



US005870550A

United States Patent [19]

[11] Patent Number: **5,870,550**

Wesinger, Jr. et al.

[45] Date of Patent: **Feb. 9, 1999**

[54] **WEB SERVER EMPLOYING MULTI-HOMED, MODULAR FRAMEWORK**

Primary Examiner—Lance Leonard Barry, Esq.
Attorney, Agent, or Firm—McDonnell Boehnen Hulbert & Berghoff

[75] Inventors: **Ralph E. Wesinger, Jr.**, San Jose;
Christopher D. Coley, Morgan Hill,
both of Calif.

[57] ABSTRACT

[73] Assignee: **Network Engineering Software**, San Jose, Calif.

A Web server is provided having a multi-homed, modular framework. The modular framework allows extensions to the Web server to be easily compiled into the Web server, allowing the extensions to run natively as part of the server instead of incurring the overhead typical of CGI scripts, for example. The multi-homing capabilities of the Web server provide the appearance to Web users of multiple distinct and independent servers, allowing a small company or individual to create the same kind of Web presence enjoyed by larger companies. In effect, multiple virtual servers run on the same physical machine. The Web server as a whole is easily extensible to allow additional capabilities to be provided natively within the Web server itself. Furthermore, each virtual server is independently configurable in order to turn different capabilities on or off or to modify operation of the virtual server. The Web server is also provided with enhanced security features, built-in animation capability, and other features that afford maximum flexibility and versatility.

[21] Appl. No.: **607,068**

[22] Filed: **Feb. 26, 1996**

[51] Int. Cl.⁶ **G06F 17/30**

[52] U.S. Cl. **395/200.48**; 395/200.49;
707/10

[58] Field of Search 707/9-10, 501,
707/513, 200, 47-48; 395/200.49, 200.57,
400

[56] References Cited

U.S. PATENT DOCUMENTS

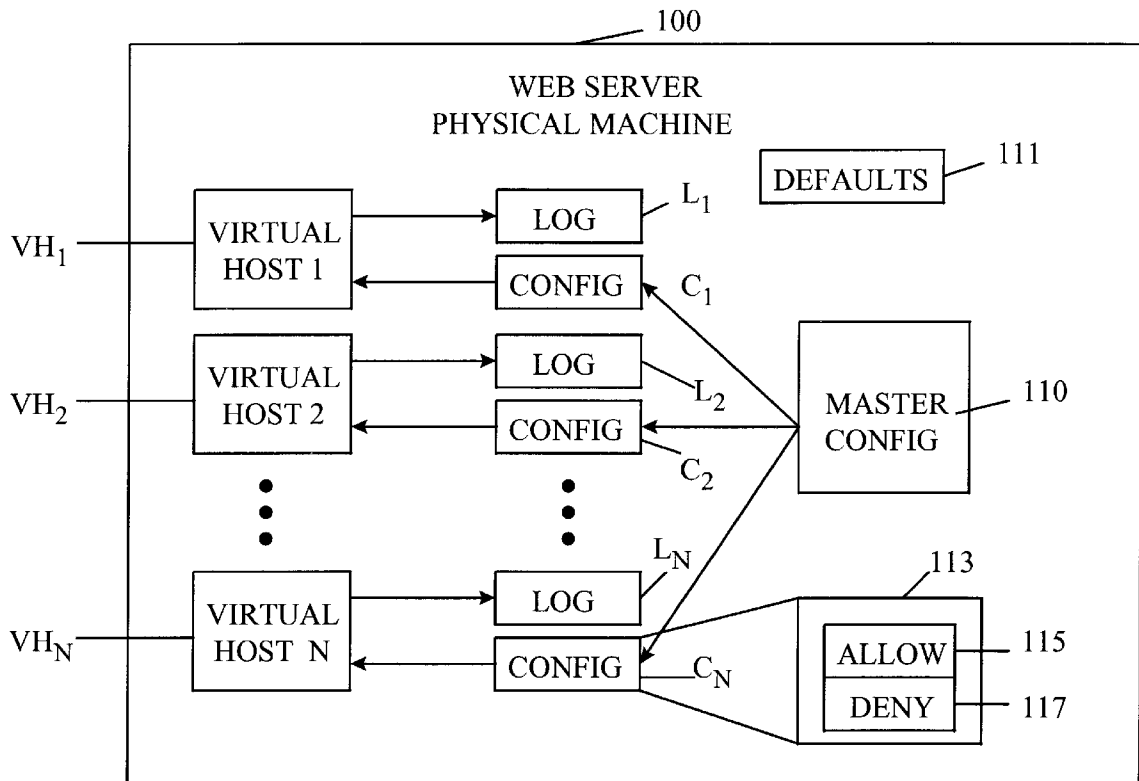
5,617,540 4/1997 Civanlar et al. 395/200.57

OTHER PUBLICATIONS

G. Stephen Nace, Search Tools on the Internet, <http://www.nashville.com/~gsnace/search.htm>, 1993.

William Cross, All-in-One Search Page, <http://www.albany.net/allinone/all1www.html>, 1995.

15 Claims, 6 Drawing Sheets



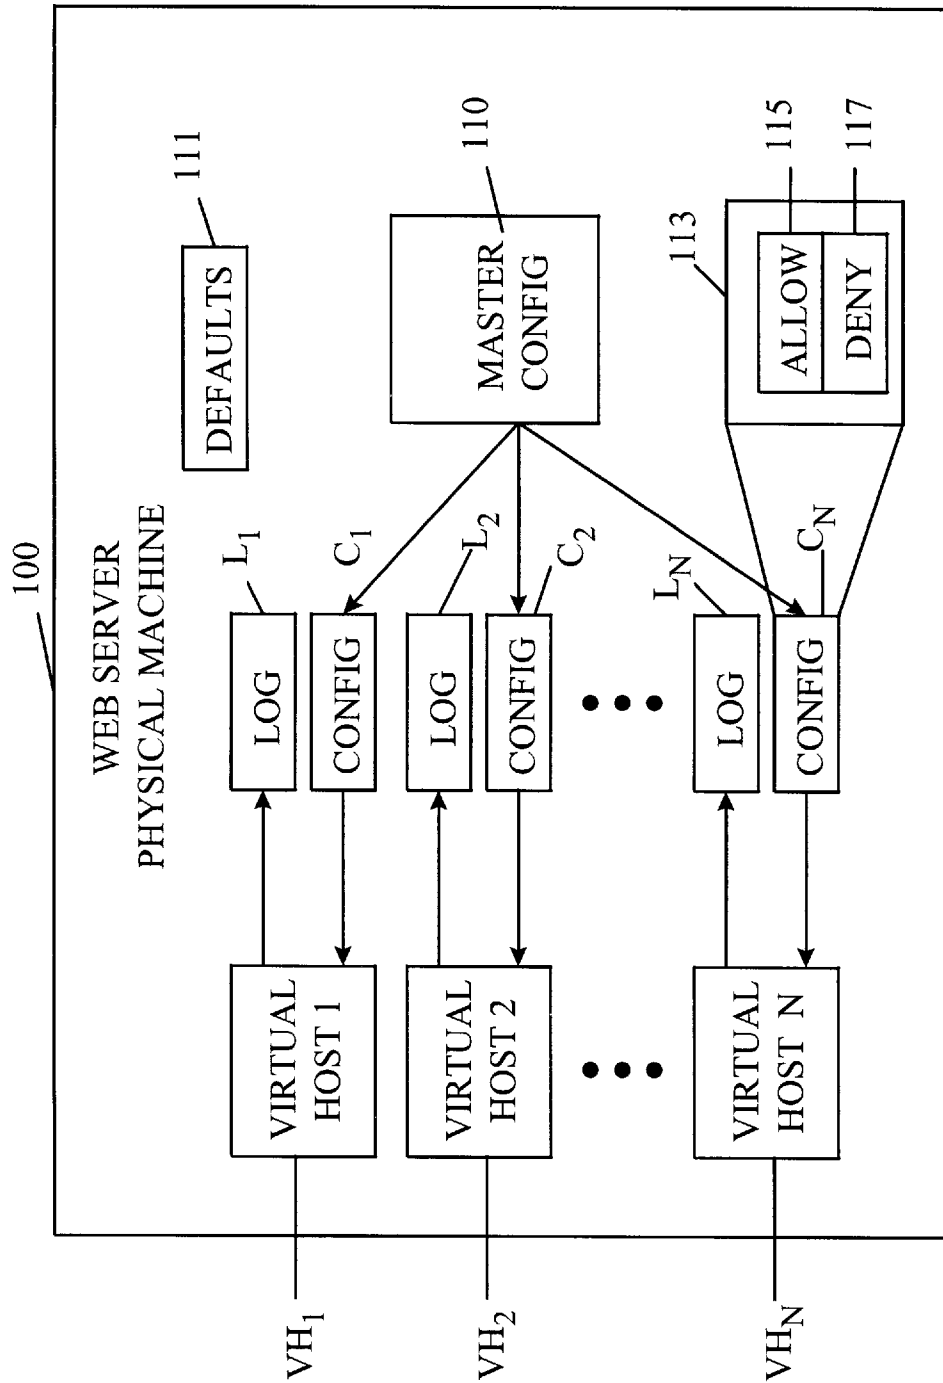


Fig. 1

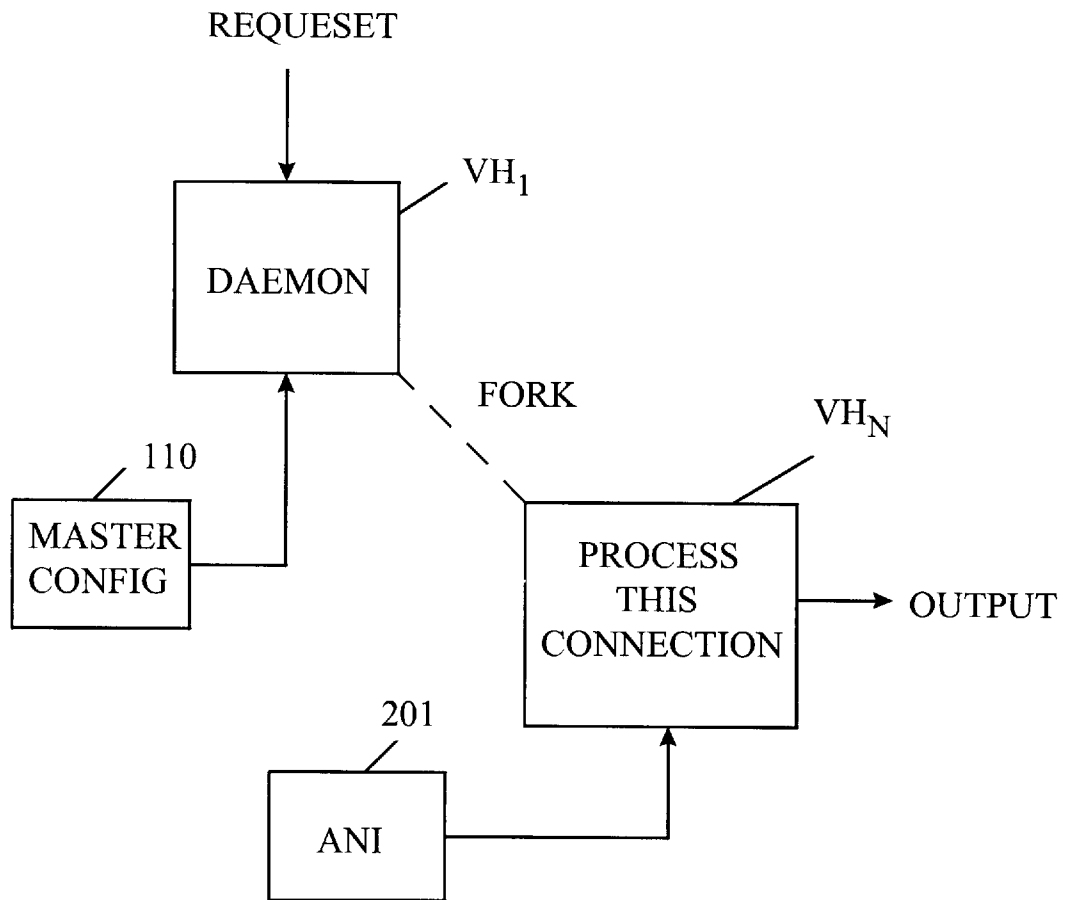


Fig. 2

```
struct PROCESS_TYPES DefaultProcessTypes[ ] = {
/* Animation scripts (Server push) */
“ani”,    TY_ANIMATION,  NULL,          ProcessANI,
“cgi”,    TY_CGI,        NULL,          ProcessCGI,
/* html's */
“html”,   TY_HTML,      “text/html”,   SendFile,
/* shtml's will be server passed */
“shtml”,  TY_SHTML,     “image/html”,  ProcShtml,
/* gif graphics Images */
“gif”,    TY_GIF,       “image/gif”,   SendFile,
/* jpeg graphics Images*/
“jpeg”,   TY_JPEG,     “image/jpeg”,  SendFile,
“jpg”,    TY_JPEG,     “image/jpeg”,  SendFile,
“jpe”,    TY_JPEG,     “image/jpeg”,  SendFile,
/* map Image Maps */
“map”,    TY_MAP,      NULL,          ProcessMap,
/* default is to send file */
“”,       TY_DEFAULT,  “text/html”,   SendFile,
/* Structure terminator */
NULL,     TY_ALLDONE,  NULL,          NULL};
```

Fig. 3

```

PORT          = 80
RULE1        = {
    TIME      = "1AM-12PM"
}
WWW.SRMC.COM = {
    .CGI     = "PROCESSCGI"
    ROOT    = "/HOME/HONOLULU/HTML"
}
WWW.HONOLULU.NET = {
    .CGI     = ""
    ROOT    = "/HOME/HONOLULU/HTML"
}
WWW.SANJOSE.NET = {
    .CGI     = "PROCESSCGI"
    ALLOW   = {
        *.SRMC.COM
        205.138.192.*
        205.138.192.0/23
    }
    DENY    = {
        MISTERPAIN.COM
    }
}
WWPROXY.SRMC.COM = {
    MODE    = RT_SERVERPROXY
}
NS.SRMC.COM   = {
    ALLOW   = {
        192.168.0.*
        192.168.1.* = RULE1
        192.168.2.* = {
            TIME      = "1AM-12PM"
        }
        192.168.3.* = 192.168.2.*
    }
}
MJU.SRMC.COM = {
    ALLOW   = {
        192.168.0.0.23 = RULE1
    }
    DENY    = {
        192.168.0.* = {
            TIME      = "12PM-1AM"
        }
    }
}
}

```

Fig. 4

501 / (REMOTE IP OR NAME) USERNAME [DATE & TIME] "GET ...HTTP/1.0" ERC BYTESXFRD 511
503 / 505 / 507 / 509 /

Fig. 5

651 / (HOST) HITBY (REMOTE IP OR NAME) USERNAME [DATE & TIME]
"GET ...HTTP/1.0" ERC BYTESXFRD REFERRER

653 /

Fig. 6




Fig. 7

WEB SERVER EMPLOYING MULTI-HOMED, MODULAR FRAMEWORK

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to Web servers, i.e., server software for delivering content through the World Wide Web.

2. State of the Art

The Internet, and in particular the content-rich World Wide Web ("the Web"), have experienced and continue to experience explosive growth. The Web is an Internet service that organizes information using hypermedia. Each document can contain embedded reference to images, audio, or other documents. A user browses for information by following references. Web documents are specified in HyperText Markup Language (HTML), a computer language used to specify the contents and format of a hypermedia document (e.g., a homepage). HyperText Transfer Protocol (HTTP) is the protocol used to access a Web document.

Web servers offer Web services to Web clients, namely Web browsers. Primarily, Web servers retrieve Web documents and send them to a Web browser, which displays Web documents and provides for user interaction with those documents. Unlike Web browsers, of which there are many, the number of commercially available Web server packages, although growing, remains small. Currently, popular Web servers include those available from Netscape Communications, National Center for Supercomputing Applications (NCSA), and CERN.

Ideally, a Web server should be able to respond to and expeditiously service every connection request it receives, regardless of the volume of requests. The job of getting the request to the server and the reply back to the requester falls to the telecommunications infrastructure. In reality, however, because of machine limitations, there is a limit to the number of requests the server can serve within a given period of time without slowing down the machine to an extent that, from the viewpoint of the users, is perceptible, irritating, or simply intolerable. Web servers have typically followed one of two extreme approaches, either "rationing service" by denying further requests once some limiting number of requests are already pending, or attempting to service all requests received and hence slowing the machine to a crawl during extremely busy periods.

Web servers, most of which are written for UNIX, often run under INETD ("eye-net-D"), an Internet services daemon in UNIX. (A daemon is a UNIX process that remains memory resident and causes specified actions to be taken upon the occurrence of specified events.) Typically, if more than sixty connection requests occur within a minute, INETD will shut down service for a period of time, usually five minutes, during which service remains entirely unavailable. Such interruptions are clearly undesirable, both on the part of the client-requester and the content provider. In the case of other Web servers, when the server is brought up, some fixed number of copies of the server, e.g. 50, are started. Up to 50 simultaneous connections may therefore be handled. A request for a 51st connection, however, will be denied, even if many or all of the 50 existing connections are temporarily idle.

Other considerations further complicate the picture of what is desired from a Web server, including considerations such as cost, perception and the very dynamic nature of the Web and Web content. Typically, only large organizations, or

smaller organizations having considerable technical expertise, have their own Web servers. Establishing and running a Web server can entail a significant and ongoing investment in time and money. The alternative is for a person or organization to pay an Internet Service Provider (ISP) to house its content on the Web and make it available through the ISP's Web server. Of course, most organizations (and people) would like to be perceived as not lacking in either money or expertise. In cyberspace, therefore, an important factor in how an organization is perceived is its Web address. A Web address of XYZCorp.com commands, in a manner of speaking, immediate attention and respect, whereas a Web address of ISP.com/XYZCorp does not, at least to the same extent.

For a person or organization to have the best of both worlds, i.e., house its content on someone else's server but have it appear to be their own, the server must provide the capability of multi-homing. A multi-homed server behaves as multiple servers in one, responding to requests to multiple addresses, e.g. ISP.com, XYZCorp.com, johnsmith.com, etc. Some existing servers are multi-homed. However, the multi-homing capabilities of existing servers are quite limited. For example, although different customers may have different needs and desire different levels of service, in existing multi-homed servers, as regards a particular physical machine, the core functionality offered by the server (whether extensive or more limited) is necessarily the same for each customer.

An entirely different question concerns the extensibility of Web servers. Because the Web is continually in flux, a Web server must provide a mechanism that allows for extensions to be added to the Web server, or face obsolescence. Presently, the commonly-accepted mechanism for extending the capabilities of existing Web servers is one called the Common Gateway Interface (CGI). The CGI specification has emerged as a standard way to extend the services and capabilities of a Web server having a defined core functionality. CGI "scripts" are used for this purpose. CGI provides an Application Program Interface, supported by CGI-capable Web servers, to which programmers can write to extend the functionality of the server. CGI scripts, however, although they may be compiled, are typically interpreted, meaning that they run at least ten times slower, typically, than compiled binary code. As the complexity of the Web, and hence the amount of time spent by a Web server running CGI scripts, increases, the Web server unavoidably suffers a significant performance hit.

What is needed, then, is a Web server that overcomes the foregoing difficulties.

SUMMARY OF THE INVENTION

The present invention, generally speaking, provides a Web server having a multi-homed, modular framework. The modular framework allows extensions to the Web server to be easily compiled into the Web server, allowing the extensions to run natively as part of the server instead of incurring the overhead typical of CGI scripts, for example. The multi-homing capabilities of the Web server provide the appearance to Web users of multiple distinct and independent servers, allowing a small company or individual to create the same kind of Web presence enjoyed by larger companies. In effect, multiple virtual servers run on the same physical machine. The Web server as a whole is easily extensible to allow additional capabilities to be provided natively within the Web server itself. Furthermore, each virtual server is independently configurable in order to turn

different capabilities on or off or to modify operation of the virtual server. The Web server is also provided with enhanced security features, built-in animation capability, and other features that afford maximum flexibility and versatility.

BRIEF DESCRIPTION OF THE DRAWING

The present invention may be further understood from the following description in conjunction with the appended drawing. In the drawing:

FIG. 1 is a block diagram of one embodiment of the Web server of the present invention;

FIG. 2 is a block diagram illustrating the manner in which the present Web server handles connection requests;

FIG. 3 is an example of a database-like data structure that allows extensions to the Web server to be easily added so as to run natively as part of the Web server;

FIG. 4 is an example of a portion of the master configuration file of FIG. 1;

FIG. 5 is an example of a conventional Web server log entry;

FIG. 6 is an example of a Web server log entry in accordance with one aspect of the present invention; and

FIG. 7 is a block diagram illustrating in greater detail the structure of the present Web server.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring now to FIG. 1, an overview of the present Web server will be provided. The Web server is a software package that runs on a physical machine 100, e.g., a super-minicomputer such as a SunServer machine available from Sun Microsystems of Menlo Park, Calif., or on any of a wide variety of suitable UNIX platforms.

Conventionally, a Web server is provided with some fixed feature set. To provide additional capabilities, CGI scripts are used. Furthermore, the logical view of the Web server on the Web is the same as the physical view of the underlying hardware. That is, a single physical machine running a conventional Web server appears on the Web as a single logical machine.

The present Web server, on the other hand, although it runs on a single physical machine 100, appears on the Web as multiple virtual hosts VH1 through VHn. Each virtual host has a separate configuration sub-file (sub-database) C1, C2, etc., that may be derived from a master configuration file, or database, 110. A defaults file 111 is used to provide plug-in extensibility of the Web server. The defaults file is compiled as part of the Web server and expresses at the program level what the program is to do to provide various kinds of functionality. The configuration sub-files, on the other hand, are text files that may be used to enable or disable different functions for each virtual host.

Each virtual host also has its own separate log file L1, L2, etc. This feature allows for users of different hosts on the same machine to have access to up-to-the minute log information without allowing access to logs of other pages—a very useful feature for ISPs, for example. Each virtual host is capable of servicing many simultaneous connections. The number of allowable simultaneous connections is configurable and may be limited to a predetermined number, or may be limited not by number but only by the load currently experienced by the physical machine. The number of maximum allowable connections or the maximum allowable machine load may be specified in the configuration file.

As described in greater detail in connection with FIG. 5, each configuration file C1, C2, etc., may have an access rules database 113, including a Allow portion 115, a Deny portion 117, or both. Using the access rules database 113, the Web server is able to perform many of the functions of a conventional firewall, thereby enhancing network security.

The Web server is self-daemoning, meaning that it is not subject to the limitations ordinarily imposed by the usual Internet daemon, INETD. Referring to FIG. 2, when the Web server is brought up, it first reads in the master configuration file and then becomes a daemon and waits for connection requests. When a connection request is received, the Web server “forks a copy” of itself, using the forking capabilities of UNIX, to create a virtual host VHn to handle that connection request. When a parent process is forked to create a child process, the child process runs off the same base code as the parent process but using its own set of variables. In particular, the child process will typically use its own sub-database from within the master configuration database to determine the configuration of that particular virtual host. Child processes are created “on demand” as connection requests are received and complete running as service of those connection requests is completed.

As described in greater detail hereinafter, animation capabilities, instead of being an add-on feature provided through a CGI script, are built into the server itself. Hence, in FIG. 2, when a request is received on a connection to process an animation file 201, the child process VHn belonging to that connection is able to process the animation file directly.

Animation is one example of a recent enhancement now supported by some Web servers, although typically through CGI and not directly. Numerous and varied other enhancements are sure to follow. Hence, the present Web server, although it provides a default feature set, allows for that default feature set to be readily expanded.

The default feature set is defined in the defaults file 111, an example of which is shown in FIG. 3. The defaults file is, in effect, a database that matches different types of requests with modules available for handling those requests. The database is a table having, in the illustrated embodiment, four columns. In the first column is the file extension received as part of the request. Examples of file extensions are “ani” (animation), “cgi,” “html,” “shtml” (server-parsed HTML), etc. In the second column is a type identifier used internally by the Web server, e.g., TY_ANIMATION, TY_CGI, etc. In the third column is type information returned to the Web browser, e.g., “text/html,” “image/gif,” etc. In some instances, this type is not clearly defined in advance but depends on the module called to process the request. The type returned to the browser may therefore be NULL. Finally, in the fourth column is the module to be called to process the request, e.g., ProcessANI (an animation player), ProcessCGI, etc.

The defaults table of FIG. 3 serves as the foundation of a modular framework that allows extensions to the Web server to be easily added so as to run natively as part of the Web server. In one embodiment, extensions are added by modifying the source code file containing the defaults table, to add a new row to the table, for example, or to modify an existing row. A new module having new capabilities might be substituted for an existing module, for example. The defaults file is then recompiled and linked to the Web server. Preferably, a software development kit is provided with the Web server to facilitate this process.

Hence, a mechanism is established that allows modules to be added without the need to rebuild the entire Web server.

In accordance with this mechanism, adding an additional server feature involves the following steps:

1. Edit the defaults file (the list of processes) and add an entry for the specific file extension. Add the procedure name to be used for processing this new type. Add an “#include” file defining the procedure at the beginning of the defaults table.
2. Write the procedure. The procedural call to the process will pass two parameters: a File Descriptor to send output to; and a file Name of the document/file to process. The procedure should be written accordingly.
3. Compile the procedure, link it into the Web server, and test.

The modular framework just described governs the capabilities of the overall Web server. All customers may not want to pay for use of a high-performance Web server having a comprehensive feature set. The multi-homing capabilities of the present Web server, however, allows the same code to be configured differently for each of multiple virtual hosts, thereby catering to the preferences (and budgets) of different classes of customers. The effect is the same as offering a range of different Web server products of different capabilities, but with the distinct advantage that only a single code package need be produced and maintained.

Referring again to FIG. 1, apart from the defaults database, each virtual host has its own configuration sub-file C1, C2, etc. Within the configuration sub-file of each virtual host is a text-file database, not compiled, that is similar in concept to the database of FIG. 3. However, the process to be called for a particular type of request may vary, virtual host to virtual host, from the default process specified in the defaults database. On a virtual host that is configured to have a more limited feature set, for example, CGI scripts may not be allowed. Therefore, the process specified for CGI, instead of ProcessCGI, might be NULL. Similarly with other capabilities. The differences between the configuration database and the defaults database, instead of limiting the capabilities of the virtual host, may actually increase the capabilities of the virtual host. For example, for HTML requests, instead of the usual SendFile process, the process ProcSHTML may be specified. This configuration would allow all HTML files to include other files, cause execution of other programs, etc., resulting in a dynamic document rather than a static document.

Furthermore, the Web server may be a “client application” of a license server, described in U.S. patent application Ser. No. 08/607,081 AUTOMATED SYSTEM FOR MANAGEMENT OF LICENSED SOFTWARE now U.S. Pat. No. 5,790,664, filed on even date herewith, incorporated herein by reference. Using the system described in the foregoing application, features of the Web server may be enabled or disabled on a feature-by-feature and virtual host-by-virtual host basis.

An example of a portion of a master configuration file is shown in FIG. 4. Within the master configuration file database, different portions of the file form sub-databases for different virtual hosts. Each sub-database may specify a root directory for that particular virtual host.

Also as part of the configuration file of each virtual host, an access rules database may be provided governing access to the virtual host, i.e., which connections will be allowed and which connections will be denied. Many of the features more commonly included in a firewall may therefore be included in the Web server itself. The syntax of the access rules database is such as to allow great flexibility in specifying not only what machines are or are not to be allowed access but also when such access is allowed to occur. The

access rules database may have an Allow portion, a Deny portion or both. If the access rules database has an Allow portion, only connections from machines matching the Allow rules will be allowed, regardless of whether there is also a Deny portion. If there is a Deny portion but no Allow portion, then connections from machines matching the Deny rules will be denied and all other connections will be allowed. Machines may be specified by name or by IP address, and may include “wildcards,” address masks, etc., for example: MisterPain.com, *.srmc.com, 192.168.0.* , 192.168.0.0/24, and so on.

Time restrictions may be included in either the Allow rules or the Deny rules. For example, access may be allowed from 1 am to 12 pm; alternatively, access may be denied from 12 pm to 1 am. Also, rules may be given identifiers, such as RULE1, RULE2, etc., and repeated elsewhere within the configuration sub-file of the virtual host.

All access rules must be satisfied in order to gain access to a virtual host. Depending on the virtual host, however, further levels of access scrutiny may be specified within the configuration sub-file. Each successive level of access scrutiny includes all previous levels. The first level of access scrutiny is that all rules must be satisfied, as previously described. The second level of access scrutiny is that the accessing machine must have a DNS (Domain Name Services) entry. Having a DNS entry lends at least some level of legitimacy to the accessing machine. The third level of access scrutiny is that the accessing machine must in addition have a reverse DNS entry. The fourth and most stringent level of access scrutiny is that the forward DNS entry and the reverse DNS entry must match.

If access is granted and a connection is opened, when the connection is later closed, a log entry is made recording information about that access. One important feature of the present Web server is that log entries identify the particular virtual host that was accessed. Another important feature is that log entries identify the “referrer,” i.e., the source of any link that may have been followed to get to the Web site that was accessed. The owner of a Web site may advertise that site through various different channels. In order to determine what works and what does not, in terms of generating hits on a Web site, one must know how someone came to access that Web site.

Referring to FIG. 5, conventional log entries are typically of a format similar to the one shown. The IP address or name **501** of the accessing machine is recorded, followed by the UserName **503**, if any, entered during the course of the access, the date and time **505**, the connection request **507** as it was received, an error code ERC (**509**) specifying any error that may have occurred, and a count **511** of the number of bytes transferred. As shown in FIG. 6, the present Web server logs in addition the identity of the virtual host (**651**) and the referrer (**653**). Log reports may be generated from this information, listing, for example, the top referrers generating hits to a Web site, etc. Log reports may be generated on a per-host basis or may be generated for the physical machine, i.e., all virtual hosts together.

Referring now to FIG. 7, the logical structure of the present Web server is shown in greater detail.

The main execution thread of the Web server is controlled by a daemon. In FIG. 7, the daemon includes elements **701**, **703** and **705**. Although the daemon mode of operation is the default mode, the same code can also be run interactively under the conventional INETD daemon. Hence, when the Web server is first brought up, command-line processing is performed in block **701** to determine the mode of operation (daemon or interactive), which configuration file to read, etc.

For purposes of the present discussion, the daemon mode of operation, which is the default, will be assumed.

Immediately thereafter, the daemon changes user in block 703 so as to become an unprivileged user. This step of becoming an unprivileged user is a security measure that avoids various known security hazards encountered, for example, when CGI scripts or other programs are allowed to run.

Only after the daemon has read the specified configuration file and become an unprivileged user does the daemon actually become a daemon. By daemonizing after the configuration file (e.g., the master configuration file) has been read in, the configuration file in effect becomes "hard coded" into the program such that the program no longer has read it in. The daemon then waits to receive a connection request.

When a connection request is received, the daemon forks a copy of itself to handle the connection request. The daemon then uses a piece of code referred to herein as an INET Wrapper 710 to check on the local side of the connection and the remote side of the connection to determine, in accordance with the appropriate Allow and Deny databases, whether the connection is to be allowed.

First the address and name (if possible) are obtained of the virtual machine for which a connection is requested. Once the local host has been identified by name or at least by IP address, the master configuration database is scanned to see if a corresponding sub-database exists for that local host. If so, the sub-database is set as the configuration database of the local host so that the master configuration database need no longer be referred to. If no corresponding sub-database is found, then by default the master configuration database is used as the configuration database. There may be any number of virtual machines, all independently configurable and all running on the same physical machine. The determination of which virtual host the daemon child process is to become is made in block 705, under the heading of "multi-homing."

Once the daemon child process has determined which host it is, the IWET Wrapper is used to do checking on the remote host, i.e., the host requesting the connection. Recalling the different levels of access scrutiny described previously, first, the configuration database is consulted to determine the level of access scrutiny that will be applied. (The default level of access scrutiny is that no DNS entry is required.) Then, the address and name (if possible) are obtained of the machine requested the connection, and the appropriate level of access scrutiny is applied as determined from the configuration database.

If the remote host satisfies the required level of access scrutiny insofar as DNS entries are concerned, the INET Wrapper gets the Allow and Deny databases for the virtual host. First the Allow database is checked, and if there is an Allow database but the remote host is not found in it, the connection is denied. Then the Deny database is checked. If the remote host is found in the Deny database, then the connection is denied. All other rules must also be satisfied, regarding time of access, etc. If all the rules are satisfied, then the connection is allowed.

Once the connection has been allowed, the daemon invokes HTTP server code 720 that operates in large part in a similar manner as conventional Web server. The HTTP server code 720 processes commands by examining the filename extension associated with the command in block 721 and calling appropriate routines such as routines 723-726 to process those commands. When processing is completed, the connection is closed, if it has not already been closed implicitly.

Several features of the HTTP server code should be noted. The server includes code for handling both client and server proxies and redirects, all as part of the Web server. The server mode is determined in accordance with the configuration file of the virtual host handling the request. The configuration file may specify Client Proxy mode for the virtual host, in which case the virtual host proxies requests directed to the outside world, e.g., Web sites other than the present Web site. The configuration file may specify Server Proxy mode for the virtual host, in which case the virtual host proxies requests to access a different virtual host. The configuration file may specify Redirect mode, in which case the request is redirected to a different specified server, either a different virtual host on the same physical machine or a different physical machine altogether.

The animation player 726 of the present Web server also differs in significant respects from conventional animation players. Fundamentally, to the knowledge of the inventors, no other Web server incorporates an animation player as part of the native Web server code. By incorporating the animation player into the Web server itself, instead of adding the animation capabilities through the use of a CGI script, for example, animations may be handled much more efficiently. Furthermore, unlike conventional animation players, which typically just send a sequence of graphics, the present animation player provides full-fledge programming capability, including the ability to send one graphic at a specified time and then a next graphic at a next specified time. When a line is reached in the animation file that calls for a graphic to be displayed at a certain time, if that time has already passed, then that line is ignored by the animation player and the graphic is not sent. Instead the animation player tries to display the next graphic at the next specified time. In this manner, regardless of the speed of the connection, the playing time of the animation is always the same. On a slow connection, the user will see fewer frames, and on a fast connection the user will see more frames. The length of the animation will be the same. Also, labels may be included in the animation, and commands may be used to go to a specified label, or go to the beginning of the animation.

It will be appreciated by those of ordinary skill in the art that the invention can be embodied in other specific forms without departing from the spirit or essential character thereof. The foregoing description is therefore considered in all respects to be illustrative and not restrictive. The scope of the invention is indicated by the appended claims, and all changes which come within the meaning and range of equivalents thereof are intended to be embraced therein.

What is claimed is:

1. A computer-implemented method performed using network server software, comprising the steps of:
 - receiving from a remote requester a connection request that a specified resource from a specified virtual host be served to the remote host;
 - identifying the requested virtual host by determining at least one of the virtual host name and host address from the received connection request;
 - configuring the virtual host in accordance with predetermined configuration parameters in response to the identification of the virtual host;
 - processing the connection request to identify the specified resource as being of one of a plurality of resource types; and
 - consulting a configuration file to identify a code module to be used by the specified virtual host to process resources of the type of the specified resource.

2. The method of claim 1, wherein the connection request is an HTTP protocol request.

3. The method of claim 1, comprising the further steps of: processing the specified resource using an identified code module; and
serving the specified resource to the remote host.

4. The method of claim 3, comprising extending the native capabilities of the network server software by:
providing a new code module to be used to process resources of a particular type when requested by a remote host;
modifying the data structure stored in the configuration file to identify the new code module as a code module to be used to process resources of the particular type, thereby producing a modified data structure; and
compiling the modified data structure and the new code module as part of the network server software.

5. The method of claim 1 wherein the step of configuring the virtual host comprises the step of initiating a child process.

6. The method of claim 1 wherein the step of configuring the virtual host further comprises the steps of:
determining the address of the virtual host from the connection request; and
configuring the virtual host using predetermined configuration parameters corresponding to the determined address.

7. The method of claim 6 wherein the configuration parameters determine the level of access scrutiny to be applied to the connection request.

8. The method of claim 1 wherein the configuration file parameters are user configurable, thereby allowing selective enablement of a plurality of security features.

9. The method of claim 1 wherein the configuration file is user configurable, thereby allowing selective enablement of a plurality of code modules.

10. A computer-implemented method performed using network server software running on a single physical machine, comprising the steps of:
receiving from a remote host a connection request requesting that a specified resource be served to the remote host from a specified one of a plurality of virtual local hosts;
processing the connection request to identify the specified virtual local host; and
consulting a data structure stored in a configuration file to identify, for the specified virtual local host, a code module to be used to process resources of the type of the specified resource, the data structure identifying for each of the plurality of virtual local hosts and for each of a plurality of types of resources a code module, if any, to be used by that virtual local host to process resources of that type.

11. The method of claim 10, wherein a null entry in the data structure indicates that resources of a corresponding type are not allowed to be processed by a corresponding virtual local host.

12. A method of providing a plurality of virtual hosts on a single computing platform, comprising the steps of:
receiving from a remote requester a connection request to a specified virtual host;
identifying the specified virtual host;
configuring the virtual host in accordance with predetermined configuration parameters in response to the identification of the virtual host; and
processing the connection request using the configured virtual host.

13. The method of claim 12 wherein the step of identifying the specified virtual host is performed by determining from the received connection request at least one of: the virtual host name and virtual host address.

14. The method of claim 12 wherein the step of configuring the virtual host further comprises the steps of:
initiating a child process; and
consulting a configuration file to determine configuration parameters used to configure the child process.

15. A computer-implemented method performed using network server software, comprising the steps of:
receiving from a remote host a connection request requesting that a specified resource be served to the remote host;
processing the connection request to identify the specified resource as being of one of a plurality of resource types;
consulting a configuration file to identify a code module to be used to process resources of the type of the specified resource, the data structure identifying for each of a plurality of types of resources a code module to be used to process resources of that type;
processing the specified resource using an identified code module;
serving the specified resource to the remote host; and
extending the native capabilities of the network server by:
providing a new code module to be used to process resources of a particular type when requested by a remote host;
modifying the data structure stored in the configuration file to identify the new code module as a code module to be used to process resources of the particular type, thereby providing a modified data structure; and
incorporating the modified data structure and the new code module as part of the network server software.