



(22) Date de dépôt/Filing Date: 2005/10/26

(41) Mise à la disp. pub./Open to Public Insp.: 2006/05/30

(30) Priorité/Priority: 2004/11/30 (US11/000,430)

(51) Cl.Int./Int.Cl. *G06F 17/00* (2006.01),
G06F 17/30 (2006.01)


(71) Demandeur/Applicant:
MICROSOFT CORPORATION, US

(72) Inventeurs/Inventors:
CHRISTIANSEN, NEAL R., US;
THIND, RAVINDER S., US;
HAVEWALA, SAROSH CYRUS, US

(74) Agent: SMART & BIGGAR

(54) Titre : METHODE ET SYSTEME POUR MAINTENIR LA CORRELATION ENTRE UN CHAMP DE NOM ET UN
SYSTEME D'ARCHIVAGE

(54) Title: METHOD AND SYSTEM FOR MAINTAINING NAMESPACE CONSISTENCY WITH A FILE SYSTEM



(57) Abrégé/Abstract:

A method and system for maintaining namespace consistency between selected objects maintained by a file system and a filter associated therewith. A filter monitors selected types of requests (or operations associated therewith) and determines whether the object is within a namespace associated with the filter. The namespace associated with the filter is updated based on a change to the object.

ABSTRACT

A method and system for maintaining namespace consistency between selected objects maintained by a file system and a filter associated therewith. A filter monitors selected types of requests (or operations associated therewith) and determines whether the object is within a namespace associated with the filter. The namespace associated with the filter is updated based on a change to the object.

METHOD AND SYSTEM FOR MAINTAINING NAMESPACE CONSISTENCY WITH A FILE SYSTEM

FIELD OF THE INVENTION

5 The invention relates generally to computers, and more particularly to file systems.

BACKGROUND

With contemporary operating systems, such as Microsoft Corporation's Windows® XP operating system with an underlying
10 file system such as the Windows® NTFS (Windows® NT File System), FAT, CDFS, SMB redirector filesystem, or WebDav file systems, one or more file system filter drivers may be inserted between the I/O manager that receives user I/O
15 requests and the file system driver. In general, filter drivers (sometimes referred to herein simply as "filters") are processes that enhance the underlying file system by performing various file-related computing tasks that users desire, including tasks such as passing file system I/O
20 (requests and data) through anti-virus software, file system quota providers, file replicators, and encryption/compression products.

For example, antivirus products provide a filter that watches I/O to and from certain file types (.exe, .doc, and

the like) looking for virus signatures, while file replication products perform file system-level mirroring. Other types of file system filter drivers are directed to system restoration (which backs up system files when changes are about to be made so that the user can return to the original state), disk quota enforcement, backup of open files, undeletion of deleted files, encryption of files, and so forth. Thus, by installing file system filter drivers, computer users can select the file system features they want and need, in a manner that enables upgrades, replacement, insertion, and removal of the components without changing the actual operating system or file system driver code.

A file system filter may maintain internal metadata for files and directories on a volume. Changes to a volume that a filter is associated with may cause the internal metadata of the filter to be out of sync with the state of the volume. This may cause the filter to behave incorrectly or render it unable to perform its desired function.

What is needed is a method and system for maintaining namespace consistency between selected objects maintained by a file system and a filter associated therewith.

SUMMARY

Briefly, the present invention provides a method and system for maintaining namespace consistency between selected objects maintained by a file system and a filter associated therewith. A process requests access to an object of the file system. A filter monitors selected types of requests (or operations associated therewith) and determines whether the object is within a namespace associated with the filter. The namespace associated with the filter is updated based on a change to the object.

In one aspect of the invention, the request comprises a rename or delete operation. An object is within the namespace associated with a filter if the object has a policy associated with it or if the object is an ancestor of such an object.

In another aspect of the invention, a rename operation may cause the object to be moved and/or renamed.

In another aspect of the invention, an object is accessed by opening a handle to the object. A handle comprises information that identifies a communication channel used to access the object. A flag may be set that indicates that the object should be deleted when all handles to the object are

closed. If the flag remains set after the last handle is closed, the object is deleted.

In one aspect of the invention, maintaining namespace consistency allows filters to apply policies on objects when
5 the objects are renamed and to stop applying policies on objects that are deleted.

Other aspects will become apparent from the following detailed description when taken in conjunction with the drawings, in which:

10

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 is a block diagram representing a computer system into which the present invention may be incorporated;

FIG. 2 is a block diagram representing an exemplary
15 arrangement of components of a system in which the present invention may operate in accordance with various aspects of the invention;

FIG. 3 is a block diagram representing another exemplary arrangement of components of a system in which the present
20 invention may operate in accordance with various aspects of the invention;

FIG. 4 is a block diagram representing another exemplary arrangement of components of a system in which the present invention may operate in accordance with various aspects of the invention;

5 FIG. 5 is a block diagram representing a data structure that may be used to monitor objects that have policies applied to them in accordance with various aspects of the invention;

FIG. 6 is a flow diagram that generally represents actions that may occur to maintain consistency between a
10 filter namespace and a file system namespace for objects of interest in accordance with various aspects of the invention;

FIG. 7 is a flow diagram that generally represents actions which correspond to block 630 of FIG. 6 which may occur when a rename operation is received in accordance with
15 various aspects of the invention;

FIG. 8 is a flow diagram that generally represents actions which correspond to block 745 of FIG. 7 that may occur in updating the data structure in accordance with various aspects of the invention;

20 FIG. 9 is a flow diagram generally representing actions which correspond to block 815 of FIG. 8 that may occur in

adding nodes when an object is moved in accordance with various aspects of the invention;

FIG. 10 is a flow diagram that generally represents actions which correspond to block 825 of FIG. 8 that may occur
5 when updating the data structure to remove prefix nodes associated with the old object name in accordance with various aspects of the invention;

FIG. 11 is a flow diagram that generally represents actions which correspond to block 640 of FIG. 6 that may occur
10 to maintain consistency between a filter namespace and a file system namespace during delete related operations in accordance with various aspects of the invention;

FIG. 12 is a flow diagram that generally represents actions which correspond to block 1115 of FIG. 11 which may
15 occur when a create for delete on close operation is received in accordance with various aspects of the invention;

FIG. 13 is a flow diagram that generally represents actions which correspond to block 632 of FIG. 6 which may
occur when a rename operation is received in accordance with
20 various aspects of the invention;

FIG. 14 is a flow diagram that generally represents actions which correspond to block 1135 of FIG. 11 which may

occur when a set disposition operation is received in accordance with various aspects of the invention; and

FIG. 15 is a flow diagram that generally represents actions which correspond to block 1145 of FIG. 11 which may occur when a cleanup operation is received in accordance with various aspects of the invention.

DETAILED DESCRIPTION

EXEMPLARY OPERATING ENVIRONMENT

Figure 1 illustrates an example of a suitable computing system environment 100 on which the invention may be implemented. The computing system environment 100 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computing environment 100 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 100.

The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for

use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microcontroller-based systems, set top boxes, programmable consumer electronics, 5 network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

The invention may be described in the general context of computer-executable instructions, such as program modules, 10 being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, and so forth, which perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where 15 tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

20 With reference to Figure 1, an exemplary system for implementing the invention includes a general-purpose computing device in the form of a computer 110. Components of

the computer 110 may include, but are not limited to, a processing unit 120, a system memory 130, and a system bus 121 that couples various system components including the system memory to the processing unit 120. The system bus 121 may be
5 any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus,
10 Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

Computer 110 typically includes a variety of computer-readable media. Computer-readable media can be any available
15 media that can be accessed by the computer 110 and includes both volatile and nonvolatile media, and removable and non-removable media. By way of example, and not limitation, computer-readable media may comprise computer storage media and communication media. Computer storage media includes both
20 volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data

structures, program modules, or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the computer.

110. Communication media typically embodies computer-readable instructions, data structures, program modules, or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer-readable media.

The system memory 130 includes computer storage media in the form of volatile and/or nonvolatile memory such as read

only memory (ROM) 131 and random access memory (RAM) 132. A
basic input/output system 133 (BIOS), containing the basic
routines that help to transfer information between elements
within computer 110, such as during start-up, is typically
5 stored in ROM 131. RAM 132 typically contains data and/or
program modules that are immediately accessible to and/or
presently being operated on by processing unit 120. By way of
example, and not limitation, Figure 1 illustrates operating
system 134, application programs 135, other program modules
10 136, and program data 137.

The computer 110 may also include other removable/non-
removable, volatile/nonvolatile computer storage media. By
way of example only, Figure 1 illustrates a hard disk drive
140 that reads from or writes to non-removable, nonvolatile
15 magnetic media, a magnetic disk drive 151 that reads from or
writes to a removable, nonvolatile magnetic disk 152, and an
optical disk drive 155 that reads from or writes to a
removable, nonvolatile optical disk 156 such as a CD ROM or
other optical media. Other removable/non-removable,
20 volatile/nonvolatile computer storage media that can be used
in the exemplary operating environment include, but are not
limited to, magnetic tape cassettes, flash memory cards,

digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 141 is typically connected to the system bus 121 through a non-removable memory interface such as interface 140, and magnetic
5 disk drive 151 and optical disk drive 155 are typically connected to the system bus 121 by a removable memory interface, such as interface 150.

The drives and their associated computer storage media, discussed above and illustrated in Figure 1, provide storage
10 of computer-readable instructions, data structures, program modules, and other data for the computer 110. In Figure 1, for example, hard disk drive 141 is illustrated as storing operating system 144, application programs 145, other program modules 146, and program data 147. Note that these components
15 can either be the same as or different from operating system 134, application programs 135, other program modules 136, and program data 137. Operating system 144, application programs 145, other program modules 146, and program data 147 are given different numbers herein to illustrate that, at a minimum,
20 they are different copies. A user may enter commands and information into the computer 20 through input devices such as a keyboard 162 and pointing device 161, commonly referred to

as a mouse, trackball or touch pad. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, a touch-sensitive screen of a handheld PC or other writing tablet, or the like. These and other input
5 devices are often connected to the processing unit 120 through a user input interface 160 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 191 or other type of display device is also
10 connected to the system bus 121 via an interface, such as a video interface 190. In addition to the monitor, computers may also include other peripheral output devices such as speakers 197 and printer 196, which may be connected through an output peripheral interface 190.

15 The computer 110 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 180. The remote computer 180 may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes
20 many or all of the elements described above relative to the computer 110, although only a memory storage device 181 has been illustrated in Figure 1. The logical connections

depicted in Figure 1 include a local area network (LAN) 171 and a wide area network (WAN) 173, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the computer 110 is connected to the LAN 171 through a network interface or adapter 170. When used in a WAN networking environment, the computer 110 typically includes a modem 172 or other means for establishing communications over the WAN 173, such as the Internet. The modem 172, which may be internal or external, may be connected to the system bus 121 via the user input interface 160 or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 110, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, Figure 1 illustrates remote application programs 185 as residing on memory device 181. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

Maintaining Namespace Consistency

FIG. 2 is a block diagram representing an exemplary arrangement of components of a system in which the present invention may operate in accordance with various aspects of the invention. The components include one or more applications 205, an applications programming interface (API) 210, an input/output (I/O) manager 215, a filter manger 220, a file system 225, and one or more filters 230-232.

The applications 205 may make file system requests (e.g., via function/method calls) through the API 210 to the I/O manager 215. The I/O manager 215 may determine what I/O request or requests should be issued to fulfill each request and send each I/O request to the filter manager 220. The I/O manager 210 may also return data to the applications 205 as operations associated with the file system requests proceed, complete, or abort.

In one implementation, filters comprise objects or the like that when instantiated register (e.g., during their initialization procedure) with a registration mechanism in the filter manager 220. For efficiency, each filter typically will only register for file system requests in which it may be interested in processing. To this end, as part of

registration, each filter notifies the filter manager 220 of the types of I/O requests in which it is interested (e.g., create, read, write, close, rename, and so forth). For example, an encryption filter may register for read and write
5 I/Os, but not for others wherein data does not need to be encrypted or decrypted. Similarly, a quota filter may be interested only in object creates and object writes.

In addition to specifying the types of I/O requests in which it is interested, a filter may further specify whether
10 the filter should be notified for pre-callbacks and post callbacks for each of the types of I/O. A pre-callback is called as data associated with an I/O request propagates from the I/O manager 215 towards the file system 225, while a post-callback is called during the completion of the I/O request as
15 data associated with the I/O request propagates from the file system 225 towards the I/O manager 215.

From each I/O request, the filter manager 220 may create a data structure in a uniform format suitable for use by the filters 230-232. Hereinafter, this data structure is
20 sometimes referred to as callback data. The filter manager 220 may then call and pass the callback data to each filter that has registered to receive callbacks for the type of I/O

received by the filter manager 220. Any filters registered to receive callbacks for the type of I/Os received by the filter manager are sometimes referred to as registered filters.

Typically, the filter manager 220 passes callback data associated with a particular type of I/O request to each registered filter sequentially in an order in which the registered filters are ordered. For example, if the filters 230 and 232 are registered to receive callbacks for all read I/O requests and are ordered such that the filter 230 is before the filter 232 in processing such requests, then after receiving a read I/O, the filter manager 220 may first call and pass the callback data to the filter 230 and after the filter 230 has processed the callback data, the filter manager 220 may then call and pass the callback data (as modified, if at all) to the filter 232.

A filter may be attached to one or more volumes. That is, a filter may be registered to be called and receive callback data for I/Os related to only one or more than one volumes.

A filter may generate its own I/O request which may then be passed to other filters. For example, an anti-virus filter may wish to read a file before it is opened. A filter may

stop an I/O request from propagating further and may instruct the filter manager to report a status code (e.g., success or failure) for the I/O request. A filter may store data in memory and persist this data on disk. In general, a filter
5 may be created to perform any set of actions that may be performed by a kernel-mode or user-mode process and may be reactive (e.g., wait until it receives I/O requests before acting) and/or proactive (e.g., initiate its own I/O requests or perform other actions asynchronously with I/O requests
10 handled by the I/O manager 215).

In one embodiment, filters may be arranged in a stacked manner as illustrated in FIG. 3, which is a block diagram representing another exemplary arrangement of components of a system in which the present invention may operate in
15 accordance with various aspects of the invention. In this embodiment, each of the filters 305-307 may process I/O requests and pass the requests (modified or unmodified) to another filter or other component in the stack. For example, in response to a read request received from one of the
20 applications 205, the I/O manager 215 may issue an I/O request and send this request to the filter 305. The filter 305 may examine the I/O request and determine that the filter 305 is

not interested in the I/O request and then pass the I/O request unchanged to the filter 306. The filter 306 may determine that the filter 306 will perform some action based on the I/O request and may then pass the I/O request (changed
5 or unchanged) to the filter 307. The filter 307 may determine that the filter 307 is not interested in the I/O request and pass the I/O request to the file system 235.

After the file system 235 services the I/O request, it passes the results to the filter 307. Typically, the results
10 pass in an order reverse from that in which the I/O request proceeded (e.g., first to filter 307, then to filter 306, and then to filter 305). Each of the filters 305-307 may examine the results, determine whether the filter is interested in the results, and may perform actions based thereon before passing
15 the results (changed or unchanged) on to another filter or component.

In another embodiment of the invention, filters may be arranged in a stacked/managed manner as illustrated in FIG. 4, which is a block diagram representing another exemplary
20 arrangement of components of a system in which the present invention may operate in accordance with various aspects of the invention. In this configuration, some filters are

associated with a filter manager while other filters are not. The filter manager 220 is placed in a stack with other filters (e.g., filters 305 and 307).

It will be readily recognized that filters may be
5 implemented in many other configurations without departing from the spirit or scope of the invention. In some embodiments, a filter comprises any object that examines I/O between an application and a file system and that is capable of changing, completing, or aborting the I/O or performing
10 other actions based thereon.

Returning to FIG. 2, the file system 235 may include one or more volumes that may be located locally or remotely to the machine or machines upon which the applications 205 execute.

Some filters monitor files in certain directories. For
15 example, a quota filter may enforce quotas on particular directories. As another example, a data screen filter may deny the creation of certain types of files (e.g., files that may be copyrighted as indicated by an extension such, .MP3, .JPG, and the like) in certain directories. These filters may
20 include a namespace that identifies the directories involved.

Certain file system requests (e.g., renames and deletes) may change the file system namespace for the objects monitored

by a filter. For example, a quota filter may be configured to enforce a quota of one gigabyte on an object called C:\DIR\QUOTA. The object's name and the quota applicable to the object may be stored in a metadata file which is persisted on non-volatile storage. An application may rename the object C:\DIR to C:\DIR2 or may move C:\DIR\QUOTA to C:\QUOTA. To continue enforcing the quota on the QUOTA object, the quota filter monitors renames and updates its metadata file each time a rename affects an object for which the quota filter is enforcing a quota.

In some operating systems, a rename operation may rename an object and/or move the object. Thus, by monitoring rename operations, a filter may capture either a name rename, a move of an object, or both. In operating systems in which renaming the file and moving the file are separate operations, a filter may need to monitor both of these operations to maintain namespace consistency with the namespace of a file system.

In addition, an application may delete an object, for example, C:\DIR or C:\DIR\QUOTA. To avoid monitoring objects that have been deleted, not enforce policies when an existing object is deleted and a new object with the same name is recreated in its place, and to reduce the size of the metadata

used in monitoring such objects, a filter may monitor deletes and update its metadata accordingly.

FIG. 5 is a block diagram representing a data structure that may be used to monitor objects that have policies applied to them in accordance with various aspects of the invention. As referred to herein, monitored objects comprise directories, files, anything that may be stored on a file system, and the like.

The data structure includes a node for each object of interest and is sometimes referred to as a prefix tree. The exemplary objects represented include the root directory "\", \DIR1, \DIR1\QUOTA1, \DIR1\QUOTA1\QUOTA2, \DIR2, \DIR2\DIR3, and \DIR2\DIR3\QUOTA3. The objects \DIR1\QUOTA1, \DIR1\QUOTA1\QUOTA2, and \DIR2\DIR3\QUOTA3 have policies P1, P2, and P3 applied to them, respectively. A policy may comprise a quota on disk space consumed by an object and its descendants (e.g., subdirectories), files allowed in an object and its descendants, what is allowed and/or not allowed with respect to an object and its descendants, and the like. When an object of interest is renamed or deleted on a file system, the data structure shown in FIG. 5 may need to be updated. Such updates may involve removing nodes, adding additional

nodes, and changing the contents of nodes as described in more detail below.

In general, an object is of interest if renaming or deleting the object would affect finding the object in the data structure or the application of a policy. For example, if \DIR2 and its subdirectories were deleted, the policy P3 would no longer need to be applied to \DIR2\DIR3\QUOTA3. Furthermore, the nodes DIR2, DIR3, and QUOTA3 would no longer be needed to apply the policy P3 and could be deleted to conserve space and decrease lookup time for any remaining policies.

As another example, if \DIR1\QUOTA1\QUOTA2 is renamed to \QUOTA5 on the file system this may affect finding the new object in the data structure. This is important because before a filter applies a policy to a particular operation, the filter needs to know whether the operation may affect any objects for which policies have been established. To find the new object in the data structure (e.g., so that a filter may determine if it needs to enforce a policy on a particular operation), a new node may be created and associated with the new object, the policies associated with the old node QUOTA2 (e.g., P2) may then be copied to the new node, and the old

node QUOTA2 may then be deleted. Policies P2 may then be enforced for the new object. Adding and deleting nodes in response to changes to objects are described in more detail in conjunction with FIGS. 8-10.

5 An application (e.g., a service or otherwise) may communicate which objects have policies associated with them to a filter. The application may provide a list of such objects or may add and/or delete policies as desired. After receiving such a list or instruction, the filter may create a
10 new data structure or modify an existing one to account for the objects of interest.

 The data structure may be reconstructed at boot time or instantiation of a filter that uses the data structure. This may be done using the metafile described above that is
15 persisted on non-volatile storage (e.g., a disk).

 The data structure may be shared between more than one filter. For example, two or more filters may enforce policies to objects that may be renamed or deleted. Instead of maintaining a separate copy of the data structure for each
20 filter, the two or more filters may share the data structure. One or more of the filters may be designated to keep the data structure in synch with the file system.

Furthermore, even though the filters may share the data structure, they may not wish to monitor all the same objects. To this end, additional data may be maintained that indicates which nodes each filter is interested in.

5 FIG. 6 is a flow diagram that generally represents actions that may occur to maintain consistency between a filter namespace and a file system namespace for objects of interest in accordance with various aspects of the invention. At block 605, the process begins. At block 610, the filter
10 receives an I/O operation of interest. In the case of managed filters, the filter may have previously registered to receive I/O operations of the type received. In the case of stacked filters, the filter may determine that the I/O is one that may affect the file system namespace. In one implementation, some
15 I/O operations that may affect the file system namespace include rename and delete.

At block 615, a determination is made as to whether the I/O operation relates to an object of interest. If so, processing branches to block 625; otherwise processing
20 branches to block 620 where the process returns. Referring to FIG. 5, an object of interest is an object represented by any node of the data structure. After receiving data associated

with an I/O operation of interest, the filter may then perform a lookup on the data structure to determine if the I/O operation involves any nodes therein. If not (e.g., the lookup fails), the filter may return.

5 At block 625, a determination is made as to whether the I/O operation is a rename operation. If so, processing branches to block 630; otherwise, processing branches to block 635. In another embodiment, the actions associated with blocks 615 and 610 are reversed. That is, first a
10 determination is made as to an I/O operation involves an object of interest, and then a determination is made as to whether the I/O operation is of interest. In this embodiment, If I/O does not involve an object of interest or if the I/O is not an I/O of interest, the process returns.

15 At block 630, the filter namespace is updated as appropriate as described in more detail in conjunction with FIGS. 7-10. At block 632, a stream context associated with the object is updated if the stream context exists as described in more detail in conjunction with FIG. 13. A
20 stream context comprises data into which information about an object may be stored and retrieved. A stream context may be created for an object, for example, when the object has been

marked for deletion. Even though an object may be marked for deletion, it may be renamed before it is deleted. To track the name of the object through renames, the new name of the object is stored in the stream context. Each time the object
5 is renamed, the stream context is updated to include the new name of the object. Then, when the object is finally deleted, the filter looks up the most recent name of the deleted object in the stream context and deletes the appropriate node or nodes from the filter's namespace.

10 At block 635, a determination is made as to whether the operation is a delete related operation. If so, processing branches to block 640; otherwise, processing branches to block 645. A delete related operation may comprise a create for delete on close operation, a set disposition operation, or a
15 cleanup operation as described in more detail in conjunction with FIG. 11. At block 640, the filter namespace is updated as appropriate as described in more detail in conjunction with FIG. 10.

At block 645, any other filter actions may be performed.
20 For example, a quota filter may enforce a quota by failing or allowing the I/O request to continue. In some implementations, the actions associated with block 645 occur

right after block 610. This may allow a filter to fail an operation and stop the changing of any object. At block 620, the process returns. The process described in conjunction with FIG. 6 may be repeated each time an I/O operation of
5 interest is received.

FIG. 7 is a flow diagram that generally represents actions which correspond to block 630 of FIG. 6 which may occur when a rename operation is received in accordance with various aspects of the invention. At block 705, the process
10 begins. At block 710, a determination is made as to whether the rename operation is in a pre-callback or post-callback phase. If the operation is in a pre-callback phase (i.e., before the I/O request has propagated to the file system), processing branches to block 715; otherwise, as the operation
15 is in a post-callback phase, processing branches to block 730. In essence, the process waits until the operation has succeeded before updating the data structure.

At block 715, a callback context is allocated and associated with the operation. The callback context may
20 include any data and provides a means for passing data regarding the operation to the post-callback phase. When the post-callback phase occurs for a particular operation, the

callback context associated with the operation may be accessed.

At block 720, the name of the object before the operation is performed (i.e., the old name) is stored in the callback
5 context. At block 725, a post-operation callback is forced which essentially means that the filter indicates that it wishes to be called after the operation has been performed on the file system and before the I/O results return to the application.

10 In the post-operation phase at block 730, a determination is made as to whether the operation failed. If so, processing branches to block 750; otherwise processing branches to block 735. If the operation failed, there is no need to update the data structure.

15 At block 735, the old name of the object is retrieved from the callback context. This old name is used to find a node in the data structure that maintains the filter namespace. At block 740, the file system is queried for the new name of the object which may include a new object name
20 and/or object path. At block 745, the data structure is updated as described in more detail in conjunction with FIG.

8. At block 750, the callback context is freed. At block 755, the process returns.

FIG. 8 is a flow diagram that generally represents actions which correspond to block 745 of FIG. 7 that may occur in updating the data structure in accordance with various aspects of the invention. At block 805, the process begins. At block 810, a determination is made as to whether the object was moved. If so, processing branches to block 815; otherwise, processing branches to block 830. If the object was moved, old nodes may be deleted and new nodes created to update the filter namespace and policies as indicated by blocks 815, 820, and 825. If only the name of the object was changed and not its path, the name of the object as indicated in the node may be changed to the new object name as indicated by block 830.

At block 835, descendant nodes of the moved node may be updated as appropriate. In one implementation, the complete path and object name are stored in each node. For example, referring to FIG. 5, the QUOTA2 node may store \DIR1\QUOTA1\QUOTA2. This may be done to speed execution or for other reasons. When an object name and/or path of a node are changed, each of the nodes that are descendants of the

node may be updated to include the new prefix path that includes the new object name and/or path associated with the changed node. In other implementations (e.g., where the complete path and object name are not stored in the node and
5 do not need to be updated in descendant nodes), block 835 may be skipped.

At block 840, the process returns.

FIG. 9 is a flow diagram generally representing actions which correspond to block 815 of FIG. 8 that may occur in
10 adding nodes when an object is moved in accordance with various aspects of the invention. At block 905, the process begins. At block 910, the first component of the new path of the object is selected. As used in a path, component refers to a portion of the path that corresponds to an ancestor
15 object of the object and may include the root directory. For example, the components of \DIR1\QUOTA1\QUOTA2 are root (i.e., "\"), DIR1, QUOTA, and QUOTA2. The first component is root.

At block 915, a determination is made as to whether a node exists that corresponds to the selected component. If
20 so, processing branches to block 925 as a new node does not need to be created; otherwise, processing branches to block

920. At block 920, a new node is created that corresponds to the component.

At block 925, a determination is made as to whether the currently selected component is the last component of the new path. If so, processing branches to block 905 where the process returns. Otherwise, processing branches to block 930 where the next component is selected and the actions associated with blocks 915-925 may be repeated until the last component is processed. It will be recognized that the process described in conjunction with FIG. 9, in essence, modifies the data structure to include nodes for any components in the new path that do not already exist in the data structure.

FIG. 10 is a flow diagram that generally represents actions which correspond to block 825 of FIG. 8 that may occur when updating the data structure to remove prefix nodes associated with the old object name in accordance with various aspects of the invention. In general, when a node is moved to a new location in the data structure, all nodes that are no longer necessary are removed from the data structure. A node is no longer necessary if it has no descendant that has a

policy associated with it and the node itself does not have a policy associated with it.

The process begins at block 1005. At block 1010, the node associated with the old object is selected. At block 5 1015, the policies of the old node are removed as they have already been copied to the new node previously. At block, 1020, a determination is made as to whether the currently selected node has a policy associated with it. If so, processing branches to block 1045; otherwise, processing 10 branches to block 1025.

At block 1025, a determination is made as to whether the node has any children. If so, processing branches to block 1045; otherwise, processing branches to block 1030. Because nodes are maintained for objects that have policies and their 15 ancestors and are not maintained otherwise, if node has any children, this indicates that at least one of the node's descendants is associated with a policy. For example, if a node that did not have a policy had two children and one of the children was renamed such that it was no longer a 20 descendant of the node, the node would not be deleted unless the second child was also removed.

At block 1030, the node is deleted as it is no longer needed. At block 1035, a determination is made as to whether the node has a parent. If so, processing branches to block 1040, where the parent is selected. From block 1040,
5 processing continues to block 1020 and the actions associated with blocks 1020-1035 may be repeated. It will be recognized, that these actions, in effect, recursively climb the data structure deleting nodes until they find a node that is either itself associated with a policy or which has a descendant
10 associated with a policy.

At block 1045, the process returns.

It will be recognized that the flow diagrams of FIGS. 6-10 represent one way in which a namespace data structure may be updated to be consistent when renames occur to a file
15 system, and that keeping the namespace data structure consistent with the namespace of the file system during renames may be implemented in a variety of ways without departing from the spirit or scope of the invention.

FIG. 11 is a flow diagram that generally represents
20 actions which correspond to block 640 of FIG. 6 that may occur to maintain consistency between a filter namespace and a file system namespace during delete related operations in

accordance with various aspects of the invention. At block 1105, the process begins.

At block 1110, a determination is made as to whether the operation is a create for delete on close operation. If so, 5 processing branches to block 1115; otherwise, processing branches to block 1130. A create operation may create an object or open an already-existing object. An object may be created for delete on close. When an object is created for delete on close, this means that the object should be deleted 10 after the object is closed. That is, the object may be created, exist, and continue to exist, until all processes that have handles to the object have closed their handles. After all the handles have been closed, the object is deleted. At block 1115, actions for the create are performed as 15 described in more detail in conjunction with FIG. 12.

At block 1130, a determination is made as to whether the operation is a set disposition operation. If so, processing branches to block 1135; otherwise, processing branches to block 1140. In one implementation, an object may have a flag 20 associated therewith that indicates whether the object should be deleted after closing all handles to the object. This flag may be set (i.e., indicate that the object should be deleted)

or cleared (i.e., indicate that the object should not be
deleted). Furthermore, the flag may be set or cleared by each
process that has an open handle to the object. Thus, one
process may set the flag only to have the flag cleared by
5 another process. The object remains or is deleted depending
on the disposition (e.g., set or cleared) of the flag at the
time that the last process with a handle to the object closes
the handle. At block 1135, actions related to setting the
disposition of the object are performed as described in more
10 detail in conjunction with FIG. 14.

At block 1140, a determination is made as to whether the
operation is a cleanup operation. If so, processing branches
to block 1145; otherwise, processing branches to block 1155.
A cleanup operation may be issued each time a process closes a
15 handle to an object. As mentioned previously, however, an
object may not actually be deleted until the last process with
a handle to the object closes the handle. Even after the last
process has closed the object, the object may not be deleted
if, for example, at the time of the last close, the flag for
20 deleting the object indicates that the object should not to be
deleted upon close. At block 1145, cleanup actions are

performed as described in more detail in conjunction with FIG.
15.

At block 1155 the process returns. The process described
in conjunction with FIG. 11 may be repeated each time a filter
5 receives a delete related I/O operation for an object of
interest.

FIG. 12 is a flow diagram that generally represents
actions which correspond to block 1115 of FIG. 11 which may
occur when a create for delete on close operation is received
10 in accordance with various aspects of the invention. At block
1205, the process begins.

At block 1210, a determination is made as to whether the
operation is in the pre-operation phase. If so, processing
branches to block 1215; otherwise, processing branches to
15 block 1220.

At block 1215, a post-operation callback is forced so
that the filter will receive a callback during the post-
operation phase of the operation. At block 1220, a
determination is made as to whether the operation failed. If
20 so, processing branches to block 1240; otherwise, processing
branches to block 1225.

At block 1225, a stream context is allocated (if it does not already exist) and associated with the object. At block 1230, the object name is stored in the stream context. At block 1235, a delete flag is updated in the stream context.

5 For objects created for delete on close, the flag may indicate that the object is to be deleted when all handles to the object are closed. This type of flag may not be modified by subsequent processes. When all handles to the object are closed, an object that has been opened for delete on close is
10 deleted.

For objects that have not been opened for delete on close, a flag may be set or cleared (via a set disposition operation) to indicate that the object should or should not be deleted on close. The flag may then be changed by subsequent
15 set disposition operations. Whether the object is deleted after all handles are closed depends on the state of the flag when the last handle to the object is closed.

At block 1240, the process returns.

FIG. 13 is a flow diagram that generally represents
20 actions which correspond to block 632 of FIG. 6 which may occur when a rename operation is received in accordance with

various aspects of the invention. At block 1305, the process begins.

At block 1310, a determination is made as to whether the operation is in the pre-operation phase. If so, processing
5 branches to block 1330; otherwise, processing branches to block 1315.

At block 1330, a post-operation callback is forced so that the filter will receive a callback during the post-operation phase of the operation. At block 1315, a
10 determination is made as to whether the operation failed. If so, processing branches to block 1335; otherwise, processing branches to block 1320.

At block 1320, a determination is made as to whether a stream context exists that is associated with the object. If
15 so, processing branches to block 1325; otherwise, processing branches to block 1335. A stream context may not exist if, for example, the object was not opened for delete on close and no dispositions have been set or cleared on the flag that indicates whether the object should be deleted on close.

20 At block 1325, the name is updated in the stream context to reflect the name to which the object has been changed.

This is useful, for example, later in determining whether the name change will affect any node.

At block 1335, the process returns.

FIG. 14 is a flow diagram that generally represents
5 actions which correspond to block 1135 of FIG. 11 which may occur when a set disposition operation is received in accordance with various aspects of the invention. At block 1405, the process begins.

At block 1410, a determination is made as to whether the
10 operation is in the pre-operation phase. If so, processing branches to block 1415; otherwise, processing branches to block 1420.

At block 1415, a post-operation callback is forced so that the filter will receive a callback during the post-
15 operation phase of the operation. At block 1420, a determination is made as to whether the operation failed. If so, processing branches to block 1440; otherwise, processing branches to block 1425.

At block 1425, a determination is made as to whether a
20 stream context exists that is associated with the object. If so, processing branches to block 1435; otherwise, processing branches to block 1430 where a stream context is created and

the name of the object stored therein. The stream context is associated with the object.

At block 1435, the delete flag is updated to reflect the delete disposition of the object. As mentioned before, a
5 process may set the flag to indicate that the object should be deleted upon close or clear the flag to indicate that the director should not be deleted on close. At block 1440, the process returns.

FIG. 15 is a flow diagram that generally represents
10 actions which correspond to block 1145 of FIG. 11 which may occur when a cleanup operation is received in accordance with various aspects of the invention. At block 1505, the process begins.

At block 1510, a determination is made as to whether the
15 operation is in the pre-operation phase. If so, processing branches to block 1515; otherwise, processing branches to block 1520.

At block 1515, a post-operation callback is forced so that the filter will receive a callback during the post-
20 operation phase of the operation. At block 1520, a determination is made as to whether the operation failed. If

so, processing branches to block 1550; otherwise, processing branches to block 1525.

At block 1525, a determination is made as to whether a stream context exists that is associated with the object. If so, processing branches to block 1530; otherwise, processing branches to block 1550. At block 1530, a determination is made as to whether the object is an object of interest to the filter. If so, processing branches to block 1535; otherwise, processing branches to block 1550.

At block 1535, a determination is made as to whether the delete flag is set to delete or if the object was opened for delete on close. If either of these conditions is met, processing branches to block 1540; otherwise, processing branches to block 1550. In one implementation, when an object is opened for delete on close, it may not be later indicated that the object will not be deleted on close.

At block 1540, a determination is made as to whether the object is deleted. As many processes may open an object and the object is only deleted after all the processes have closed handles to the object, an object may not be deleted even though all the other criteria have been met. In one embodiment of the invention, an operating system may inform

the filter when an object has been deleted (instead of the filter checking for itself). In such an embodiment, the method described in conjunction with FIG. 15 may comprise determining if the object is of interest and deleting nodes as
5 appropriate.

At block 1545, the node and any related nodes are deleted as appropriate. If the node is an ancestor of other nodes in the data structure described in conjunction with FIG. 5, the other nodes are also deleted. Furthermore, actions similar to
10 those described in conjunction with FIG. 10 may be performed to delete nodes that are ancestor nodes of the node and that are no longer needed. At block 1550, the process returns.

It will be recognized that the actions described in conjunction with FIGS. 11-15 may be replaced with other
15 actions and may be performed in a different order depending on the way an operating system operates without departing from the spirit or scope of the invention. For example, some operating systems may not have a flag to indicate whether the file should be deleted after all handles to the file are
20 closed. Rather, such operating systems may provide a separate delete operation to delete an object. In such operating systems, a method for detecting changes for deletes may

comprise monitoring when an object is deleted and deleting nodes (if any) that are associated with the object.

The force of post-operation callbacks of FIGS. 12-14 may be skipped in implementations in which a post-operation
5 callback occurs by default.

As can be seen from the foregoing detailed description, there is provided a method and system for maintaining namespace consistency with a file system. While the invention is susceptible to various modifications and alternative
10 constructions, certain illustrated embodiments thereof are shown in the drawings and have been described above in detail. It should be understood, however, that there is no intention to limit the invention to the specific forms disclosed, but on the contrary, the intention is to cover all modifications,
15 alternative constructions, and equivalents falling within the spirit and scope of the invention.

WHAT IS CLAIMED IS:

1. A computer-readable medium having computer-executable instructions, comprising:
receiving a request to access an object of a file system;
5 determining that the object is within a namespace associated with a filter; and
updating the namespace associated with the filter based on a change to the object.

10 2. The computer-readable medium of claim 1, wherein the request to access an object comprises a rename operation.

3. The computer-readable medium of claim 2, wherein the rename operation causes the object to move to a different
15 location in the file system.

4. The computer-readable medium of claim 2, wherein the rename operation causes a name associated with the object to change.

20

5. The computer-readable medium of claim 1, wherein the request to access an object comprises a delete operation.

6. The computer-readable medium of claim 5, wherein the delete operation comprises setting a flag to indicate that the object is to be deleted after all handles associated with the object are closed.

7. The computer-readable medium of claim 6, further comprising clearing the flag to indicate that the object is not to be deleted.

8. The computer-readable medium of claim 1, wherein the request to access an object comprises a cleanup operation.

9. The computer-readable medium of claim 8, wherein a cleanup operation is issued each time a process closes a handle to the object.

10. The computer-readable medium of claim 1, wherein determining that the object is within a namespace associated with the filter comprises searching for a name of the object in a data structure maintained by the filter.

11. The computer-readable medium of claim 10, wherein the data structure is arranged hierarchically.

12. The computer-readable medium of claim 10, wherein
5 the data structure includes a node for each object that has a policy associated therewith.

13. The computer-readable medium of claim 12, wherein the data structure further includes a node for each ancestor
10 of each object that has a policy associated therewith.

14. The computer-readable medium of claim 13, wherein an ancestor comprises a directory in which the object resides.

15 15. The computer-readable medium of claim 1, wherein the object comprises a directory of the file system.

16. The computer-readable medium of claim 1, wherein the object comprises a file of the file system.

20

17. The computer-readable medium of claim 1, wherein the change to the object comprises a change to a name of the object.

5 18. The computer-readable medium of claim 1, wherein the change to the object comprises a change to a location of the object in the file system.

19. In a computing environment, a method, comprising:
opening a handle to an object of a file system;
determining whether the object is related to a filter
namespace; and

5 if the object is related to the filter namespace,
updating the filter namespace based on a change to the object.

20. The method of claim 19, wherein opening a handle to
an object of the file system is part of a create.

10 21. The method of claim 20, wherein the create comprises
an open operation.

22. The method of claim 20, wherein the create comprises
15 an open for delete on close operation, wherein the object is
deleted after all handles to the object have been closed.

23. The method of claim 20, wherein the create causes
the object to be created on the file system.

20

24. The method of claim 19, wherein the filter comprises a process that has an opportunity to examine data related to a request that propagates to the file system before the request has propagated to the file system.

5

25. The method of claim 19, wherein the filter comprises a process that has an opportunity to examine result data related to a request that propagates to the file system before the result data returns to a requestor that issued the request.

10

26. The method of claim 19, wherein the filter registers to receive data for certain operations related to a volume.

15

27. The method of claim 19, wherein the filter namespace includes data that identifies an object for which a policy exists.

28. The method of claim 27, wherein the filter namespace further includes data that identifies each ancestor of the object.

20

29. The method of claim 28, wherein the data is stored
in nodes that are arranged hierarchically.

30. The method of claim 19, wherein the filter comprises
5 a kernel-mode process.

31. The method of claim 19, wherein the change to the
object comprises deleting the object.

10 32. The method of claim 19, further comprising
determining whether the change is successful.

33. The method of claim 19, further comprising changing
a file system name of the object.

15

34. The method of claim 19, further comprising changing
a delete disposition of the object.

35. An apparatus for synchronizing a namespace,
comprising:

a file system arranged to store an object, the file
system having a first namespace;

5 a process arranged to request access to the object; and

a filter arranged to examine selected communication
between the file system and the process and to update a second
namespace therefrom for any change to the first namespace that
affects the second namespace.

10

36. The apparatus of claim 35, wherein a change to the
first namespace affects the second namespace if the object is
included in the second namespace and the object is deleted or
renamed.

15

37. The apparatus of claim 36, wherein renaming the
object comprises changing a name of the object in the first
namespace.

20

38. The apparatus of claim 36, wherein renaming the
object comprises moving the object to another location within
the first namespace.

**Smart & Biggar
Ottawa, Canada
Patent Agents**





FIG. 1

2/15

FIG. 2

3/15


FIG. 3

FIG. 4

4/15






FIG. 5

5/15



6/15

FIG. 6

7/15



FIG. 7

FIG. 8

9/15





FIG. 9

FIG. 10


11/15

FIG. 11

12/15

FIG. 12

13/15

FIG. 13

14/15


FIG. 14

FIG. 15