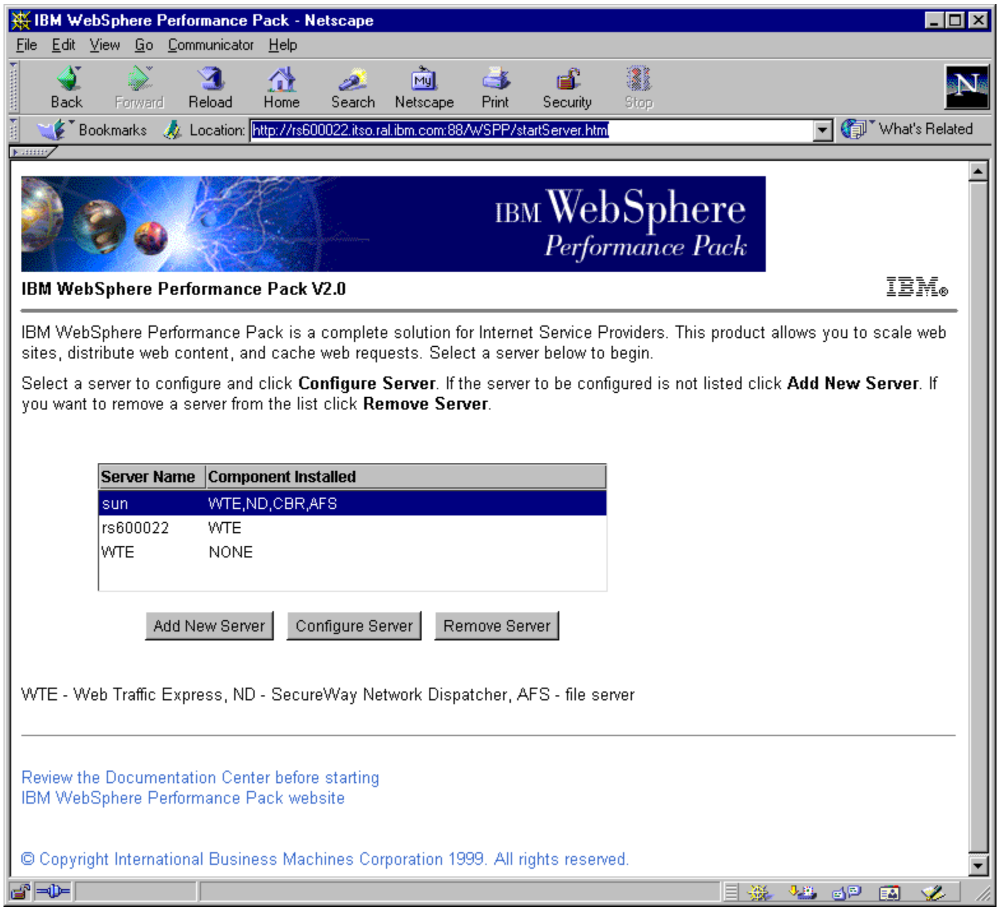
IBM WebSphere Performance Pack - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com:88/WSPP/startServer.html
What's Related
IBM WebSphere Performance Pack
IBM

IBM WebSphere Performance Pack V2.0

IBM WebSphere Performance Pack is a complete solution for Internet Service Providers. This product allows you to scale web sites, distribute web content, and cache web requests. Select a server below to begin.

Select a server to configure and click **Configure Server**. If the server to be configured is not listed click **Add New Server**. If you want to remove a server from the list click **Remove Server**.

Server Name	Component Installed
sun	WTE,ND,CBR,AFS
rs600022	WTE
WTE	NONE

WTE - Web Traffic Express, ND - SecureWay Network Dispatcher, AFS - file server

Review the Documentation Center before starting
IBM WebSphere Performance Pack website

© Copyright International Business Machines Corporation 1999. All rights reserved.

Figure 230. Common Configuration with Multiple WebSphere Performance Pack Servers

With a server selected, click **Configure Server** to begin the configuration. Once activated, the Common Configuration does not differ too much from the graphical user interface (GUI) tools that can be used to configure AFS, ND and WTE respectively. The advantage of the Common Configuration utility is that it allows centralized, remote, and secure administration of entire WebSphere Performance Pack sites. Only authenticated users can administer a WebSphere Performance Pack server through a Common Configuration server.

Appendix A. Special Notices

This publication is intended to help you to plan for, install, configure, use, tune and troubleshoot IBM Web Traffic Express, the Caching and Filtering component of IBM WebSphere Performance Pack. The information in this publication is not intended as the specification of any programming interfaces that are provided by IBM WebSphere Performance Pack. See the PUBLICATIONS section of the IBM Programming Announcement for IBM WebSphere Performance Pack for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX
AS/400
eNetwork
NetView
OS/2
SecureWay
SP2
SystemView
VisualAge
APPN
Deep Blue
IBM
Operating System/2
RS/6000
SP
System/390
TXSeries
WebSphere

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

SET and the SET logo are trademarks owned by SET Secure Electronic Transaction LLC.

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

Other company, product, and service names may be trademarks or service marks of others.

Appendix B. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

B.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How to Get ITSO Redbooks” on page 317.

- *IBM WebSphere Performance Pack: Web Content Management with IBM AFS Enterprise File System*, SG24-5857
- *IBM WebSphere Performance Pack: Load Balancing with IBM SecureWay Network Dispatcher*, SG24-5858
- *IBM WebSphere Performance Pack Usage and Administration*, SG24-5233
- *Internet Security in the Network Computing Framework*, SG24-5220
- *Network Computing Framework Component Guide*, SG24-2119
- *Load Balancing Internet Servers*, SG24-4993
- *RS/6000 Performance Tools in Focus*, SG24-4989
- *Understanding IBM RS/6000 Performance and Sizing*, SG24-4810
- *TCP/IP Tutorial and Technical Overview*, GG24-3376
- *Java 2 Network Security*, SG24-2109
- *RS/6000 Web Server - Powered by Apache,* SG24-5132

B.2 Redbooks on CD-ROMs

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at `http://www.redbooks.ibm.com/` for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	**Collection Kit Number**
System/390 Redbooks Collection	SK2T-2177
Networking and Systems Management Redbooks Collection	SK2T-6022
Transaction Processing and Data Management Redbooks Collection	SK2T-8038
Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
AS/400 Redbooks Collection	SK2T-2849
Netfinity Hardware and Software Redbooks Collection	SK2T-8046
RS/6000 Redbooks Collection (BkMgr Format)	SK2T-8040
RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

B.3 Other Publications

The following publication is also relevant as a further information source:

- *SecureWay Network Dispatcher User’s Guide Version 2.1 for Solaris, Windows NT and AIX*, GC31-8496

B.4 Referenced Web Sites

- http://www.redbooks.ibm.com
- http://www.ibm.com
- http://rs6000.itso.ibm.com
- http://sphinx.sg.ibm.com
- http://w3.ibm.com
- http://wtr05178.itso.ral.ibm.com/PICSxmp.age1.html
- http://wtr05178.itso.ral.ibm.com/PICSxmp/age2.html
- http://wtr05178.itso.ral.ibm.com/PICSxmp/age3.html
- http://wtr05178.itso.ral.ibm.com/PICSxmp/age4.html
- http://w3.itso.ibm.com
- http://www.javasoft.com
- http://www.software.ibm.com/webservers/appserv
- http://eggplant.rte.microsoft.com/wpad/wpad.txt
- http://www.w3.org/PICS
- http://www.ibm.com/java/jdk/download/index.html
- http://www.microsoft.com/downloads
- http://www.netscape.com/computing/download/index.html
- http://developer.netscape.com/software/jdk/download.html
- http://service.boulder.ibm.com/asd-bin/doc/en_us/catalog.htm
- http://www.sun.com/solaris/jdk/download.1.1.7
- http://www.javasoft.com/products/jdk/1.1/download-jdk-windows.html
- http://www.support.tivoli.com/sva/shasdk.html
- http://www.lotus.com
- http://www.software.ibm.com/webservers/wte
- http://rs600022.itso.ral.ibm.com
- http://wtr05178.ral.itso.ibm.com
- http://www.ics.raleigh.ibm.com/WebTrafficExpress/main.html
- http://sphinx.sg.ibm.com/sphinx.html
- http://sphinx.sg.ibm.com/sphinx11.html
- http://rs6000.itso.ral.ibm.com/pacfiles/autoproxy
- http://ADMIN.uconn.edu
- http://sun.itso.ral.ibm.com
- http://aixafs.itso.ral.ibm.com
- http://rs600030/afs/webstone/html/sphinx.html
- http://home.netscape.com
- http://www.w3.org/Amaya/User/Put.html

- http://www.apacheweek.com/features/put
- http://rs600022.itso.ral.ibm.com/admin-bin/webexec/frameset.html
- http://www.w3.org/TR/REC-PICS-labels
- http://www.w3.org/pub/WWW/PICS/selfrat.htm
- http://wtr05178.itso.ral.ibm.com/PICSxmp:
- http://www.w3.org/pub/WWW/PICS/services.html
- http://rs600030.itso.ral.ibm.com/wteRatings/newEnzo.html
- http://rs600030.itso.ral.ibm.com/wteoRatings/V1-0.html
- http://rs600030.itso.ral.ibm.com/wteRatings/V1-0.html
- http://WTR05178.itso.ral.ibm.com/PICSxmp/age4.html
- http://wtr05178.itso.ral.ibm.com/PICSxmp/age1.html
- http://www.coolratings.com/CoolSite
- http://wtr05178.itso.ral.ibm.com/PICSxmp/testage4.html
- http://9.67.133.18/purchase.html
- http://www.redbooks.ibm.com/redpieces.html
- http://www.elink.ibmlink.ibm.com/pbl/pbl

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at `http://www.redbooks.ibm.com/`.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Redbooks Web Site on the World Wide Web**

  `http://w3.itso.ibm.com/`

- **PUBORDER** – to order hardcopies in the United States
- **Tools Disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLCAT REDPRINT
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get BookManager BOOKs of redbooks, type the following command:

  ```
  TOOLCAT REDBOOKS
  ```

  To get lists of redbooks, type the following command:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  ```

  To register for information on workshops, residencies, and redbooks, type the following command:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1998
  ```

- **REDBOOKS Category on INEWS**
- **Online** – send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (`http://www.redbooks.ibm.com/redpieces.html`). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** – send orders to:

	IBMMAIL	**Internet**
In United States	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** – send orders to:

IBM Publications
Publications Customer Support
P.O. Box 29570
Raleigh, NC 27626-0570
USA

IBM Publications
144-4th Avenue, S.W.
Calgary, Alberta T2P 3N5
Canada

IBM Direct Services
Sortemosevej 21
DK-3450 Allerød
Denmark

- **Fax** – send orders to:

United States (toll free)	1-800-445-9269
Canada	1-800-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States)** or **(+1) 408 256 5422 (Outside USA)** – ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **On the World Wide Web**

Redbooks Web Site	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (`http://www.redbooks.ibm.com/redpieces.html`). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	**Order Number**	**Quantity**

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

□ Invoice to customer number

□ Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

List of Abbreviations

ACL	access control list
AFS	Enterprise File System
AIX	advanced interactive executive
ATM	Asynchronous Transfer Mode
API	application programming interface
APPN	Advanced Peer-to-Peer Network
ARP	Address Protocol Request
CARP	Cache Array Routing Protocol
CB	Component Broker
CBR	Content Based Routing
CORBA	Common Object Request Broker Architecture
DMZ	demilitarized zone
DNS	Domain Name System
EJB	Enterprise JavaBeans
EJS	Enterprise Java Services
FDDI	Fiber Distributed Data Interface
FTP	File Transfer Protocol
GC	garbage collector
GEM	Global Enterprise Manager
GIF	graphic interchange format
GUI	graphical user interface
HACMP	High Availability Cluster Multiprocessing
HP-UX	Hewlett-Packard UNIX
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
IBM	International Business Machines Corporation
IP	Internet Protocol
ISP	Internet service provider
ISS	Interactive Session Support
ITSO	International Technical Support Organization
JDK	Java Development Kit
JRE	Java Runtime Environment
JSP	JavaServer Pages
JVM	Java Virtual Machine
LAN	local area network
LDAP	Lightweight Directory Access Protocol
MAC	Media Access Control
Mbps	megabits per second
MBps	megabytes per second
MIB	management information base
MVS	multiple virtual storage
NAP	network access points
NCSA	National Center for Supercomputing Applications
ND	Network Dispatcher
NNTP	NetNews transfer protocol
OS/2	Operating System/2
PICS	Platform for Internet Content Selection
POP	points of presence
POP3	Post Office Protocol 3
RAM	random access memory
RCA	Remote Cache Access
RFC	Request for Comment
RMI	Remote Method Invocation
RPC	remote procedure call
RSACi	Recreational Software Advisory Council on the Internet
SDA	Server Directed Affinity
SSL	Secure Sockets Layer
SMTP	Simple Mail Transfer Protocol
SOCKS	software common knowledge IR system
SRC	System Resource Controller
TCP/IP	Transmission Control Protocol/Internet Protocol
TEC	Tivoli Enterprise Console
UDP	User Datagram Protocol
URL	Universal Resource Locator, Uniform Resource Locator
WAN	wide area network
WAND	Wide Area Network Dispatcher
WTE	Web Traffic Express

WLM	workload manager
WPAD	Web Proxy Auto Discovery
WWW	World Wide Web
W3C	World Wide Web Consortium
XML	Extensible Markup Language

Index

ITSO Redbook Evaluation

IBM WebSphere Performance Pack: Caching and Filtering with IBM Web Traffic Express
SG24-5859-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at http://www.redbooks.ibm.com
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?
_ **Customer** _ **Business Partner** _ **Solution Developer** _ **IBM employee**
_ **None of the above**

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction __________

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

SG24-5859-00

Printed in the U.S.A.

IBM WebSphere Performance Pack: Caching and Filtering with IBM Web Traffic Express

SG24-5859-00