

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
15 November 2007 (15.11.2007)

PCT

(10) International Publication Number
WO 2007/130192 A1

- (51) International Patent Classification:
G06F 15/16 (2006.01) G06F 12/16 (2006.01)
- (21) International Application Number:
PCT/US2007/004183
- (22) International Filing Date:
15 February 2007 (15.02.2007)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/405,236 17 April 2006 (17.04.2006) US
- (71) Applicant (for all designated States except US): MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) Inventors: MICHAEL, Michael L.; One Microsoft Way, Redmond, WA 98052-6399 (US). SCHEIDEL, William L.; One Microsoft Way, Redmond, WA 98052-6399 (US). LUBER, Paul B.; One Microsoft Way, Redmond, WA 98052-6399 (US). OLTEAN, P. Adrian; One Microsoft Way, Redmond, WA 98052-6399 (US). KALACH, Ran; One Microsoft Way, Redmond, WA 98052-6399 (US).

AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

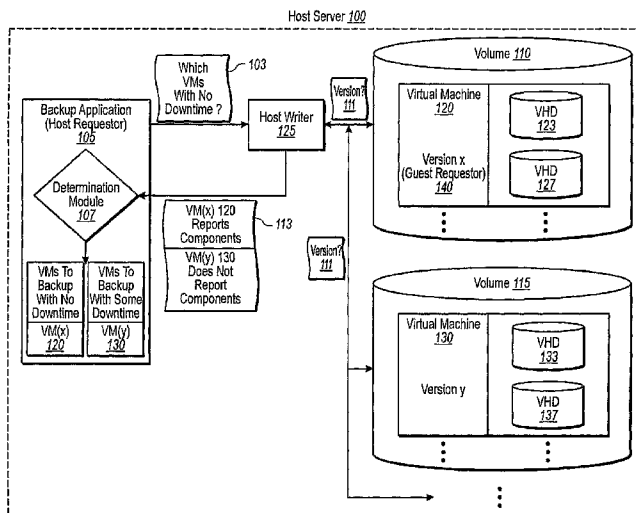
Published:

- with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

(54) Title: CREATING HOST-LEVEL APPLICATION-CONSISTENT BACKUPS OF VIRTUAL MACHINES



(57) Abstract: A host server hosting one or more virtual machines can back up host volumes and the one or more virtual machines installed thereon in an application-consistent manner. In one implementation, a host-level requestor instructs a host-level writer to identify which virtual machines qualify for application-consistent backups. The host-level requestor then instructs the host-level writer to initiate virtual machine backups through guest-level requestors in each appropriately-configured virtual machine, wherein the virtual machines create application-consistent backups within the virtual machine volumes. The host-level requestor then initiates snapshots of the server volumes on the host-level. The virtual machine-level snapshots can thus be retrieved from within the host-level snapshots of the server volumes.

WO 2007/130192 A1

CREATING HOST-LEVEL APPLICATION-CONSISTENT
BACKUPS OF VIRTUAL MACHINES

5

BACKGROUND

1. Background and Relevant Art

[0001] As computerized systems have increased in popularity, so have the needs to store and back up electronic files and other communications created by the users and applications associated therewith. In general, computer systems and related devices create files for a variety of reasons, such as in the general case of creating a word processing document in a work setting, as well as creating a file used for more sophisticated database purposes. In addition, many of these documents can include valuable work product, or sensitive information that should be protected. One will appreciate, therefore, that there are a variety of reasons why an organization will want to backup electronic files on a regular basis, and thereby create a reliable restoration of an originally created file when needed.

[0002] As computerized systems are continually developed into more complex software and hardware configurations, the complications associated with backing up these configurations has also increased. In its simplest form, backing up a computer involves simply copying information from a computer disk or volume, and passing that information to a backup location. Simply copying data from a disk, however, becomes much more complex when backing up large sets of data on multiple servers on a network, particularly when the sets of data are continually changing during the backup process. For example, along with simply backing up data, some concerns for large organizations can include a need to store data in a quick, reliable, and relatively uninterrupted fashion so that the data are also quickly recoverable. Additional concerns include how to access various server

data in the first instance, particularly as each server may have different functions that may make the server more or less accessible to a given backup program than another server.

[0003] These types of complications in backing up server data can be particularly difficult in the virtualized server environment. Generally, a virtual server environment involves the use of a host server, on which one or more “virtual machines” can be installed. These virtual machines, comprise unique instances of an operating system that are installed on an allocated portion of the host server disk space and resources (e.g., host memory), and thus comprise a distinct machine boundary. Thus, each virtual machine could be represented by a unique network identifier, and thus appear separate and distinct on a network from the host server itself. In addition, each virtual machine can also have additional storage partitions made on the host space allocated for the virtual machine. Furthermore, in addition to appearing to be a distinct machine from the host server, the virtual machine can also appear to be hosting multiple drives or volumes, despite existing on the physical disk(s) of a host server. Accordingly, one can appreciate that virtual machines can be used in a number of important ways to more efficiently distribute hardware resources in an organization.

[0004] Despite these advantages, it is not a simple matter to simply copy the host volumes on which the virtual machines are installed and have that data be as useful (e.g., consistent and valid) as necessary. One of the reasons for this difficulty stems from the nature of the virtual machines themselves, whereby they are deemed to have distinct identities (e.g., including distinct operating systems) with respect to other virtual machines, and with respect to the corresponding host server(s). For example, the host server cannot simply instruct applications inside the virtual machine to create an “application-consistent” snapshot of their data, since such applications are generally under the control of the virtual machine. Thus, when a host creates a snapshot or a backup copy

of the volume (or volumes) on which the virtual machine is installed, processes inside the virtual machine may continue to run, and as such, a copy of the virtual machine could likely include data and file state that are valid for different instances of time. That is, the copy of the virtual machine data would not be "application-consistent." If you were to
5 restore such a virtual machine, it may not even function properly, if at all.

[0005] Generally, an "application-consistent" backup means that applications take part in the backup preparation processes, and produce application files to be backed up based on what is on the physical disk, as well as what is in memory. As such, these files or snapshots to be backed up are consistent with the current state of the application, and are
10 recoverable by the application at a later point. By contrast, "crash-consistent" data refer to data that are backed up generally without the benefit of involving the applications used to create the data in the first instance. As a result, the primary production during a crash-consistent backup is a set of data that are the same as what are on the physical disk during the backup process without consideration for what is in memory, and without a
15 consideration for application state. This is similar to how file data might exist during a total system crash, such as during a loss of power or reboot, and thus has no guarantees of being, for example, application-consistent. In particular, for certain complex operating environments, particularly where it is important to ensure as seamless a transition as possible after a backup, a consistent (i.e., application and/or file system)-consistent backup
20 approach is preferred.

[0006] There are a variety of ways for creating a consistent backup of volume data. Unfortunately, conventional systems that attempt to create application-consistent backups of virtual machines (and corresponding volumes managed thereby) generally cannot effectively create an application-consistent backup without first pausing or turning off the
25 virtual machine. This typically results in some unwanted downtime, which, in some cases

with certain “mission-critical” software could create particular difficulties for an administrator. One reason for this is so that the volume of interest can be copied while no additional writes are being made to the volume, such that the data generated through the application are all consistent for the same instance of time. Of course, in an environment
5 where ready, continuous and efficient data access is important, pausing or turning off a virtual machine to make a backup copy is less desirable, and can create a drain on an organization’s resources.

[0007] Other ways that an organization might attempt to create application-consistent backups of a virtual machine is by installing a specific backup agent within each instance
10 of a given virtual machine. Generally, a backup agent will be configured to interact with the application writers within the operating system to create an application-consistent backup, just as might normally be done on the host level for the host server’s main volume. Unfortunately, it is not a simple matter to just install new backup agents for each virtual machine. For example, an organization will typically need to purchase a new
15 separate backup agent (or license) for each virtual machine, and then install each backup agent within the virtual machine. One will appreciate that this can present a fairly significant difficulty in terms of cost and resource expenditures (e.g., including management overhead) to large organizations that might be running tens, hundreds, or even thousands of virtual machines.

20 [0008] Accordingly, there are a number of difficulties associated with backing up virtual machines that can be addressed.

BRIEF SUMMARY

[0009] Implementations of the present invention provide systems, methods, and computer program products configured to create consistent host-level backups of one or more virtual machines. In particular, implementations of the present invention provide the ability to backup a host server and corresponding one or more virtual machines using existing backup requestors and writers in a consistent fashion, without necessarily incurring significant interruptions to the operation of the one or more virtual machines. In one implementation, for example, a host server writer (e.g., a virtual server VSS writer) instructs each virtual machine to create one or more application-consistent snapshots of its own virtual machine-level volume data. The backup application on the host server also creates snapshots of the host-level volumes on which one or more virtual machines (e.g., a Virtual Hard Disk file thereof) are installed. The host server can then retrieve the previously created virtual machine-level snapshots from within the host-level snapshots, and finish backup processes.

[0010] For example, a method from the perspective of a host server of creating a consistent backup of virtual machine volume data without necessarily requiring stopping or rebooting of the one or more virtual machines, can involve identifying at least one virtual machine having one or more components for writer-involved backup processes. In addition, the method can involve sending an instruction to each of the one or more virtual machines to prepare a virtual machine-level, application-consistent snapshot. The method can also involve identifying that snapshot operations at the at least one virtual machine have completed. Furthermore, the method can involve creating one or more host-level snapshots of one or more host volumes on which the identified one or more virtual machines are installed. The method can still further involve retrieving the application-

consistent snapshot of the virtual machine volumes made by the at least one virtual machine.

[0011] By contrast, a method from the perspective of a virtual machine of creating a consistent backup of one or more virtual machine volumes in response to instructions from a host writer of the host server can involve receiving a request from a host writer to identify available software components. The method can also involve receiving a request from the host writer to take a snapshot of one or more volumes hosted by the virtual machine with at least one of the available software components. In addition, the method can involve sending instructions to one or more application writers on the virtual machine to prepare corresponding one or more virtual machine applications for backup. Furthermore, the method can involve sending a signal to the host writer that snapshot preparations for each of the one or more volumes hosted inside the virtual machine have completed.

[0012] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

[0013] Additional features and advantages of the invention will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by the practice of the invention. The features and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other features of the present invention will become more fully apparent from the following description and appended claims, or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] In order to describe the manner in which the above-recited and other advantages and features of the invention can be obtained, a more particular description of the invention briefly described above will be rendered by reference to specific
5 embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

10 [0015] Figure 1A illustrates an overview schematic diagram in accordance with an implementation of the present invention in which a host-level backup application identifies which one or more virtual machines can be backed up in a consistent manner;

[0016] Figure 1B illustrates the components of Figure 1A, in which the host server initiates backup processes in conjunction with a virtual machine-level guest requestor in a
15 virtual machine, in accordance with an implementation of the present invention;

[0017] Figure 1C illustrates the components of Figures 1A-1B, in which the host server requestor creates snapshots of the host volumes on which are installed the one or more virtual machines, such that the host-level snapshots also contain the virtual machine volume data, which in turn contain the virtual-machine level snapshots performed earlier
20 by the virtual machines;

[0018] Figure 1D illustrates an overview schematic diagram in accordance with an implementation of the present invention in which the backup application illustrated in Figures 1A-1C retrieves host-level snapshot data having virtual machine-level snapshot data contained therein; and

[0019] Figure 2 illustrates flowcharts of methods comprising a series of acts in accordance with implementations of the present invention from the perspectives of a host server and of a virtual machine for providing application-consistent backups of virtual machines installed on one or more host server volumes.

DETAILED DESCRIPTION

[0020] Implementations of the present invention extend to systems, methods, and computer program products configured to create consistent host-level backups of one or more virtual machines. In particular, implementations of the present invention provide the ability to backup a host server and corresponding one or more virtual machines using existing backup requestors and writers in a consistent fashion, without necessarily incurring significant interruptions to the operation of the one or more virtual machines. In one implementation, for example, a host server writer (e.g., a virtual server VSS writer) instructs each virtual machine to create one or more application-consistent snapshots of its own virtual machine-level volume data. The backup application on the host server also creates snapshots of the host-level volumes on which one or more virtual machines (e.g., a Virtual Hard Disk file thereof) are installed. The host server can then retrieve the previously created virtual machine-level snapshots from within the host-level snapshots, and finish backup processes.

[0021] One will appreciate after reading this specification and claims that a host server (e.g., through a virtual server writer) can also enable the creation of nevertheless consistent backups of virtual machines even if they cannot be readily identified as being configured for consistent (e.g., application and/or file system)-consistent backup processes. For example, the virtual machines may be turned off or otherwise not running, or the virtual machines may not be installed with the appropriate software or hardware components. Nevertheless, the host server can use different components to copy the host volumes and corresponding virtual machines in a manner that preserves state, and then pauses or stops the virtual machines at least momentarily. The host server writer can then allow the virtual machines to resume after the snapshot of the host volume(s) has/have been made.

[0022] One will appreciate after reading this specification and claims, that implementations of the present invention can provide for consistent backups of virtual machines in a way that minimizes downtime, and without requiring the purchase and installation of new backup agents. As such, organizations that implement virtual machine solutions can provide better host server, virtual machine, and data availability in a network system in a manner that minimizes an organization's resource expenditures.

[0023] Accordingly, Figure 1A illustrates an exemplary host server 100 configured to manage volumes 110 and 115. Installed on volumes 110 and 115 are virtual machines 120 and 130, respectively. By way of explanation, although each virtual machine 120, 130 can be thought of as a separate computerized system on one level, each virtual machine 120, 130 can also be thought of as a collection of "files" (e.g., a Virtual Machine Configuration, or "VMC," file, and one or more virtual hard disk – "VHD" – files) when viewed from the level of host 100. In any event, although Figure 1A illustrates one virtual machine per volume, this is not required, and there may be several virtual machines installed on any given host server 100 volume. Similarly, each virtual machine could span multiple volumes on one or more host servers.

[0024] Furthermore, each virtual machine can manage additional drives, which are effectively additional volume allocations within the volume allocation(s) on which the given virtual machine is installed. For example, Figure 1A shows that virtual machine 120 also manages one or more virtual physical disks, which are nevertheless part of host volume 100. Generally, each virtual physical disk can also be represented by a file, such as a Virtual Hard Disk file (i.e., VHD 123, 127, etc.). In particular, the VHD file is exposed as a physical disk inside the virtual machine that can further comprise additional one or more volumes (not shown) contained therein, where each volume has its own Global Unique Identifier ("GUID"). Thus, VHD file 123 (which in this illustration

contains only one volume for simplicity) might thus have a volume that appears as drive “m:\” – or some other unique identifier appropriate for a given operating system – while VHD 127 (which also in this case includes only one volume) might have a volume that appears on the network as drive “n:\” and so on.

5 [0025] In addition, Figure 1A shows that host 100 comprises backup application 105, which can also be referred to as a “host requestor,” or “host-level requestor.” Generally, backup application 105 comprises a series of computer-executable instructions configured to direct backup events at host 100. In one implementation, such as in the MICROSOFT operating environment for example, backup application 105 comprises a Volume Shadow
10 Copy Service (“VSS”) Requestor. Figure 1A also shows that backup application 105, in turn, can be configured to direct its backup instructions through host writer 125, which can also be referred to as a “host writer,” “host-level writer,” or “host-level virtual server VSS writer.”

[0026] Generally, host writer 125 comprises a series of computer-executable
15 instructions configured to implement backup instructions received from backup application 105. In one implementation, such as in the MICROSOFT environment for example, host writer 125 can comprise a VSS Writer, such as the VSS writer used with MICROSOFT VIRTUAL SERVER. By way of explanation, reference herein to MICROSOFT components is exemplary only. In particular, one will appreciate after
20 reading this specification and claims that the components, modules, systems, and functions described herein can be applied to a wide variety of components, modules, and functions used in other operating environments.

[0027] In order to perform consistent backups of the virtual machines, host 100 will generally need to determine which of the hosted virtual machines can report appropriate
25 components configured to make consistent backups possible in the first place. As a

preliminary matter, continued reference is made herein to “application-consistent” backup or snapshot processes. One will appreciate, however, that application-consistent backup processes are only one example of “consistent” backup operations in accordance with implementations of the present invention. Other examples of consistent backup processes
5 include file system and/or crash-consistent backup processes.

[0028] In any event, and with reference to application-consistent backups, for example, some virtual machines might be installed with the appropriate writers and requestors for making application-consistent backups, while other virtual machines might be installed without the appropriate writers and requestors. For those virtual machines
10 running without such appropriate components, host 100 can still backup those virtual machines, but might not necessarily do so in the same manner as those with the appropriate components. In particular, host 100 can be configured to backup those virtual machines that report the appropriate components with no downtime or interruption (or virtually none), and alternatively configured to backup those virtual machines that do not
15 report (i.e., missing or not in operation) the appropriate components with at least some downtime or interruption.

[0029] Accordingly, Figure 1A shows that backup application 105 (i.e., “host requestor”) begins backup services at least in part by sending request 103 to host writer 125. In this case, request 103 instructs host writer 125 to identify which virtual machines
20 can be backed up with “no downtime.” For example, request 103 instructs host writer 125 to ascertain which of virtual machines 120, 130, etc. comprise the appropriate guest writers and/or requestors for performing an internal application-consistent backup. In one implementation in the MICROSOFT environment, for example, such components can include those compatible with “VM Additions” as used within a MICROSOFT VIRTUAL
25 SERVER infrastructure.

[0030] These and other similarly configured virtual machine components are configured to interact from within a virtual machine with (and respond to instructions from), for example, a host-level virtual server VSS writer. Figure 1A also shows that, upon receiving request 103, host writer 125 can communicate via message 111 (which can
5 be, for example, a private Application Program Interface – “private API”) with virtual machine 120, and identifies that virtual machine 120 reports components pursuant to “version x.” For example, Figure 1A shows that virtual machine 120 includes “Guest Requestor 140.”

[0031] Figure 1A also shows that host writer 125 further communicates (e.g., via
10 message 111) with virtual machine 130, and that host writer 125 identifies that virtual machine 130 reports components pursuant to “version y.” In this particular case, “version y” signifies that virtual machine 130 does not have the appropriate components for consistent backup processes. In response to its communications with the virtual machines (e.g., 120, 130, etc.), host writer 125 can send one or more messages of its gathered
15 responses back to backup application 105. For example, writer sends message 113, which indicates that virtual machine 120 is of version “x,” and reports the appropriate components, and further indicates that virtual machine 130 is of version “y,” but does not report the appropriate components.

[0032] Upon receipt, Figure 1A shows that backup application 105 can take message
20 113, and make its own determinations about which virtual machines to backup, and which virtual machines not to backup. For example, Figure 1A shows that determination module 107 parses the information of message 113, places virtual machine 120 into the “VMs to Backup with No Downtime” category, and places virtual machine 130 into the “VMs to Backup with Some Downtime” (or VMs not to backup) category. In alternative
25 implementations, host writer 125 simply makes its own determinations about what should

or should not be backed up (or backed up with some downtime), and then reports such categorizations back to backup application 105. In any event, the illustrated categorizations do not necessarily mean that virtual machine 130 will not be backed up, per se. In most cases, this simply means that virtual machine 120 cannot be backed up in an appropriately consistent fashion, and virtual machine 130 might only be configured to be backed up a consistent fashion with some downtime. After categorizing each virtual machine, backup application 105 begins implementing backup processes.

[0033] As shown in Figure 1B, for example, backup application 105 sends instructions 117a to host writer 125. Instructions 117a, in turn, tell host writer 125 to begin application-consistent backup processes at least with respect to virtual machine 120. Host writer 125 then prepares its own communication 117b, which tells a virtual machine to begin “writer-involved” backup processes. Host writer 125 then sends message 117b to each virtual machine (e.g., 120) indicated by backup application 105, and each indicated virtual machine in turn receives message 117b through its corresponding in-guest requestor.

[0034] For example, Figure 1B shows that guest requestor 140 of virtual machine 120 receives message 117b. In one implementation, host writer 125 can communicate (e.g., message 111, 117b, etc.) with guest requestor 140 through one or more private APIs, through Remote Procedure Calls (“RPCs”), or the like, though this is not required in all implementations. In other implementations, such as if host writer 125 and guest requestor 140 are located at separate network locations (or even different domains), host writer 125 can communicate with other appropriate communication interfaces and or components. In particular, one will appreciate that implementations of the present invention can also be practiced even when needing to backup a virtual machine from a remote location over a Wide Area Network.

[0035] In any event, and once triggered, guest requestor 140 thus begins backup processes of its volumes in accordance with backup application's 105 initial instructions. As shown in Figure 1B, for example, guest requestor 140 sends a virtual machine-internal message 117c to applications installed on any volumes managed by virtual machine 120. Message 117c can include instructions for each application writer (e.g., a VSS writer, not shown) to prepare for application-consistent backup processes of each volume managed within the volumes contained in its physical disks (i.e., represented by VHD files 123, 127). Generally, each application writer can comprises computer-executable instructions that can be included in applications and services within the virtual machine, and which help provide application-consistent backups of the application data.

[0036] For applications that are running when receiving request 117c, the application writer might respond, for example, by preparing its data stores, and ensuring that no writes occur on the volume (e.g., 123, 127) while the snapshot is being created. To make data on the disk consistent, the application writer might also flush its buffers to disk, or write out in-memory data to disk. In addition, the application writer might provide information about the application name, icons, files to include or exclude, and a strategy to restore the files. For applications that are not running, the corresponding application writer may not respond to message 117c, and guest requestor 140 might thus assume that all data handled by the application writers on the volume are consistent, the databases are closed, and no additional effort is required to perform the backup.

[0037] Generally, an application writer in accordance with one or more implementations of the present invention can be associated with one or more components. Each component, in turn, can comprise a group of files (e.g., a database and a set of log files) that are to be copied up as a whole. Thus, each application writer will need to provide data of each component and each corresponding file of the component to an

appropriate backup service (e.g., a virtual machine requestor – e.g., guest requestor 140). Each application writer can further provide information about restoring the data on a component-by-component (and hence component-file-by-component-file) basis. In one implementation, for example, host writer 125 can provide a list of files used to keep persistent information about the managed virtual machines (e.g., 120, 130, etc.) Host writer 125 might report, for example, for each virtual machine a path of a particular virtual machine configuration file, any paths of a virtual hard disk file, etc.

[0038] In any event, and in response to instructions 117c, Figure 1B shows that the corresponding application writers within volumes 123 and 127 produce an application-consistent state of their volume data therein. The volume data reported by the relevant applications for the volume within the corresponding physical disk file (e.g., VHDs 123, 127) are encompassed in volume snapshots. For example, Figure 1B shows that volume data 145 in VHD 123 are copied as “clean copy of data” 155, and data 150 in VHD 127 are copied as “clean copy of data” 160. Generally, these “clean” copies of data 155, 160 are consistent copies (e.g., volume shadow copy) of data within the guest volume snapshot, and generally do not constitute separate copies from the snapshot itself. In particular, “clean copies” 155, 160 are essentially application-consistent copies of volume data within the snapshots contained in the Virtual Hard Disk files (VHDs) 123 and 127 (Figure 1A), respectively.

[0039] Once each appropriately configured virtual machine (e.g., 120) has made its own internal, application-consistent (or “clean”) copy of its physical disk volumes (e.g., copy 155 of data 145), host writer 125 can allow backup application 105 to proceed making snapshots of the host volumes (e.g., 110) on which the appropriately configured virtual machines (e.g., 120) are installed. The virtual machine volume data inside of these

host-level snapshots, however, are not necessarily application-consistent (i.e. “dirty” or “crash consistent”).

[0040] Generally, “dirty” copies are those that cannot be guaranteed as application-consistent with respect to virtual machine volume data, at least in part since they do not
5 invoke application writer-involved snapshot processes within the virtual machines. For example, host 100 can use host-level application writers (e.g., host writer 125) to backup volume 110 data with writer-involved processes, but, as previously mentioned, these host-level application writers will only copy the entire virtual machine files as they are viewed by the host. As a result, even if using host-level application writers to create application-
10 consistent copies of host-level volume data, each underlying virtual machine could be undergoing various data changes at the time of making the host-level backup of volume 110.

[0041] For example, Figure 1C shows that host writer 125 sends message 157 to backup application 105, wherein message 157 indicates that backup of the virtual
15 machines (e.g., 120) is complete. As such, backup application 105 can proceed making host-level snapshots of host volume 110 (and also volume 115, etc., if appropriate). In response, Figure 1C shows in this example that backup application 105 initiates (e.g., signal 135) backup processes at volumes 110, 115, etc. In one implementation, for example, host requestor 105 initiates writer-involved backup processes and produces an
20 application-consistent snapshot of the data on host-level volumes 110, 115 (i.e., application-consistent apart from virtual machine files) through a host-level VSS component (not shown).

[0042] For example, Figure 1C shows that host 100 creates a snapshot 167, 173 of host volumes 110, 115. Snapshots 167, 173 also include copies of the virtual machine
25 VHD files contained therein (i.e., VHD file 123 inside volume 110, etc.). Since, during

the host-level snapshot, the virtual machines 120, 130 did not have their corresponding writers involved in the snapshot processes, the newer, host-level snapshots include VHD volume data (e.g., 156, 161) for the virtual machine-level volumes that are not application-consistent. Nevertheless, since the relevant virtual machine has already made an application-consistent copy of its volumes, VHD data 156 within host-level snapshot 167 includes the previously created (i.e., virtual machine-created) application-consistent snapshot 155. Similarly, VHD data 161 within host-level snapshot 167 includes the previously created (i.e., virtual machine-created) application-consistent snapshot 160. As mentioned above, the VHD data 156, 161 within host-level snapshot 167 (also 173) of VHD files 123, 127 are generally crash-consistent as a whole.

[0043] Upon creating host-level snapshots 167, 173 of the host-level volumes, host 100 will need to retrieve the virtual-machine level snapshot from within the host-level volume snapshots 167, 173. In one implementation, this involves reverting the host level VHD snapshot data 156, 161 back to the corresponding virtual machine-level snapshot data 155, 160. To do so, host writer 125 can first mount these host-level and virtual machine-level snapshots within the file system using any number of appropriate components to expose the underlying data. For example, backup application 105 can configure host writer 125 to make host snapshots 167, 173, etc. (and underlying VHD files snapshots) writable (rather than read-only). Furthermore, host writer 125 can obtain those device IDs for the host-level snapshots, and use those to identify the different "Device IDs" corresponding to the encapsulated VHD files (e.g., 123, 127).

[0044] Upon exposing the virtual machine VHD files from within the writable host-level snapshots 167, 173, host server 100 can then mount each device ID corresponding to a VHD file, so that each VHD file appears as a physical disk at the host level. In particular, host writer 125 can use the device IDs for VHD files 123 and 127 within the

host snapshot in conjunction with an image mounter component to identify each volume within VHD files 123 and 127. For example, Figure 1C shows that host server, (e.g., via host writer 125) sends instructions 175 to mount device IDs, including as arguments the device ID for VHD file 123 and VHD file 127. In one implementation with particular
5 respect to the MICROSOFT operating environment for example, host 100 can use a “VHDImageMounter” component to receive and execute mounting instructions 175.

[0045] Mounting VHD files 123, 127 from within the host snapshot 167 in turn can expose the operating systems within those VHD files, which in turn allows the virtual-machine-level snapshots 155, 160 to be exposed on host server 100. In particular,
10 mounting VHD files 123, 127 from within the host snapshot 167 of volume 110 in this case exposes the virtual machine-level snapshots 155, 160 inside the different VHD data 156, 161 of host snapshot 167. Thus, Figure 1C also shows that host 100 mounts VHD file 123 inside host snapshot 167 as drive “x:\,” which exposes VHD data 156, and further exposes the earlier-created virtual machine-level snapshot data 155 therein; and further
15 mounts VHD file 127 from within host snapshot 167 as drive “y:\,” where host snapshot 167 also includes VHD data 161, which has the earlier-created virtual machine-level snapshot data 160 exposed therein.

[0046] Now that backup application 105 has mounted VHD files 123, 127 from within the host snapshot 167, backup application 105 can now revert the VHD data 156, 161 of
20 the host-level snapshot 167 back to the VHD data 155, 160 of the virtual machine-level volume snapshots. As shown in Figure 1D, for example, backup application 105 uses component interface 185 (e.g., a VSS component interface) to revert the dirty and clean snapshot data within drives “x” and “y.” In one implementation, backup application 105 uses interface 185 to simply delete or overwrite VHD data 156, 161, and replace it with
25 corresponding clean VHD data 155, 160.

[0047] Once making this replacement, backup application 105 has not only corrected or adjusted (i.e., reverted, replaced) the snapshot data of VHD files 123, 127, but also changed the overall host-level snapshot 167 so that its underlying virtual machine-level volume data are now fully consistent. That is, all virtual-machine data within the host-level snapshots can now also be considered application-consistent. Furthermore, control of the virtual machine-level snapshots 155, 160 has been placed into the control of host 100 backup processes. Thus, when host 100 makes its own application-consistent backup of its volume snapshots, this host-level backup can now include an application-consistent backup of each appropriately-configured virtual machine.

10 [0048] The above-identified components and processes are described primarily for the case in which the identified virtual machines (e.g., 120) provide a response that indicates configuration with the appropriate software components for creating a consistent backup. Specifically, such virtual machines (e.g., 120) can be backed up with little or no downtime. As previously mentioned, however, there may be alternative instances in which a given virtual machine provides no response since it is not running, or provides a response that indicates that it is not configured to make consistent backups. For example, virtual machine 130 might be running an operating system that does not include certain needed software components, such as the appropriate application writers (e.g., correct version of a shadow copy service), an appropriate guest requestor (e.g., 140) or the like.

15 In such alternative instances, implementations of the present invention can still provide at least crash-consistent backups of the virtual machines (e.g., 130) that may or may not require at least some subsequent downtime if later needing to ensure the data are consistent.

[0049] For example, host writer 125 can call a communication interface (not shown) and instruct the interface to save all of the state of virtual machine 130 (and/or other

25

similarly-limited virtual machines). In one implementation with particular respect to the MICROSOFT operating environment, for example, host writer 125 can call a Virtual Server Communication interface (i.e., "VS COM API," not shown). The communication interface can then report to host writer 125 when it has finished saving the virtual machine
5 130 state. In particular, the communication interface can report this to host writer 125, which in turn allows backup application 105 to create a snapshot of the corresponding host volume on which virtual machine 130 resides (i.e., volume 115).

[0050] In addition to saving virtual machine 130 state, the communication interface can also momentarily pause or stop virtual machine operations using any appropriate
10 virtual and/or private APIs. If the communication interface identifies, for example, any changes to the virtual machine 130 state, the communication interface can report this information. Such a report can, in turn, fail the backup operations, and cause restart at a more appropriate time. In particular, the backup application 105 can continue to repeat these steps until backup application 105 has finished creating snapshots of volume 115 in
15 a state of desired consistency.

[0051] Once backup operations are complete, host 100 has thus created an application-consistent copy of volume 110 and its corresponding virtual machines (e.g., 120), and has further created an consistent (i.e., at least crash-consistent) copy of volume 115 and its corresponding virtual machines (e.g., 130). One will appreciate that this copy of volume
20 115 can not only be crash-consistent, but also be "VHD-consistent" with respect to the state of the VHD file. In particular, saving virtual machine state can provide a valid and consistent VHD file that can be successfully recovered at a later point, as desired.

[0052] In still another implementation, components in accordance with the present invention can further include use of a hardware provider to provide many of the consistent
25 snapshot functions described previously. When using a hardware provider, for example a

“VSS Hardware Provider” (e.g., a plex/mirror provider), volume 110, 115 storage can reside on a “disk array” that is being “mirrored” in a separate location. When it is time to create the host-level snapshot (i.e., non-application-consistent for the virtual machine), the hardware provider can simply “break” the mirror between the volume and its
5 corresponding mirror. This can expose the host-level snapshot as a separate, standalone disk device. Backup application 105 can then call any number of Virtual Disk Service (“VDS”) and/or Volume Shadow Copy Server (“VSS”) APIs to make the mirror of the volume writable, manipulate the mirror of the volume, and then expose and mount the underlying clean (i.e., application-consistent) snapshot data 155, 160. This type of
10 hardware solution could also be used with remotely connected storage volumes, such as in a Storage Area Network (“SAN”) environment to create application-consistent backups of virtual machines distributed over several storage mediums.

[0053] Along these lines, implementations of the present invention can further include use of, for example, guest requestor 140 to create a hardware-transportable, application-
15 consistent snapshot inside a particular virtual machine (e.g., 120). For example, a component inside, for example, virtual machine 120 configured for snapshots of SAN volumes (e.g., an appropriately configured VSS component) can create a writer-involved snapshot (not shown) of a remote volume attached to the guest (e.g., virtual machine 120, 130), for which there is thus no host server 100 VHD file. Guest requestor 140 can then
20 import the snapshot and surface the snapshot as a separate Logical Unit Number (“LUN”) on the host server. Host writer 125 can then ask backup application 105 (i.e., a host requestor) to backup the LUN as part of backing up virtual machine 120. Backup application 105 can thus backup the entire LUN (i.e. the guest or virtual machine-level snapshot), such that reversion with, for example, a host-level snapshot of a virtual machine
25 volume, would not be necessary.

[0054] Accordingly, Figures 1A-1D and the corresponding text provide a number of systems, components, and interfaces for providing efficient host-level backups of virtual machines in a consistent manner. In addition to the foregoing, implementations of the present invention can also be described in terms of methods comprising one or more acts
5 for accomplishing a particular result. For example, Figure 2 illustrates flow charts from the perspective of host 100 and virtual machine 120 for creating host-level, application-consistent backups of virtual machines. The methods and corresponding acts of Figure 2 are described below with reference to the systems and components of Figures 1A-1D.

[0055] For example, Figure 2 shows that a method from the perspective of host 100 of
10 creating an application-consistent backup of each of one or more virtual machines without necessarily requiring stopping or pausing of the one or more virtual machines, comprises an act 210 of identifying one or more virtual machines on a host. Act 210 includes identifying at least one virtual machine having one or more components for writer-involved backup processes. For example, backup application 105 (i.e., "host requestor"
15 105) asks (via message 103) host writer 125 what virtual machines are present on host 100, and which of these virtual machines can be backed up in an application-consistent manner (i.e., with little or no downtime).

[0056] Host writer 125, e.g., a VSS writer, sends corresponding version request 111 to each virtual machine (e.g., 120, 130), and identifies what software components exist to aid
20 the backup processes (e.g., guest requestor 140), such as components that can be used for creating application-consistent backups. Host writer 125 then sends corresponding response 113 to backup application 105. Backup application 105 then determines, e.g., via determination module 107, which of the identified virtual machines can be backed up using application-consistent methods (i.e., "VMs to Backup," which include the
25 appropriate components), and those to be backed up using other methods (i.e., "VMs to

Backup with Some Downtime”) which do not include the appropriate components, or are turned off).

[0057] Accordingly, Figure 2 further shows that a method from the perspective of virtual machine 120 of creating a consistent backup of one or more virtual machine volumes comprises an act 210 of receiving a request to identify components. Act 210 includes receiving a request from a host writer to identify available software components. For example, virtual machine 120 (e.g., via guest requestor 140) receives version request 111, which includes a request to identify if virtual machine 120 includes a guest requestor and one or more application writers (and/or appropriate versions thereof), as these would be appropriate components for creating an application-consistent backup. Virtual machine 120 then returns a corresponding response, which host writer 125 formats as message 113 and sends to backup application 105.

[0058] In addition, Figure 2 shows that the method from the perspective of host 100 comprises an act 220 of sending a backup request to the set of virtual machines. Act 220 includes sending an instruction to each of the one or more virtual machines to prepare one or more machine-level snapshots of one or more virtual machine volumes. As shown in Figure 1B, for example, backup application 105 sends instructions 117a to backup virtual machine 120 to host writer 125, where virtual machine 120 is part of a set of virtual machines that has appropriate components for creating application-consistent snapshots with little or no downtime. Host writer 125, in turn, sends instructions 117b to initiate writer-involved backup preparations to guest requestor 140 at virtual machine 120.

[0059] Accordingly, Figure 2 shows that the method from the perspective of virtual machine 120 also comprises an act 230 of receiving a snapshot request from a host writer. Act 230 includes receiving a request from the host writer to create a snapshot of one or more volumes hosted by the virtual machine with at least one of the available software

components. For example, guest requestor 140 receives message 117b from host writer 125, where message 117b includes a request to begin writer-involved snapshot processes. As such, virtual machine will begin preparing a snapshot of its volumes managed by each VHD file (e.g., 123, 127, etc.).

5 [0060] In addition, Figure 2 shows that the method from the perspective of virtual machine 120 comprises an act 240 of sending backup instructions to one or more application writers. Act 240 includes sending instructions to one or more application writers on the virtual machine to prepare corresponding one or more virtual machine applications for a snapshot. For example, Figure 1B shows that guest requestor 140 sends
10 message 117c to the application writers (not shown) on volumes within VHD files 123, 127. In one implementation, this can involve sending request to writers corresponding to such applications as database, e-mail, or other similar applications to prepare for a snapshot of the relevant volume data.

[0061] Furthermore, Figure 2 shows that the method from the perspective of virtual
15 machine 120 comprises an act 250 of sending a signal indicating completion of the virtual machine snapshots. Act 250 includes sending a signal to the host writer that snapshot operations for each of the one or more volumes hosted within the virtual machine have completed. For example, upon completion of virtual machine snapshot activities, guest requestor sends a completion signal (not shown) to host writer 125. In response to this
20 signal, host writer 125 can then send message 157 to backup application 105, telling backup application 105 to finish host-level snapshot operations.

[0062] Accordingly, Figure 2 also shows that the method from the perspective of host 100 comprises an act 260 of identifying completion of the virtual machine snapshots. Act 260 includes identifying that snapshot operations at the at least one virtual machine have

completed. For example, as previously mentioned, host writer 125 sends message 157, which indicates that virtual-machine level snapshots have completed.

[0063] In addition, Figure 2 shows that the method from the perspective of host 100 comprises an act 270 of creating one or more host-level snapshots. Act 270 includes
5 creating one or more host-level snapshots of one or more host volumes on which the at least one virtual machine is installed. For example, backup application 105 sends instructions (e.g., 135) to continue backup processes at host volumes 110, 115. As shown in Figure 1B, this can result in a host snapshot 167 that includes non-application-consistent VHD file 123, 127 data 156, 161. As discussed herein, the different VHD data
10 of the host-level snapshot 167, while not necessarily being application-consistent for the VHD files 123, 127, nevertheless include the previously created application-consistent VHD snapshot data 155, 160. In particular, the host-level snapshot 167 of volume 110 could be application-consistent in all other respects except for VHD files 123, 127, etc.

[0064] In another implementation, such as with hardware providers, host-level
15 snapshots can be made in the case of mirrored host-level volumes by breaking the mirror between the host volume(s) on which a virtual machine is installed and the corresponding mirror volume. The mirror volume thus comprises a snapshot of the host volume. This snapshot also comprises the previously taken virtual machine-level snapshots of its one or more volumes, where the virtual machine-level snapshots are application-consistent. The
20 virtual-machine level snapshot can thus be retrieved from the host's mirror volume.

[0065] Furthermore, Figure 2 shows that the method from the perspective of host 100 comprises an act 280 of retrieving the virtual machine-level snapshots encapsulated inside the VHD files of the host snapshot. Act 280 includes retrieving the one or more virtual machine-level snapshots of the one or more virtual machine volumes from the one or more
25 host-level snapshots of the one or more host volumes. For example, as shown in Figure

1D, host 100 can replace (e.g., via a VSS component) host-level VHD snapshot data 156 within VHD 123 with virtual machine-level VHD snapshot data 155. Similarly, host 100 can replace host-level VHD data 161 within VHD 127 with virtual machine-level VHD snapshot data 160. Alternatively, in the case of a hardware provider (i.e., breaking a mirror between a host volume and a mirror volume), the virtual machine-level snapshot can be exposed and retrieved from the mounted mirror volume, as previously described. In either case, the host-level backup data (e.g., 195) can thus be based at least in part on virtual machine-level snapshot data.

[0066] Accordingly, Figures 1A-2, and the corresponding text, provide a number of systems, components, and mechanisms that can be used to create appropriately consistent backup copies of virtual machines on one or more host volumes. As described herein, components in accordance with the present invention can be used to create consistent backups of virtual machines that do not require downtime of the virtual machine, nor require the purchase of additional backup agents that need to be separately installed and managed. Furthermore, components as described herein can take advantage of conventional shadow copy services to generate a consistent snapshot of a virtual machine, and then retrieve the virtual machine snapshot from within the host-level snapshot.

[0067] The embodiments of the present invention may comprise a special purpose or general-purpose computer including various computer hardware, as discussed in greater detail below. Embodiments within the scope of the present invention also include computer-readable media for carrying or having computer-executable instructions or data structures stored thereon. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer.

[0068] By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk

storage or other magnetic storage devices, or any other medium which can be used to carry or store desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or a combination of hardwired or wireless) to a computer, the computer properly views the connection as a computer-readable medium. Thus, any such connection is properly termed a computer-readable medium. Combinations of the above should also be included within the scope of computer-readable media.

10 [0069] Computer-executable instructions comprise, for example, instructions and data which cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, 15 the specific features and acts described above are disclosed as example forms of implementing the claims.

[0070] The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be 20 considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

CLAIMS

We claim:

1. At a host server in a computerized environment, the host server having one or more virtual machines installed on one or more host volumes, the virtual machines configured to manage one or more virtual machine volumes on the one or more host volumes, a method of creating an application-consistent backup of virtual machine volume data without necessarily requiring stopping or rebooting of the one or more virtual machines, comprising the acts of:

10 identifying at least one virtual machine having one or more components for writer-involved backup processes;

sending an instruction to the at least one virtual machine to prepare corresponding one or more virtual machine-level snapshots of one or more virtual machine volumes;

15 identifying that snapshot operations at the at least one virtual machine have completed;

creating one or more host-level snapshots of one or more host volumes on which the at least one virtual machine is installed; and

20 retrieving the one or more virtual machine-level snapshots of the one or more virtual machine volumes from the one or more host-level snapshots of the one or more host volumes.

2. The method as recited in claim 1, wherein the virtual machine-level snapshots of the one or more virtual machine volumes are application-consistent.

3. The method as recited in claim 2, wherein the host level snapshot includes one or more different snapshots of the one or more virtual machine volumes, the method further comprising an act of replacing data of the one or more different snapshots within the one or more host-level snapshots with data of the one or more virtual machine-level snapshots
5 of one or more virtual machine volumes.

4. The method as recited in claim 1, further comprising an act of converting the one or more host-level snapshots from read-only data to writable data.

10 5. The method as recited in claim 1, wherein a host requestor relays the backup instruction through a host writer to a guest requestor in each of the one or more virtual machines.

6. The method as recited in claim 1, wherein the one or more virtual machine-level
15 volumes are managed within one or more virtual disk files, such that each different snapshot comprises a snapshot of a virtual disk file having a virtual machine-level snapshot contained therein.

7. The method as recited in claim 6, further comprising an act of mounting a device
20 identifier corresponding to each of the one or more different snapshots, wherein each virtual hard disk file of the at least one virtual machine is exposed as a physical disk on the host.

8. The method as recited in claim 1, further comprising an act of identifying one or more different virtual machines for which the one or more components for writer-involved backup processes cannot be identified.

5 9. The method as recited in claim 8, further comprising an act of identifying that the one or more different virtual machines have been one of saved or turned off.

10 10. The method as recited in claim 8, wherein the act of identifying one or more different virtual machines for which the one or more components cannot be identified further comprises an act of identifying that the one or more different virtual machines do not include at least one of an appropriately configured guest requestor and an appropriately configured virtual machine-level application writer.

15 11. The method as recited in claim 8, further comprising an act of saving state for each of the one or more different virtual machines.

12. The method as recited in claim 11, further comprising an act of at least momentarily pausing operations of the one or more different virtual machines.

20 13. The method as recited in claim 12, further comprising an act of, upon identifying that the state has been saved, creating a snapshot of each host volume on which the one or more different virtual machines are installed.

14. The method as recited in claim 13, further comprising a step for instructing each of the one or more different virtual machines to resume operations.

15. The method as recited in claim 1, wherein the one or more host volumes on which
5 the at least one virtual machine is installed are implemented in one or more disks hosted by one or more external storage components, the one or more external storage components being configured to create one or more snapshots of the one or more external disks using at least one of:

- (i) a hardware-based snapshot creation method;
- 10 (ii) a snapshot creation method exposed by one or more storage arrays.

16. The method as recited in claim 15, wherein the act of creating one or more host-level snapshots of one or more host volumes allows the one or more host-level snapshots to be managed independently on the host server, wherein the one or more host-level
15 snapshots can be transformed in separate, read-write volumes, and wherein the one or more virtual machine-level snapshots can be retrieved from within the one or more host-level snapshots.

17. At a virtual machine in a computerized environment, wherein the virtual machine is installed on a host server volume, a method of creating an application-consistent backup of one or more virtual machine volumes in response to instructions from a host writer of the host server, comprising the acts of:

5 receiving a request from a host writer to identify available software components;

 receiving a request from the host writer to create a snapshot of one or more volumes hosted by the virtual machine with at least one of the available software components;

10 sending instructions to one or more application writers on the virtual machine to prepare corresponding one or more virtual machine applications for the snapshot; and

 sending a signal to the host writer that snapshot operations for each of the one or more volumes hosted by the virtual machine have completed.

15

18. The method as recited in claim 17, wherein the at least one available software component is a guest requestor installed with the virtual machine, wherein the guest requestor receives the request from the host writer, and sends the instructions to the one or more application writers.

20

19. The method as recited in claim 17, wherein at least one of the one or more of the virtual machine volumes resides on a network storage medium, the method further comprising the acts of:

importing a snapshot of the at least one volume from the network storage
5 medium to the host server over a network; and

surfacing the snapshot at the host server as a logical unit number, wherein
the host server can backup the snapshot corresponding to the logical unit number
as consistent data.

20. At a host server in a computerized environment, the host server having one or more virtual machines installed thereon, a computer program product having computer-executable instructions stored thereon that, when executed, cause one or more processors at the host server to perform a method of creating an application-consistent backup of each
5 of the one or more virtual machines without necessarily requiring stopping or rebooting of the one or more virtual machines, comprising the acts of:


identifying at least one virtual machine having one or more components for writer-involved backup processes;

10 sending an instruction to the at least one virtual machine to prepare corresponding one or more virtual machine-level snapshots of one or more virtual machine volumes;

identifying that snapshot operations at the at least one virtual machine have completed;

15 creating one or more host-level snapshots of one or more host volumes on which the at least one virtual machine is installed; and

retrieving the one or more virtual machine-level snapshots of the one or more virtual machine volumes from the one or more host-level snapshots of the one or more host volumes.






Fig. 1B




Fig. 1C

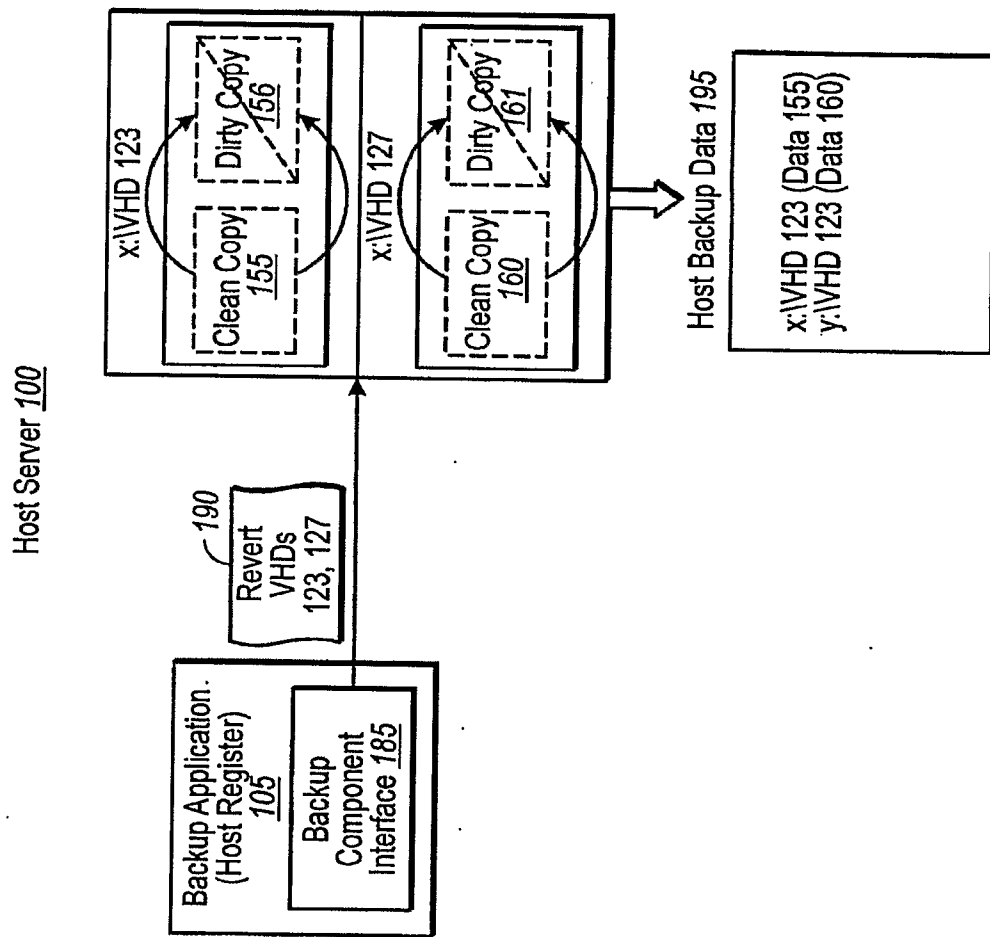


Fig. 1D

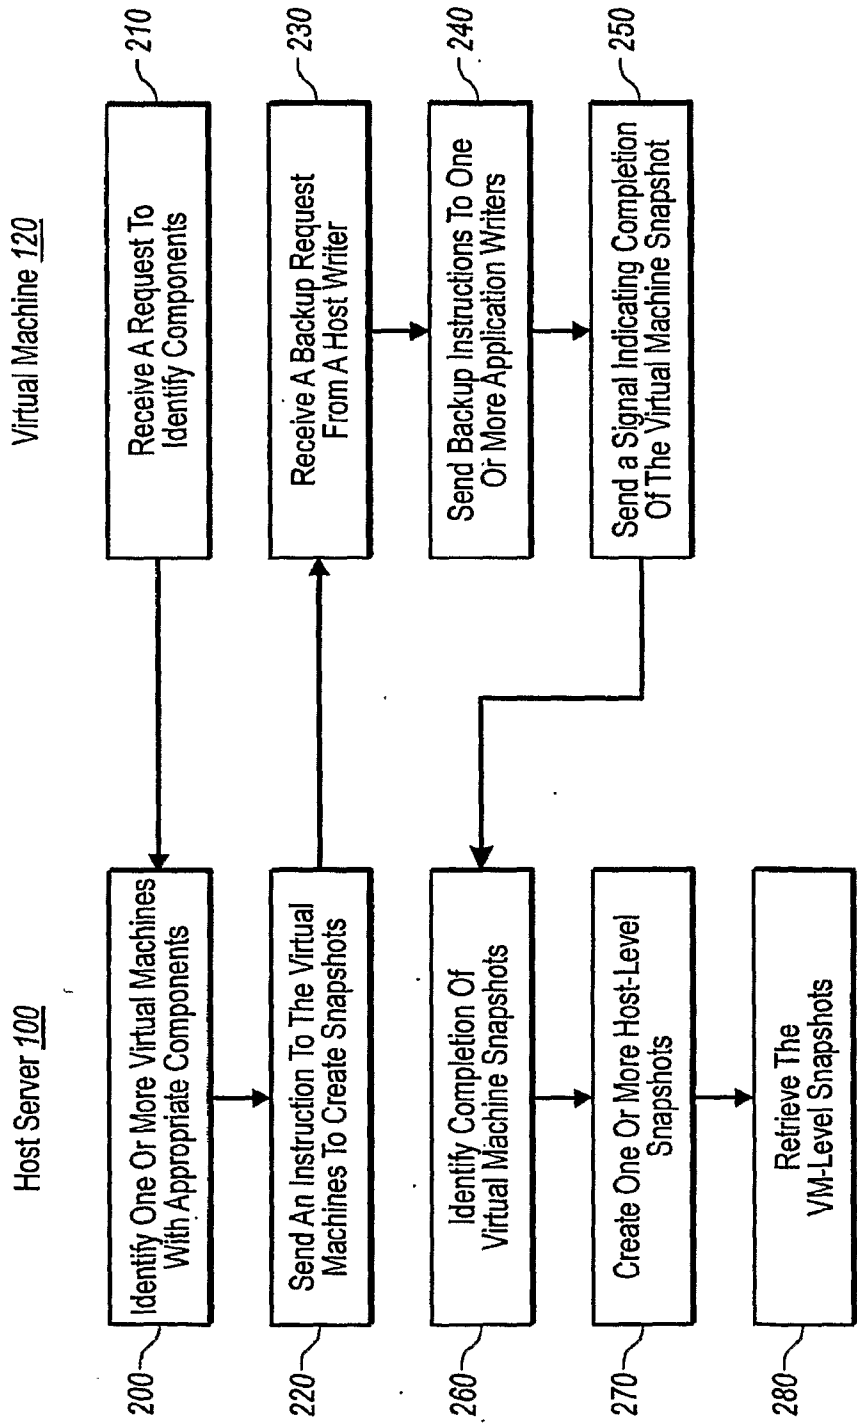


Fig. 2

A. CLASSIFICATION OF SUBJECT MATTER*G06F 15/16(2006.01)i, G06F 12/16(2006.01)i*

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 8 : G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean Utility models and applications for Utility models since 1975

Japanese Utility models and applications for Utility models since 1975

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKIPASS(KIPO internal) "host", "virtual", "machine", "backup", "snapshot"

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 20030131207 A1(HIROSHI ARAKAWA et al.) 10 July 2003 See abstract; figures 1-3, 12-24, 27.	1-20
A	US 20050033930 A1(YUMIKO HARUMA et al.) 10 February 2005 See abstract; figures 6-12.	1-20
A	US 20040083345 A1(YOUNG HO KIM et al.) 29 April 2004 See abstract; figures 5-7.	1-20
A	US 20030204597 A1(HIROSHI ARAKAWA et al.) 30 October 2003 See abstract; figures 2, 5, 9, 10.	1-20

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"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

"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

"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

"&" document member of the same patent family

Date of the actual completion of the international search

20 AUGUST 2007 (20.08.2007)

Date of mailing of the international search report

21 AUGUST 2007 (21.08.2007)

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
920 Dunsan-dong, Seo-gu, Daejeon 302-701,
Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

HAN, Seon Kyoung

Telephone No. 82-42-481-8523



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2007/004183

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 20030131207 A1	10.07.2003	JP 2003202964 A2 US 6883073 B2	18.07.2003 19.04.2005-----
US 20050033930 A1	10.02.2005	JP 2005055972 A2 US 2006236051 A1	03.03.2005 19.10.2006
US 20040083345 A1	29.04.2004	KR 2004036788 A US 07085899 B2 US 2006206677 A1	03.05.2004 01.08.2006 14.09.2006
US 20030204597 A1	30.10.2003	None	