




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 17/30	A1	(11) International Publication Number: WO 97/46956 (43) International Publication Date: 11 December 1997 (11.12.97)
<p>(21) International Application Number: PCT/US97/09295</p> <p>(22) International Filing Date: 6 June 1997 (06.06.97)</p> <p>(30) Priority Data: 60/019,303 7 June 1996 (07.06.96) US</p> <p>(71) Applicant: AT & T CORP. [US/US]; 32 Avenue of the Americas, New York, NY 10013-2412 (US).</p> <p>(72) Inventor: RAO, Chung-Hwa, Herman; 4304 Springbrook Drive, Edison, NJ 08820 (US).</p> <p>(74) Agents: DWORETSKY, Samuel, H. et al.; AT & T Corp., P.O. Box 4110, Middletown, NJ 07748 (US).</p>	<p>(81) Designated States: CA, JP, MX, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).</p> <p>Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>	

(54) Title: INTERNET FILE SYSTEM



(57) Abstract

A method and a system in which a computing system transparently accesses resources connected to the Internet. A memory of the computing system contains an operating system, a cache associated with the operating system, a shared library and an access server. The shared library is responsive to a system call for a file by determining whether a path name for the file is located under a personal name space in the memory and by issuing a request for retrieving the file from an Internet resource based on the path name located under the personal name space when the file is not stored in the cache and has a path name located under the personal name space. An access server, in response to the request from the shared library, selects an appropriate access protocol for retrieving the file from the Internet resource and retrieves the file from the Internet resource. The shared library then issues the system call to the operating system when the access server retrieves the file. The access server restores the file to the Internet resource when the application closes the file.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

INTERNET FILE SYSTEM

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application claims priority to United States provisional patent application Serial No. 60/019,303 to C.H. Rao, filed June 7, 1996, and entitled "Integrating The Internet Into A File System."

5

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to the field of computing. More particularly, the present invention relates to a method and a system for accessing resources that are available over the Internet.

10

2. Description of the Related Art

As more resources become available over the Internet, it is becoming increasingly difficult to locate, manage and integrate resources such as Web pages, Gopher information, Network News, released documents and technical reports, public domain software, and collaborative sources and drafts. Resources connected to the Internet are accessed using a naming scheme that is defined by the Uniform Resource Locator (URL) protocol, which has become a de facto standard. Internet resources are also accessed using other standard access protocols, such as the File Transfer Protocol (FTP) for file servers, the Hypertext Transport Protocol (HTTP) for web servers, the gopher protocol

15
20

(GOPHER) for gopher servers, the Network News Transfer Protocol (NNTP) for news servers, and the Remote Shell (RSH) for file servers. Each of these particular access protocols define access mechanisms for retrieving resources from servers that are connected to the Internet. Internet browsers, such as

5 Mosaic and Netscape, have been introduced for conveniently searching networks and retrieving information, but while perfectly adequate for searching and retrieving resources, such conventional browsers are limited and do not lend themselves for integration into larger toolsets.

Several systems have been developed for making access to files

10 over the Internet easier and more efficient. For example, the Andrew File System provides a wide area file service by extending the scope of file systems from a Local Area Network (LAN) to the Internet. However, the Andrew File System uses its own protocol, not the standard Internet protocols previously mentioned and, consequently, requires an Andrew file server.

15 Another system that has been developed for accessing files over the Internet is the Bell Laboratory Plan 9 which permits mounting of Internet file servers to local name spaces by introducing a new operating system. Plan 9 supports a per-process name space and a message-oriented file system protocol so that a file system-like interface is implemented. While different

20 access protocols are accommodated, the Plan 9 operating system is not transparent to existing application tools and kernels and the applications running on the kernels require modification for use with Plan 9.

The Alex File System provides transparent read access to Internet resources by being built on top of the existing Network File System (NFS), with an Alex server being added as a Virtual File System (VFS) interface. Hosts on a LAN use the NFS protocol for sending calls to the Alex server.

- 5 The calls pass through an NFS client on a kernel that is local to the NFS server kernel on which the Alex server is running. The calls are then upcalled from the NFS server kernel to the Alex server. Nevertheless, the NFS protocol causes the Alex system to be limited because, while the NFS protocol implements file services, no information has been provided regarding a process
- 10 initiating a request, which is important for implementing an authenticated access. Consequently, the Alex system only supports limited file-oriented file access, such as an anonymous FTP access.

- The Jade File System provides a uniform way for naming and accessing files in an Internet environment using a new personal name server
- 15 that integrates a heterogeneous collection of underlying existing file systems that may not be modified because of autonomy. A private name space can be defined by each user that supports two features: multiple file systems are allowed to be mounted under one directory, and one logical name space is permitted for mounting other logical name spaces. However, the Jade File
- 20 System does not use URL naming. Instead, a separate name server is provided for a pathname. Further, individual access servers must be mounted and the program running on the host must be recompiled before a user can access a file system.

The Multiple Dimensional File System (n-DFS) is a logical file system that allows new services to be added to underlying file systems without requiring applications or the operating system kernel to be modified. The n-DFS is layered on top of existing physical file systems, but does not provide
5 convenient Internet access using existing protocols.

What is needed is an Internet file system that transparently uses existing protocols in connection with existing operating system kernels, applications and file servers for accessing Internet resources.

10 SUMMARY OF THE INVENTION

The present invention provides an Internet file system that transparently uses existing protocols in connection with existing operating system kernels, applications and file servers for accessing Internet resources. Consequently, an application programmer can use the same system calls to
15 handle files, whether stored locally or at Internet resources. The advantages of the present invention are provided by a method and a system in which a computing system having a processor and a memory transparently accesses resources connected to the Internet. The memory of the computing system contains an operating system, a cache associated with the operating system, a
20 shared library and an access server. The shared library is responsive to a system call for a file by determining whether a path name for the file is located under a personal name space in the memory and by issuing a request for retrieving the file from an Internet resource based on the path name located

under the personal name space when the file is not stored in the cache and has a path name located under the personal name space. In response to the request from the shared library, the access server selects an appropriate access protocol for retrieving the file from the Internet resource and retrieves the file from the Internet resource. The shared library then issues the system call to the operating system when the access server retrieves the file. The access server restores the file to the Internet resource when the application closes the file.

BRIEF DESCRIPTION OF THE DRAWING

10 The present invention is illustrated by way of example and not limitation in the accompanying figures in which like reference numerals indicate similar elements and in which:

 Figure 1 is a block diagram of the architecture of a system according to the present invention;

15 Figure 2 is a block diagram of a system having an architecture that supports coherent sharing for different applications according to the present invention;

 Figure 3 illustrates the concept of a personal name space (PNS) according to the present invention;

20 Figure 4 shows a diagram illustrating the mounting of remote file systems according to the present invention;

 Figure 5 shows a flow diagram of a validation process according to the present invention; and

Figure 6 shows a flow diagram of the details of a step of the validation process of Figure 5.

DETAILED DESCRIPTION

5 The present invention relates to an Internet File System (IFS) that allows users to manage Internet files in a Personal Name Space (PNS) by transparently integrating existing protocols, while requiring no modification of software or change in management of remote servers. The present invention introduces a logical layer between an operating system and applications running
10 on the operating system that presents the applications with exactly the same system call interface that the underlying operating system provides. The logical layer of the present invention runs in the same address space as the application without requiring operating system modifications.

 According to the invention, each user selects Internet files for
15 access by organizing the files within a PNS and thereby avoids the complexity of maintaining a system-wide global name space. According to the present invention, a PNS is private and remains undefined to all other users. Only the owner the PNS has access to the PNS. A PNS can be "mounted" or attached to a shared name space or to another name space that supports URL naming.
20 Internet files can be accessed directly using the URL names of the files as pathnames. For example, a home page identified by `http://www.att.com/research.html` is named `~/IFS/http://www.att.com/research.html`, where `~/IFS` is the mount point on a

local file system for a PNS. Additionally, a PNS supports mount operations that attach name spaces of Internet file servers to nodes of the PNS using access protocols, such as File Transport Protocol (FTP). Thus, a PNS allows a file system to name and access files, and to perform mounting operations that

5 support location-independent naming.

Figure 1 is a block diagram of the architecture of a system according to the present invention. System 10 includes a processor 11 coupled to a memory 12. Memory 12 includes memory space for storing an operating system kernel 13, a cache 14 for kernel 13, at least one application 15, such as

10 a browser, a shared library 16 and an access server 17. Processor 11 and memory 12 can be physically located within a single device or can be configured as a plurality of processors and memory devices that are distributed using well-known LAN techniques. A plurality of applications can also be run by system 10 using well-known techniques.

15 When application 15 sends a request or system call to kernel 13 for opening, reading, writing, closing, etc., a file, for example, the request is received by a logical layer 22 that includes shared library 16 and access server 17. Preferably, shared library 15 presents applications with a complete POSIX system interface. Shared library 16 is linked by application 15 and runs in the

20 address space of application 15. All of the necessary functionality for handling system calls is provided by shared library 16. According to the invention, shared library 16 includes a capability for handling a PNS, that is, shared library 16 is capable of determining whether the pathname for a file specified in

a system call is located under a PNS for a user of the application. The determination is made by examining the mount point (the iroot). If the specified file is located under a PNS, shared library 16 determines whether the requested file is an Internet file or is stored locally. If the file is an Internet
5 file, shared library 16 communicates the request from the application 15 to access server 17 through an interprocess communication (IPC) in a well-known manner. For computing systems that provide dynamic shared libraries, application 15 may invoke shared library 16 by simply redefining a library search path to include shared library 16. For computing systems without
10 dynamic shared libraries, applications must be relinked with shared library 16.

Access server 17 selects the appropriate access protocol for the specified file, connects to remote server 19 via the Internet 18, and obtains the specified file which is stored in memory 20 from remote server 19. Access server 17 is capable of performing an authenticated access of an Internet server
15 "on behalf of" a user who is running application 15. That is, access server 17 can access remote file servers using the user's password information. To do this, shared library 16 is capable of authenticating itself to access server 17.

Once the specified file is obtained, access server 17 caches the file locally in cache 14. Access server 17 contains a plurality of agents, with
20 each agent implementing a particular access protocol. Figure 1 shows exemplary agents 17a, 17b and 17c. Agents may contact remote servers directly or through proxy servers, such as proxy server 21. Access server 17 performs three basic functions of mounting a remote file system on a local

name space, retrieving a remote file and restoring a file to its primary server.

The functions of retrieving and restoring are mapped to corresponding commands supported by specific protocols, while the function of mounting is a local operation.

5 Access server 17 maintains a mounting database for each PNS defined within system 10. Information required for accessing an Internet file, such as the access protocol, the host, authentication information, remote path information, etc., are stored in the database. Each entry for a PNS also includes the mount point for the PNS. Access server 17 refers to the mounting
10 information in the mounting database for locating, retrieving and storing the file in cache 14 of the operating system kernel 13. Access server 17 sends an IPC message to shared library 16 when the specified file has been cached. Shared library 16 then retrieves the specified file from cache 14. When the file is closed by application 15, access server 17 restores the file back to remote
15 server 19. When the specified file is a local file, that is, stored locally elsewhere within system 10, shared library 16 directly accesses the file.

 Preferably, an Internet file is cached in its entirety on the local file system when it is opened because remote file servers are contacted only for system file "opens" and "closes" and not for individual reads and writes. In
20 this manner, the total network overhead incurred by transmitting a file is lower when the entire file is retrieved rather than as a series of system requests and responses for individual pages. Since accesses to files over the Internet are expensive in terms of overall system resource consumption, when a cached

copy of a file is accessed, the present invention determines whether the data of the cached are consistent with master copy data stored in the remote server.

While a single access server 17 is shown in layer 22 in Figure 1, a plurality of access servers can be used. For example, a basic configuration of the present invention can include an access server 17 for each system user and a system wide name server for locating each user's access server.

Alternatively, groups of users can share a particular access server 17. The present invention also allows file servers located on the Internet to be mounted as user-level implementations, as shown in Figure 2, rather than as a kernel implementation as in conventional systems. In Figure 2, each shared library 16 that is associated with a process 23 directly accesses access server 17, rather than the application being processed through an associated kernel, the NFS protocol (not shown) and an access server (not shown), as is presently done in the prior art.

Figure 3 illustrates the concept of a PNS within a local file system 30. A first PNS 31 is shown mounted to local file system 30 at /home/bob/IFS and owned by the user Bob. A second PNS 32 is shown mounted at /home/stephanie/tmp/IFS and owned by the user Stephanie. A PNS is defined by the directory pair (MountPoint, CacheDirectory), where MountPoint is a local pathname in the file system to which PNS is mounted or attached, and CacheDirectory is a local directory in which remote files and directories are cached. Essentially, MountPoint indicates where the mount is pointing. In Figure 3, the mount point for PNS 31 is /home/bob/IFS and the

mount point for PNS 32 is /home/stephanie/IFS. CACHEDIRECTORY is needed in a file system structure as a place to put a cache file. To simplify the task of maintaining multiple PNSs, the present invention uses a single directory called an IRoot for a MountPoint and a CacheDirectory. An IRoot is a mount point
5 for a PNS and a physical directory on the local file system where remote files and directories are cached. That is, an IRoot includes the mount point and the cache directory. A user defines a PNS by specifying the IRoot using the command iroot:

```
$ iroot pathname
```

10 As shown in Figure 3, the IRoot for PNS 31 is defined to be /home/bob/IFS, while the IRoot for PNS 32 is defined to be /home/stephanie/tmp/IFS. PNSs 31 and 32 are each mounted on local name space 30. Files and directories located under PNS 31 (/home/bob/IFS) are accessible only to the particular user defining PNS 31. To other users, such as
15 the user who defined PNS 12, the files and directories under /home/bob/IFS are undefined.

The Universal Resource Locator (URL) naming scheme is embedded within the system of the present invention. The URL naming scheme has the form:

20 protocol://user:password@host:port/path

where "protocol" is the access protocol, "user" is an optional user name, "password" is an optional password, "host" is the fully qualified domain name of a network host or its IP address, "port" is an optional port number to which

to connect to, and "path" is the path used by the remote server for accessing the desired resource. The access protocol not only provides the key to accessing Internet resources located on remote servers, but also hides the heterogeneity of a system on which the accessed server is located. That is, once an access protocol becomes available, it is possible to access resources provided by a server using the protocol without regard to the machine type or the operating system of the server. Consequently, the present invention supports a heterogeneous collection of access protocols.

Table I shows exemplary pathnames and corresponding URL names with "~ IFS" designating the IRoot that includes the mount point and cache directory.

TABLE I

IFS Pathname	URL Name
~/IFS/ftp://ftp.ai.mit.edu/pub/README	FTP://ftp.ai.mit.edu/pub/README
~/IFS/rsh://king/src	rsh://king/src
5 ~/IFS/http://research.att.com/people/list.html	http://research.att.com/people/list.html
~/IFS/nntp://ulysses/comp.os.research/5364	nntp://ulysses/comp.os.research/5364
~/IFS/gopher://csie.nctu.edu.tw	gopher://csie.nctu.edu.tw
~/IFS/ftp://john:xxx@gryphon/home/john	ftp://john:xxx@gryphon/home/john
~/IFS/http://jade:8000/project/DAF.html	http://jade:8000/project/DAF.html
10 ~/IFS/http/192.127.159.32:8000/IFS.html	http://192.127.159.32:8000/IFS.html

The following exemplary pseudo code opens a home page using the system call `open()`, reads a file located at a URL name using the system call `read()`, processes the contents of the file as raw data, and closes the file using the system call `close()`:

```

15         fd = open("~/IFS/http://research.att.com/books.html",
0_RDONLY);
        while ((num = read(fd, buf, nbytes)) > 0)
        {
20             /* process data in buf */
        }
        close(fd)

```

The present invention also provides that an Internet resource can be located, managed and interacted with using existing commands and tools.

For example, the command line

```
$ ls ~/IFS/nntp:/ulysses/comp.os.research/*
```

5 lists articles in the news group "comp.os.research" that are stored on the News server "ulysses." Additionally, the command line

```
$ grep thread ~/IFS/nntp:/ulysses/comp.os.research/*
```

searches for the keyword "thread" in the News group "comp.os.research".

Mounting allows shortcuts to be created when naming files by
10 providing indirect naming and also maintains authentication information for contacting remote servers. Additionally, mounting allows a proxy server to be specified through which the present invention accesses a desired file server. Both PNSs and mount points within a PNS are defined on a per-user basis. According to the invention, a mount operation allows a user to attach a name
15 space, such as PNS, to another name space. The mount operation can attach a PNS on top of a global name space, and can attach a remote file system to a PNS.

A mount point is specified by the command mount:

```
$ mount URLName MountPoint [ProxyServer:Port]
```

20 [ValidPeriod]

where "URLName" is the URL name of a file/directory accessible over the Internet, "MountPoint" is the defined mountpoint in a PNS (for example, IFS in

Figure 3), "ProxyServer" is an optional proxy server parameter, and "Port" is an optional parameter specifying the tcp port of the proxy server.

When the ProxyServer and Port parameters are specified, the present invention accesses the specified remote server through the specified proxy server. The parameter ValidPeriod (in seconds) specifies a time period during which cached files under the mount point are considered to be valid. Specifically, when a file is accessed, the cached file can be used when its age is still within the specified ValidPeriod. If ValidPeriod time has expired, the file must be obtained from the network. The URLName may contain a user name and a password, and when present, the present invention authenticates access to the remote server on behalf of the user using the password. If the value of password is "-", the mount command preferably requires the user to enter the password at the next prompt.

Figure 4 illustrates how remote file systems are mounted according to the present invention. In Figure 4, remote file systems that use different protocols are mounted to the IRoot /home/bob/IFS (PNS 31 in Figure 3). File server 40 at gryphon.att.com is mounted to /home/bob/IFS as follows:

```
$ mount ftp://bob:-@gryphon.att.com.home/bob
/home/bob/IFS/research
Password: *****
```

This command mounts the directory bob, located in the file server 40 on gryphon.att.com, on PNS 31 at the node ~/home/bob/IFS/research using the ftp protocol. In this example, the command mount prompts Bob to enter a

password because "-" has been specified in the password field. Whenever the user Bob access files under /home/bob/IFS/research, the present invention retrieves the corresponding files from gryphon.att.com on behalf of Bob, using Bob's entered log-in and password. Another example of the mount command

5 is:

```
$ mount rsh://king/home/rao/project /home/bob/IFS/project
```

For this example, the directory project, located in the file server 41 at king, on PNS 31 at the node ~ /home/bob/IFS/project via the rsh protocol, which provides authenticated accesses. Users specify user names and passwords
10 inside URLName. The system takes authentication information from the .rhosts file in the user's home directory. This avoids passing password information through the network.

In the following example of the mounting command, a mount point is created using an anonymous ftp access:

```
15 $ mount ftp://tsx-11.mit.edu/pub/linux /home/bob/IFS/src/linux  
radish:8000
```

where the proxy server on radish on the port 8000 is used for accessing tsx-11.mit.edu. The remote directory linux in the file server 42 on tsx-11.mit.edu for this example is mounted on the node linux in PNS 31.

20 The present invention provides two methods for controlling cache validation processes. First, an attribute ValidPeriod, assigned using mount commands, is associated with each directory. Each directory inherits ValidPeriod from its parent directory. A cached copy of a file under a created

directory is treated as valid, without checking its primary copies, if the copy was retrieved during ValidPeriod. For example, if ValidPeriod for a cached file is set to be 7200 seconds, the cached copy is considered as valid for 7200 seconds. When ValidPeriod is "-1", the cached copy is always considered
5 valid. When ValidPeriod is "0", the cached copy is always invalid. ValidPeriod can be changed dynamically like many other attributes associated with directories.

For the second validation control method, a per-process mask, referred to as ValidMask, is provided for overwriting the ValidPeriod of a
10 directory. A child process inherits ValidMask from its parent process and a built-in shell command vmask is used for defining and updating the process ValidMask. The per-process ValidMask has higher priorities than the per-directory ValidPeriod. That is, when dealing with a particular type of file, the valid time is determined according to the mounting, i.e., the file desired to be
15 accessed. When considering the purpose for using the file, the valid time is the time of the process according to the application, i.e., the time the application uses the file. Thus, with ValidPeriod, cache validation processes can be controlled according to types of Internet files, such as co-authored papers versus Network News, whereas, ValidMask allows overwriting ValidPeriod on
20 the process level in special circumstances, such as refreshing cache 14.

Figure 5 shows a flow diagram for a validation process 50 according to the present invention. At step 51, it is determined whether the requested information is in the cache. If the requested information is in the

cache, flow continues to step 52 where it is determined whether the cache is valid under ValidPeriod or ValidMask. If so, flow continues to step 53 where the information in the cache is retrieved for the system call. If the information is not in the cache (step 51) or the cache is not valid (step 52), flow continues
5 to step 54 where a message is sent from the shared library to the access server. Flow continues to step 55, where the Internet is accessed, the information is retrieved and then is stored in the cache.

Figure 6 shows a flow diagram 60 of the details of step 52 of process 50 shown in Figure 5. At step 52A, it is determined when the cache
10 was last updated. Then, at step 52B, the last update time is compared with ValidPeriod or ValidMask for determining whether the cached copy is valid.

The present invention identifies itself to a server for authenticating the user. Consequently, shared library 16 and access server 17 are preferably configured inside system 10 as a separate entity. In this way,
15 shared library 16 is aware of the identity of a user. Access server 17 obtains information from shared library 16 for informing servers connected to the Internet because access server 17 is not configured for determining the identity of the user making a system call request. Consequently, shared library 16 identifies itself to access server 17 before the access server can perform user
20 authentication.

The present invention also supports both anonymous and authenticated accesses to Internet servers directly or via proxy servers. According to the invention, users can authenticate themselves to remote servers

using a proper login name and password information for accessing files. An authenticated access includes a handshaking process in which shared library 16 first communicates with access server 17 that a request is to be authenticated. Access server 17 then communicates to shared library 16 that a file is to be
5 created for obtaining a time stamp. Shared library 16 identifies the file to access server 17 once the file has been created. Access server 17 then accesses the file created by shared library 16 for confirming the information.

While the present invention has been described in connection with the illustrated embodiments, it will be appreciated and understood that
10 modifications may be made without departing from the true spirit and scope of the invention.

THE CLAIMS

What is claimed is:

1 1. A computing system comprising:
2 a processor; and
3 memory coupled to the processor, the memory containing,
4 an operating system,
5 a cache associated with the operating system,
6 a shared library responsive to a system call for a file by
7 determining whether a path name for the file is located under a personal name
8 space in the memory and by issuing a request for retrieving the file from an
9 Internet resource based on the path name located under the personal name space
10 when the file is not stored in the cache and has a path name located under the
11 personal name space, and
12 an access server, responsive to the request from the
13 shared library, retrieving the file from the Internet resource,
14 the shared library issuing the system call to the operating
15 system when the access server retrieves the file from the Internet resource.

1 2. The computing system according to claim 1, wherein the access
2 server includes a plurality of agents, each agent implementing a predetermined
3 Internet access protocol, and
4 wherein the access server selects an appropriate access protocol
5 for retrieving the file from the Internet resource.

1 3. The computing system according to claim 2, wherein the shared
2 library and the access server operate transparently in response to the system
3 call.

1 4. The computing system according to claim 3, wherein the shared
2 library receives the system call from an application running in the memory.

1 5. The computing system according to claim 4, wherein the access
2 server restores the file to the Internet resource when the application closes the
3 file.

1 6. The computing system according to claim 4, wherein the access
2 server access the file from the Internet resource using a user's password
3 information.

1 7. The computing system according to claim 6, wherein the shared
2 library authenticates itself to the access server when the access server accesses
3 the file from the Internet resource using the user's password information.

1 8. The computing system according to claim 2, wherein the access
2 server retrieves the file from the Internet resource by directly contacting a
3 remote server associated the Internet resource.

1 9. The computing system according to claim 2, wherein the access
2 server retrieves the file from the Internet resource by contacting a remote
3 server associated with the Internet resource through a proxy server.

1 10. A method for accessing resources connected to the Internet, the
2 method comprising the steps of:
3 receiving a system call for a file at a shared library;
4 determining whether the file is cached locally;
5 determining whether a path name for the file is located under a
6 personal name space;
7 retrieving the file from an Internet resource based on the path
8 name located under the personal name space when the file is not stored in the
9 cache and has a path name located under the personal name space; and
10 issuing the system call to an operating system of a computing
11 system when the file is retrieved from the Internet resource.

1 11. The method according to claim 10, further comprising the step of
2 restoring the file to the Internet resource when the file is closed.

1 12. The method according to claim 11, wherein the step of retrieving
2 the file from the Internet resource includes the step of accessing the file from
3 the Internet resource using a user's password information.

- 1 13. The method according to claim 10, further comprising the step of
- 2 determining whether the file is valid when the file is cached locally based on a
- 3 last time that the file was cached and a valid file cache time.

FIG. 1




FIG. 2

2/4




FIG. 3

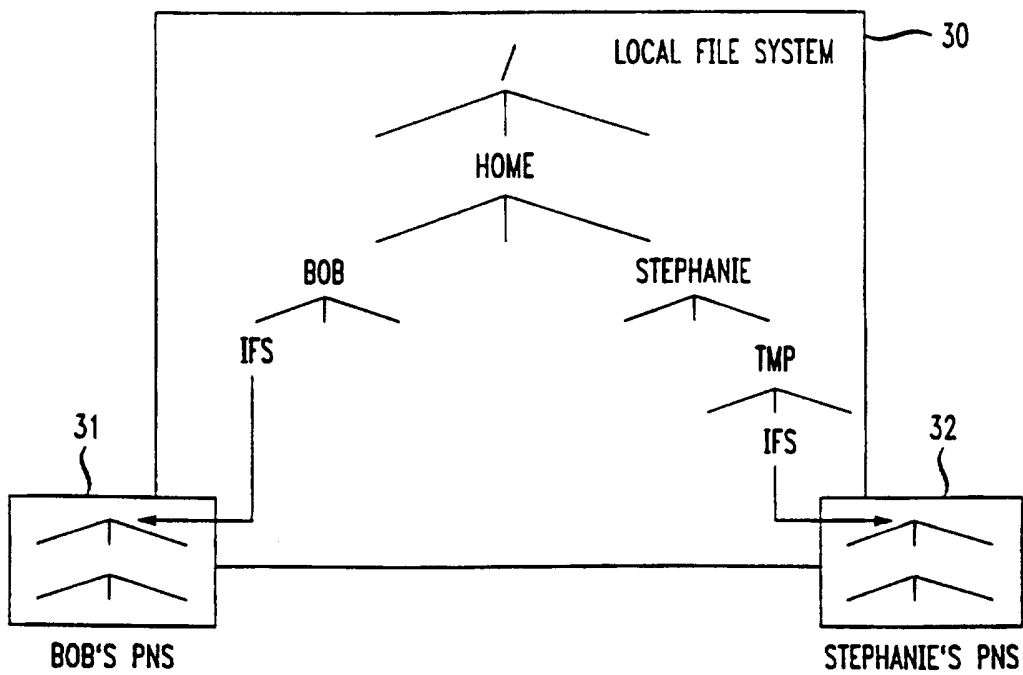


FIG. 4




FIG. 5

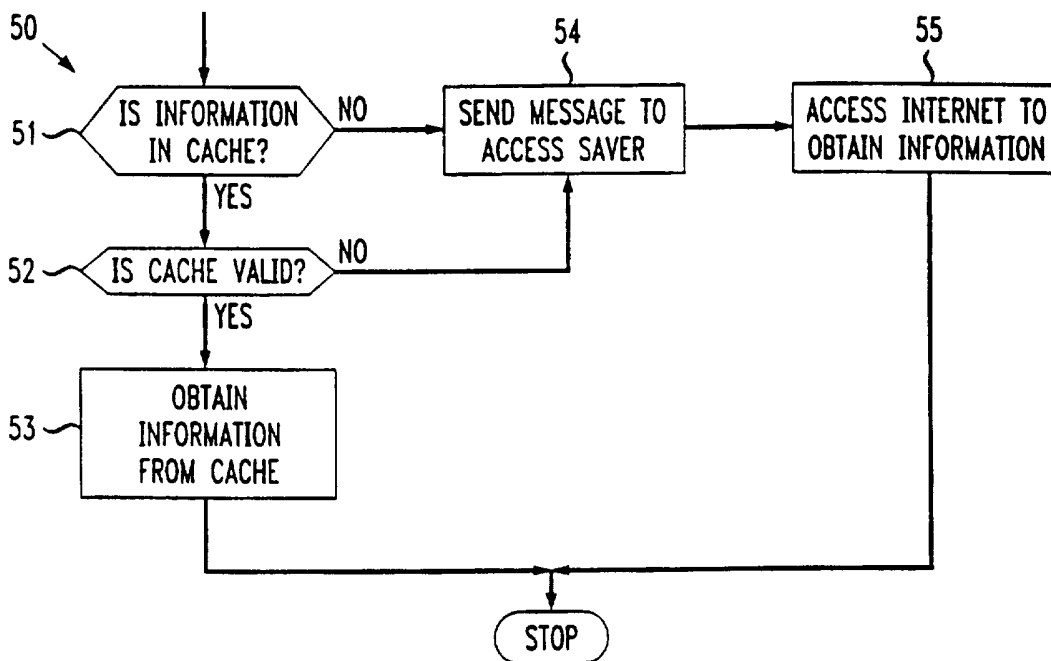



FIG. 6



INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 97/09295

A. CLASSIFICATION OF SUBJECT MATTER IPC 6 G06F17/30		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols) IPC 6 G06F		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practical, search terms used)		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category ^o	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	HERMAN C., RAO AND LARRY L. PETERSON: "Accessing Files in an Internet: The Jade File System" IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, vol. 19, no. 6, June 1993, NEW YORK, USA, pages 613-624, XP000394147	1-4,10, 11
A	see abstract see page 613, column 1 - page 614, column 3 see page 617, column 9 - page 618, column 12 see figures 1-5 --- -/--	5-9,12, 13
<input checked="checked" type="checkbox"/> Further documents are listed in the continuation of box C.		
<input type="checkbox"/> Patent family members are listed in annex.		
^o Special categories of cited documents :		
"A" document defining the general state of the art which is not considered to be of particular relevance		"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier document but published on or after the international filing date		"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)		"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
"O" document referring to an oral disclosure, use, exhibition or other means		"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed		
Date of the actual completion of the international search 11 November 1997		Date of mailing of the international search report 18.11.97
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016		Authorized officer Adkhis, F

4

INTERNATIONAL SEARCH REPORT

Interr. nal Application No
PCT/US 97/09295

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>WELCH B: "A COMPARISON OF THREE DISTRIBUTED FILE SYSTEM ARCHITECTURES: VNODE, SPRITE, AND PLAN 9" COMPUTING SYSTEMS, vol. 7, no. 2, 1 January 1994, pages 175-199, XP000577569 see the whole document ---</p>	1-13
A	<p>ROWE K E ET AL: "Reliability of WWW name servers" COMPUTER NETWORKS AND ISDN SYSTEMS, vol. 27, no. 6, April 1995, page 773-780 XP004013179 see page 774, paragraph 2.1 see the whole document -----</p>	1-13