(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
16 June 2011 (16.06.2011)

PCT

(10) International Publication Number
WO 2011/069705 A1

(51) International Patent Classification:
G06F 3/06 (2006.01)

(21) International Application Number:
PCT/EP2010/064959

(22) International Filing Date:
7 October 2010 (07.10.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09178621.0 10 December 2009 (10.12.2009) EP

(71) Applicant *(for all designated States except US)*: **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **REICHERT, Christoph** [DE/DE]; IBM Deutschland Research & Development GmbH, Hechtsheimer Str. 2, 55131 Mainz (DE). **NOLL, Dietmar** [DE/DE]; IBM Deutschland Research & Development GmbH, Hechtsheimer Str. 2, 55131 Mainz (DE).

(74) Agent: **KLEIN, Hans-Joerg**; IBM Deutschland Management & Business Support GmbH, IBM-Allee 1, 71139 Ehningen (DE).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, 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, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, 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, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

(54) Title: DATA STORAGE SYSTEM AND METHOD

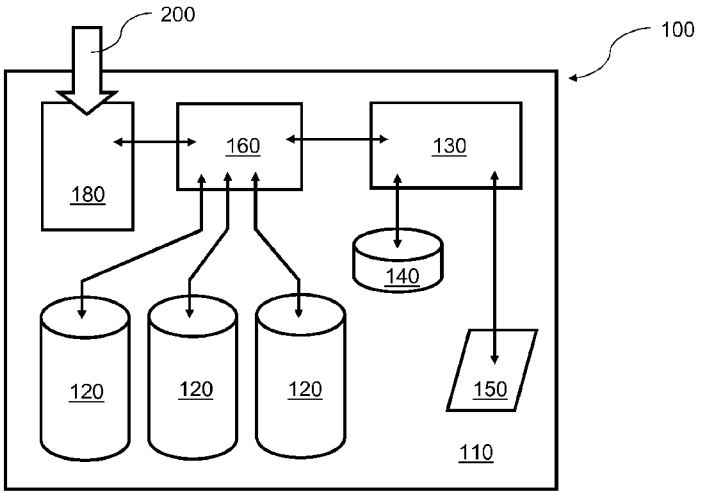

Fig. 2

(57) Abstract: The invention relates to a method of storing one or more volumes of data in a data storage system (100) including one or more data storage containers (120) provided by data storage hardware (110). The one or more data storage containers (120) are thin-provisioned to provide virtual data storage capacity which is greater than a real data storage capacity of the data storage hardware (60). The method includes: (a) determining a real data storage capacity of the data storage system (100) for accommodating the one or more new volumes; (b) determining over-allocation information relating to the one or more data storage containers (120); (c) determining extrapolated future anticipated use of the one or more containers (120) of the data storage system (100) from historical data storage use information; (d) selecting one or more candidate data storage containers (120) on the basis of information from steps (a) to (c). The method enables the data storage system (100) to be used more efficiently for storing the one or more new volumes.

WO 2011/069705 A1

D E S C R I P T I O N

DATA STORAGE SYSTEM AND METHOD

FIELD OF THE INVENTION

The invention relates to data storage systems, for example to data storage systems involving storage volume selection in thin-provisioned data storage sub-systems. Moreover, the present invention relates to methods of storing data in systems, for example to methods of selecting containers for creation of one or more data storage volumes in data storage volume in thin-provisioned data storage sub-systems. Furthermore, the present invention relates to software products recorded on machine-readable data media, wherein said software products are executable on computing hardware for implementing these aforementioned methods.

BACKGROUND OF THE INVENTION

Referring to Fig. 1, a contemporary computing system 10 has a need for considerable data storage capacity, for example for storing large data files such as video files, sound files, signal data samples, transaction files, text database files and so forth. Moreover, it is known in the computer system 10 that several software applications 20 can be executing concurrently and have a requirement to access data memory 30, namely data storage. Conceptually, in respect of a given software application 20, it is convenient to consider its requirement for storing data to be likened to a data "container" 40. However, it has also been appreciated that concurrently executing software applications 20 do not need to use their associated data containers 40 to fullest extent simultaneously. Thus, by suitable data storage management executed by data storage management 50, it is possible for a given data storage capacity

in the data memory 30 to increase the number of software applications 20 that can be supported by permitting a given real data storage hardware 60 to be shared between the data containers 40 of different software applications 20. In other words, the data storage capacity is referred to as being "thin provisioned".

Beneficially, the data storage management 50 can utilize indications communicated from the software applications 20 regarding their dynamic data storage requirements and/or monitor operation if the containers 40 and/or the real data storage hardware 60 to anticipate where data memory can be reused to serve a plurality of software applications 20.

US-A-2008/0091748 discloses a capacity monitoring method utilizable in a computer system including one or more application computers, one or more storage systems, and a management computer. The storage system includes a physical disk and a disk controller. In the capacity monitoring method, a storage area of the physical disk belongs to a storage pool, the storage system provides a volume, and the management computer monitors a used capacity of the storage pool, judges whether or not a storage capacity required for operating the application computer for a predetermined time period is present in the storage pool based upon an increasing speed of the used capacity of the storage pool to be monitored, and executes a predetermined process operation when the storage capacity required for operating the application computer for the predetermined time period is not present in the storage pool. The capacity monitoring method utilizes a volume capacity history data table representing a history related to a used capacity of a thin partitioning "volume", namely data storage area, and also used capacity of applications programs.

Thus, thin provisioning (TP) is concerned with an allocation of more data storage for data storage consumers, for example software applications, than is actually in practice available in a storage sub-system, for example in the real data storage hardware 60, thereby providing an ability in practice to over-allocate the storage sub-system. Contemporary implementations of thin provisioning do however have a requirement to configure a data storage container to accommodate thin-provisioned volumes, which often also involve defining maximum values for real as well as virtual storage capacity. One major challenge for storage administrators, for example the aforementioned data storage management 50, introduced as a consequence of thin partitioning is an additional level of complexity to cope with real as well as well as virtual capacities; in other words, it is necessary not only to cope with managing free and used capacity, but also to manage virtual free and virtual used data storage capacity as well as a degree of over-allocation which is utilized. In contemporary systems, it is a responsibility of a storage administrator, for example the data storage management 50, to identify best suited data storage containers, for example included within a data storage pool provided by a data storage subsystem, and also a particular storage subsystem out of a candidate list of storage subsystems to provide data storage capacity for a new volume, taking into account an existing configuration of candidate containers. Thus, appropriate container placement within a data storage system is a very complex and tedious task.

SUMMARY OF THE INVENTION

It is an object of the invention to provide a data storage system which is operable to identify optimal placement of one or more volumes of data within thin-partitioned data storage for

fulfilling one or more corresponding storage capacity provisional requests.

It is a further object of the invention to provide a method in a data storage system of identifying optimal placement of one or more volumes of data within thin-partitioned data storage for fulfilling one or more corresponding storage capacity provisional requests.

These objectives are achieved by the features of the independent claims. The other claims, the drawings and the specification disclose advantageous embodiments of the invention.

A method of storing one or more volumes of data in a data storage system including one or more data storage containers provided by data storage hardware, said one or more data storage containers being thin-provisioned to provide virtual data storage capacity which is greater than a real data storage capacity of the data storage hardware
said method includes:
(a) determining a real data storage capacity of the data storage system for accommodating the one or more new volumes;
(b) determining over-allocation information relating to said one or more data storage containers;
(c) determining extrapolated future anticipated use of the one or more containers of the data storage system from historical data storage use information;
(d) selecting one or more candidate data storage containers on the basis of information from steps (a) to (c).

Favorably, the present invention provides data storage systems which are operable to employ a more efficient method of identifying optimal placements of one or more new data volumes in response to a request for storage capacity provisioning, for example from one or more software applications executing on

computing hardware. The more efficient method involves utilizing information regarding current configuration and current usage of data storage sub-systems of the data storage system for proposing one or more suitable containers for use in creating new storage volumes for storing the one or more new data volumes while being synchronized with general planning for a configuration adopted for the data storage sub-systems implementing thin-provisioning. Such utilization of information involves defining a set of policies describing intended setup for various data containers and/or from a definition of a basic container configuration itself. Such basic container configuration concerns, for example, maximum real and maximum virtual memory sizes.

According to a favorable embodiment, said method may include selecting one or more candidate containers for the one or more new volumes, based upon results from steps (a) to (c) being compared with one or more data storage thresholds and/or being compared with policy rules defining operation of the data storage system.

Another favorable embodiment of said method may include:
- finding one or more containers which have sufficient virtual data storage capacity for the one or more new volumes;
- finding one or more containers which have sufficient real data storage capacity or the one or more new volumes;
- finding one or more containers which are below their over-allocation capacity with regard to thin partitioning; and
- finding one or more containers in respect of (a) to (c) whose data storage utilization does not exceed thresholds and best practice values.

According to a further favorable embodiment, information regarding utilization of the data storage capacity of the system may be stored in persistent memory for use in selection of the

one or more candidate containers for storing the one or more new volumes.

A data storage system for storing one or more volumes of data, said system including one or more data storage containers provided by data storage hardware, said one or more data storage containers being thin-provisioned to provide virtual data storage capacity which is greater than a real data storage capacity of the data storage hardware, characterized in that said system includes a memory management arrangement which is operable:
(a) to determine a real data storage capacity of the data storage system for accommodating the one or more new volumes;
(b) to determine over-allocation information relating to said one or more data storage containers;
(c) to determine extrapolated future anticipated use of the one or more containers of the data storage system from historical data storage use information; and
(d) to select one or more candidate data storage containers on the basis of information from steps (a) to (c).

Expediently, said memory management arrangement may be operable to select one or more candidate containers for the one or more new volumes, based upon results from over-allocation information, and extrapolated future anticipated use being compared with one or more data storage thresholds and/or being compared with policy rules defining operation of the data storage system.

According to a favorable embodiment, said memory management arrangement may be operable:
- to find one or more containers which have sufficient virtual data storage capacity for the one or more new volumes;
- to find one or more containers which have sufficient real data storage capacity or the one or more new volumes;

- to find one or more containers which are below their over-allocation capacity with regard to thin partitioning; and
- to find one or more containers in respect of over-allocation information, and extrapolated future anticipated use whose data storage utilization does not exceed thresholds and best practice values.

Further, said system may include persistent memory for storing information regarding a utilization of the data storage capacity of the system, said information for use in selection of the one or more candidate containers for storing the one or more new volumes.

According to another aspect of the invention, a program product is proposed comprising a computer useable medium having a computer readable program, wherein the computer readable program when executed on a computer causes the computer to perform a method according to any one of the method features described above, particularly
(a) determining a real data storage capacity of the data storage system for accommodating the one or more new volumes;
(b) determining over-allocation information relating to said one or more data storage containers;
(c) determining extrapolated future anticipated use of the one or more containers of the data storage system from historical data storage use information;
(d) selecting one or more candidate data storage containers on the basis of information from steps (a) to (c).

Further, a program product is proposed comprising a computer useable medium including a computer readable program, wherein the computer readable program when executed on a computer causes the computer to execute a method according to any one of the method features described above, particularly

(a) determining a real data storage capacity of the data storage system for accommodating the one or more new volumes;
(b) determining over-allocation information relating to said one or more data storage containers (120);
(c) determining extrapolated future anticipated use of the one or more containers (120) of the data storage system (100) from historical data storage use information;
(d) selecting one or more candidate data storage containers on the basis of information from steps (a) to (c).

Further, a data processing system is proposed for execution of a data processing program comprising software code portions for performing a method when said program is run on said computer, wherein the method includes any one of the method features described above, particularly
(a) determining a real data storage capacity of the data storage system for accommodating the one or more new volumes;
(b) determining over-allocation information relating to said one or more data storage containers;
(c) determining extrapolated future anticipated use of the one or more containers of the data storage system from historical data storage use information;
(d) selecting one or more candidate data storage containers on the basis of information from steps (a) to (c).

Further, a data processing program for execution in a data processing system comprising software code portions for performing a method when said program is run on a computer, wherein the method steps include any one of the method features described above, particularly
(a) determining a real data storage capacity of the data storage system for accommodating the one or more new volumes;
(b) determining over-allocation information relating to said one or more data storage containers (120);

(c) determining extrapolated future anticipated use of the one or more containers (120) of the data storage system (100) from historical data storage use information;
(d) selecting one or more candidate data storage containers on the basis of information from steps (a) to (c).

Advantageously, the present invention is concerned with tying together a large set of information describing data storage allocation, especially in respect of historical information about data memory capacity usage of already existing storage volumes and containers. Such tying up of data pursuant to the present invention involves extrapolating changes in capacity configuration into the future for the data storage system. In other words, the present invention is involved with a method including:

(a) monitoring use of data storage capacity to create historical data describing data storage use;
(b) extrapolating on the basis of this historical data future expected data storage requirements;
(c) creating new volumes in containers in response to request for storage in combination with the future expected data storage requirements.

Step (c) takes into account one or more policies pertaining to operation of the storage containers, capacity limits for the storage containers and policies regarding best practice. Other constrains of priority can also optionally pertain wherein software applications controlling time-critical processes are given higher priority when requesting allocation of data storage capacity for new volumes in comparison to software applications which are non-critical and can be temporarily halted or delayed. In step (c), processing of information relating to all candidate containers allows for an optimal container for one or more new data volumes to be selected, thereby resulting in more optimized

utilization of storage infrastructure and hence more effective utilization of storage capacity in the data storage system.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention together with the above-mentioned and other objects and advantages may best be understood from the following detailed description of the embodiments, but not restricted to the embodiments, wherein is shown in:

Fig. 1 a known computer system executing computer applications having access to thin-partitioned data memory;

Fig. 2 a data storage system implemented pursuant to the present invention; and

Fig. 3a, 3b a method employed for operating the system of Fig. 2.

In the drawings, like elements are referred to with equal reference numerals. The drawings are merely schematic representations, not intended to portray specific parameters of the invention. Moreover, the drawings are intended to depict only typical embodiments of the invention and therefore should not be considered as limiting the scope of the invention.

DETAILED DESCRIPTION OF EXAMPLE EMBODIMENTS

When implementing the present invention, the following information is taken into consideration when deciding placement of one or more new data volumes and/or one or more new containers for such volumes:

(i) the types of candidate containers that are available, namely for determining whether or not the candidate

containers are suitable for storing one or more new volumes;

(ii) current configuration limits for existing data storage containers;

(iii) specific policies relating to data storage use of the containers and/or future policies to be implemented in relation thereto;

(iv) a current real utilization of containers of the system for storing data volumes;

(v) historical data of existing volumes stored in the containers, and also extrapolated trends from the historical data;

(vi) existing volumes and volumes that would be created in response to request for storage space for one or more new volumes, for example including a measure of over-allocation utilized for each container in the data storage system.

A practical example of implementation of the present invention will now be described with reference to Fig. 2. In Fig. 2, a data storage system is indicated generally by 100 including a data sub-system 110 comprising one or more data containers 120 which are used in operation by the sub-system 100 for storing one or more volumes of data. Moreover, the sub-system 100 employs thin-provisioning such that data memory associated with the containers 120 is subject to potential over-allocation for rendering a virtual data storage capacity provided by the system 100 larger than its real actual data storage capacity. The system 100 also includes a management controller 160 which is responsible for managing data storage hardware of the sub-system 110, for example its data storage disc drives, solid state data memory and so forth. Moreover, the management controller 160 is operable to control and execute all operations within the sub-system 110, for example for handling external management requests 200 invoked through an external interface 180 of the

sub-system 110. Such requests 200 can include, for example, a command line interface (CLI), and can provide all necessary information about a configuration of the sub-system 110 and its status in respect of its internal storage containers 120.

The sub-system 110 additionally includes a data collection and evaluation unit 130 which is responsible for periodically collecting information about the configuration of the one or more containers 120, especially with regard to their data storage capacity key metrics. Moreover, the evaluation unit 130 is provided with it own persistent data store 140 for later evaluation, for example for generating historical data describing utilization of the one or more containers 120 for use in executing extrapolation in respect of future anticipated data storage requirements. Thus, in this manner, data stored in the persistent data store 140 is operable to provide a history of key capacity metrics for the sub-system 110.

Once a request to create a new volume is issued against the system 100 through its management interface 180, the management controller 160 invokes the collection and evaluation unit 130 which is operable to pull data from the data store 140, compute trends, pull in additional information about constraints and best practices from the configuration data store 150 to determine a best-suited container 120 for use in storing one or more new volumes; the constraints concern, for example, maximum over-allocation limits for the containers 120.

Operation of the system 100 will now be described with reference to Fig. 2, Fig. 3a and Fig. 3b. In Fig. 3a and Fig.3b, there is shown a flow chart of a method according to the invention.

The method in Figs. 3a and 3b, in order to determine a suitable container to use in the system 100 for storing one or more new volumes, one or more checks have to be undertaken in a list of

candidate containers 120 in order to dismiss those containers 120 which do not meet basic requirements for accommodating the one or more new volumes. These one or more checks are performed to determine:

(a) whether or not a given container 120 is capable of providing enough virtual space for the requested size of the one or more new volumes, see steps 310, 320 in Fig. 3a;

(b) whether or not there is enough real capacity for an initial requested real size of data to be stored; optionally, the system 100 is operable to allow a degree of additional buffer memory capacity to cope with future growth based upon configuration data stored in the configuration data stored 150, see steps 330 and 340 in Fig. 3a; and

(c) whether or not a current over-allocation factor is already higher or close to a specified maximum over-allocation factor for the given container 120 based upon configuration data stored in the configuration data stored 150, see steps 350, 360 in Fig. 3a.

After start step 300, in step 310 containers 120 are checked for sufficient virtual space. In step 320 it is checked whether any container 120 has enough virtual capacity. If no container 120 has enough virtual capacity ("n" in the flow chart), the test ends at step 390 as no container 120 is found. If there is one or more containers 120 with enough virtual capacity ("y" in the flow chart) it is continued with step 330 where containers 120 are checked for sufficient starting real capacity. In subsequent step 340 it is checked whether any container 120 has enough real capacity. If no container has enough real capacity ("n" in the flow chart) the test ends at step 390 as no container 120 is found. If any container 120 has enough real capacity ("y" in the flow chart) it is continued with step 350, where the containers 20 are checked for over-allocation. In step 360 it is checked whether any container 120 is below its maximum over-allocation

factor. If no ("n" in the flow chart) the test ends at stop step 390 as no container 120 is found.

If yes ("y" in the flow chart) a list of all containers 120 is created which are ordered by distance of over-allocation factor from its maximum value in step 370. From containers 120 in the sub-system 110 which pass tests in (a) to (c), a candidate container 120 list is generated by the method, the list being then subsequently sorted in respect of a degree to which its container 120 current over-allocation factors are away from the maximum over-allocation factor for the containers in the list; this is implemented in step 370. If none of the containers 120 pass such a test pursuant to the present invention ("n" in the slow chart), the method is operable to provide a message in return informing of no suitable container 120 having been found (step 390).

If the list is not empty ("y" in the flow chart), namely one or more candidate containers 120 have been found for which further iteration included in the method can be applied as in step 380, where it is checked whether or not there are more containers on the created list, a container 120 is selected from the list in step 400. For the selected container 120, a hypothetical over-allocation factor is computed in an event that one or more new volumes were to be stored in the selected container 120 in step 410.

Moreover, the hypothetical over-allocation factor is checked against the maximum over-allocation factor for the selected container 120 and a further check is made whether ("y" in the flow chart) or not ("n" in the flow chart) the computed hypothetical over-allocation factor violates best practices as defined by the configuration data stored in the configuration data store 150, namely as provided in steps 410, 420. If this test based upon the hypothetical over-allocation is passed ("n"

in the flow chart), a future trend for virtual and real utilization of data memory in the sub-system 110 is computed based upon historical data memory usage data retained in the persistent data store 140. Moreover, assumed values for a specified time-range are calculated, in step 430; in other words, a time is specified for when the one or more new volumes are to be stored, and optionally a time period during which the one or more new volumes need to be stored in the sub-system 110, namely in step 430.

In subsequent step 440 the assumed space is checked against the thresholds defined for the sub-system 110 and best practice values and if there the threshold or best practice values are not exceeded ("n" in the flow chart), the capacity of the one or more new volumes is included in the trend calculation in step 450.

The computer values are checked against thresholds defined for the sub-system 110 and best rules for best practices in step 460. If all checks at steps 420, 440, 460 show best practices and thresholds are not violated ("n" in the flow chart), then the selected container 120 is successful and the process ends at step 470. The selected container 120 is the one to be used to accommodate the one or more new volumes. Conversely, in an event that one or more of the checks at steps 420, 440, 460 determine a violation of best practices or thresholds ("y" in the flow chart), the method progresses to consider a next possible candidate container 120 in the aforementioned list.

In conclusion, the present invention is concerned with a task of finding data memory of a data storage sub-system for accommodating one or more new volumes of data, for example to be generated by one or more software applications executing upon computing hardware. The one or more new volumes are to be stored in one or more data containers of the sub-system, wherein

the one or more data containers occupy data memory which is thin-provisioned, namely effectively over-allocated. Such over-allocation needs to be managed to avoid a situation wherein data memory is exceeded in the sub-system such that loss of access or even of data can occur. Moreover, management of data storage resources in the sub-system are subject to various policy rules. In order to improve operation of the sub-system, the present invention concerns a method of identifying one or more eligible containers for storing the one or more new data volumes and generating a list of potentially eligible containers. The method involves:

(a) finding one or more containers which have sufficient virtual data storage capacity;
(b) finding one or more containers which have sufficient real data storage capacity;
(c) finding one or more containers which are below their over-allocation capacity with regard to thin partitioning;
(d) finding one or more containers in respect of (a) to (c) whose data storage utilization does not exceed thresholds and best practice values.

If conditions (a) to (d) are satisfied for a given container in the list, wherein future predictions of real and/or virtual capacity are determined by extrapolating historic values for example, the given container is then adopted for use in storing the one or more new data volumes. The present invention enables data storage capacity in data storage sub-systems to be utilized more efficiently, for example for enabling a larger number of concurrently executing software applications requiring data memory access to be accommodated in the computing hardware, thereby, for example, increasing functionality offered in operation by the computing hardware.

The present invention is, for example, distinguished in that it provides a method of placing one or more new volumes based on

real storage capacity information, over-allocation information and extrapolated anticipated data storage requirements. By taking all these factors into consideration, more efficient utilization of thin-partitioned data storage memory is possible.

The invention can take the form of an entirely hardware embodiment, an entirely software embodiment or an embodiment containing both hardware and software elements. In a preferred embodiment, the invention is implemented in software, which includes but is not limited to firmware, resident software, microcode, and so forth.

Furthermore, the invention can take the form of a computer program product accessible from a computer-usable or computer readable medium providing program code for use by or in connection with a computer or any instruction execution system. For the purposes of this description, a computer-usable or computer readable medium can be any apparatus that can contain, store, communicate, propagate, or transport the program for use by on in connection with the instruction execution system, apparatus, or device.

The medium can be an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system (or apparatus or device) or a propagation medium. Examples of a computer-readable medium include a semiconductor or solid state memory, magnetic tape, a removable computer diskette, a random access memory (RAM), a read-only memory (ROM), a rigid magnetic disk and an optical disk. Current examples of optical disks include compact disk-read-only memory (CD-ROM), compact disk-read/write (CD-R/W) and DVD.

A data processing system suitable for storing and/or executing program code will include at least one processor coupled directly or indirectly to memory elements through a system bus.

The memory elements can include local memory employed during actual execution of the program code, bulk storage, and cache memories which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution.

Input/output or I/O-devices (including, but not limited to keyboards, displays, pointing devices, etc.) can be coupled to the system either directly or through intervening I/O controllers.

Network adapters may also be coupled to the system to enable the data processing system or remote printers or storage devices through intervening private or public networks. Modems, cable modem and Ethernet cards are just a few of the currently available types of network adapters.

C L A I M S

1. A method of storing one or more volumes of data in a data storage system (100) including one or more data storage containers (120) provided by data storage hardware (110), said one or more data storage containers (120) being thin-provisioned to provide virtual data storage capacity which is greater than a real data storage capacity of the data storage hardware (100), wherein said method includes:
(a) determining a real data storage capacity of the data storage system (100) for accommodating the one or more new volumes;
(b) determining over-allocation information relating to said one or more data storage containers (120);
(c) determining extrapolated future anticipated use of the one or more containers (120) of the data storage system (100) from historical data storage use information;
(d) selecting one or more candidate data storage containers (120) on the basis of information from steps (a) to (c).

2. The method according to claim 1, wherein said method includes selecting one or more candidate containers for the one or more new volumes, based upon results from steps (a) to (c) being compared with one or more data storage thresholds and/or being compared with policy rules defining operation of the data storage system (100).

3. The method according to claim 1 or 2, wherein said method includes:
(i) finding one or more containers which have sufficient virtual data storage capacity for the one or more new volumes;
(ii) finding one or more containers which have sufficient real data storage capacity or the one or more new volumes;
(iii) finding one or more containers which are below their over-allocation capacity with regard to thin partitioning; and

(iv) finding one or more containers in respect of (a) to (c) whose data storage utilization does not exceed thresholds and best practice values.

4. The method according to in any one of the preceding claims, wherein information regarding utilization of the data storage capacity of the system (100) is stored in persistent memory (140) for use in selection of the one or more candidate containers (120)for storing the one or more new volumes.

5. A data storage system (100) for storing one or more volumes of data, said system (100) including one or more data storage containers (120) provided by data storage hardware (110), said one or more data storage containers (120) being thin-provisioned to provide virtual data storage capacity which is greater than a real data storage capacity of the data storage hardware (100), characterized in that said system (100) includes a memory management arrangement (160, 130, 140, 150)which is operable:
(a) to determine a real data storage capacity of the data storage system (100) for accommodating the one or more new volumes;
(b) to determine over-allocation information relating to said one or more data storage containers (120);
(c) to determine extrapolated future anticipated use of the one or more containers (120) of the data storage system (100) from historical data storage use information; and
(d) to select one or more candidate data storage containers (120) on the basis of information from steps (a) to (c).

6. The system according to claim 5, wherein said memory management arrangement (160, 130, 140, 150) is operable to select one or more candidate containers for the one or more new volumes, based upon results from over-allocation information, and extrapolated future anticipated use being compared with one or more data storage thresholds and/or being compared with

policy rules defining operation of the data storage system (100).

7. The system according to claim 5 or 6, wherein said memory management arrangement (160, 130, 140, 150)is operable:

(i) to find one or more containers which have sufficient virtual data storage capacity for the one or more new volumes;

(ii) to find one or more containers which have sufficient real data storage capacity or the one or more new volumes;

(iii)to find one or more containers which are below their over-allocation capacity with regard to thin partitioning; and

(iv) to find one or more containers in respect of over-allocation information, and extrapolated future anticipated use whose data storage utilization does not exceed thresholds and best practice values.

8. The system according to any one of claims 5 to 7, wherein said system includes persistent memory (140) for storing information regarding a utilization of the data storage capacity of the system (100), said information for use in selection of the one or more candidate containers (120)for storing the one or more new volumes.

9. A program product comprising a computer useable medium having a computer readable program, wherein the computer readable program when executed on a computer causes the computer to perform a method to any one of claims 1 to 4.

10. A data processing system for execution of a data processing program comprising software code portions for performing a method according to any one of the claims 1 to 4 when said program is run on said computer.

11. A data processing program for execution in a data processing system comprising software code portions for

performing a method when said program is run on a computer, wherein the method steps are defined in any one of claims 1 to 4.

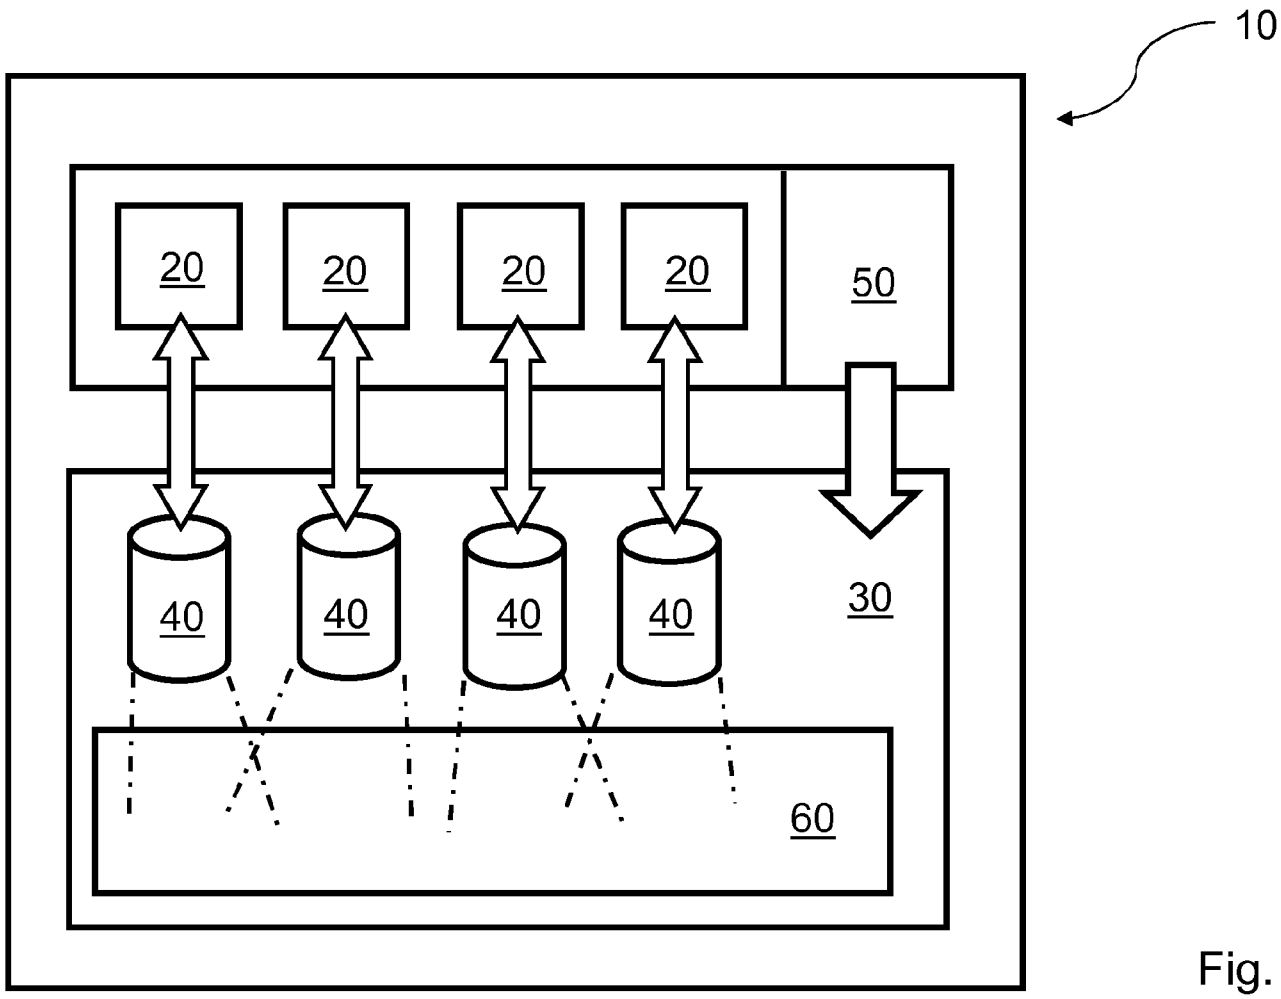

Fig. 1

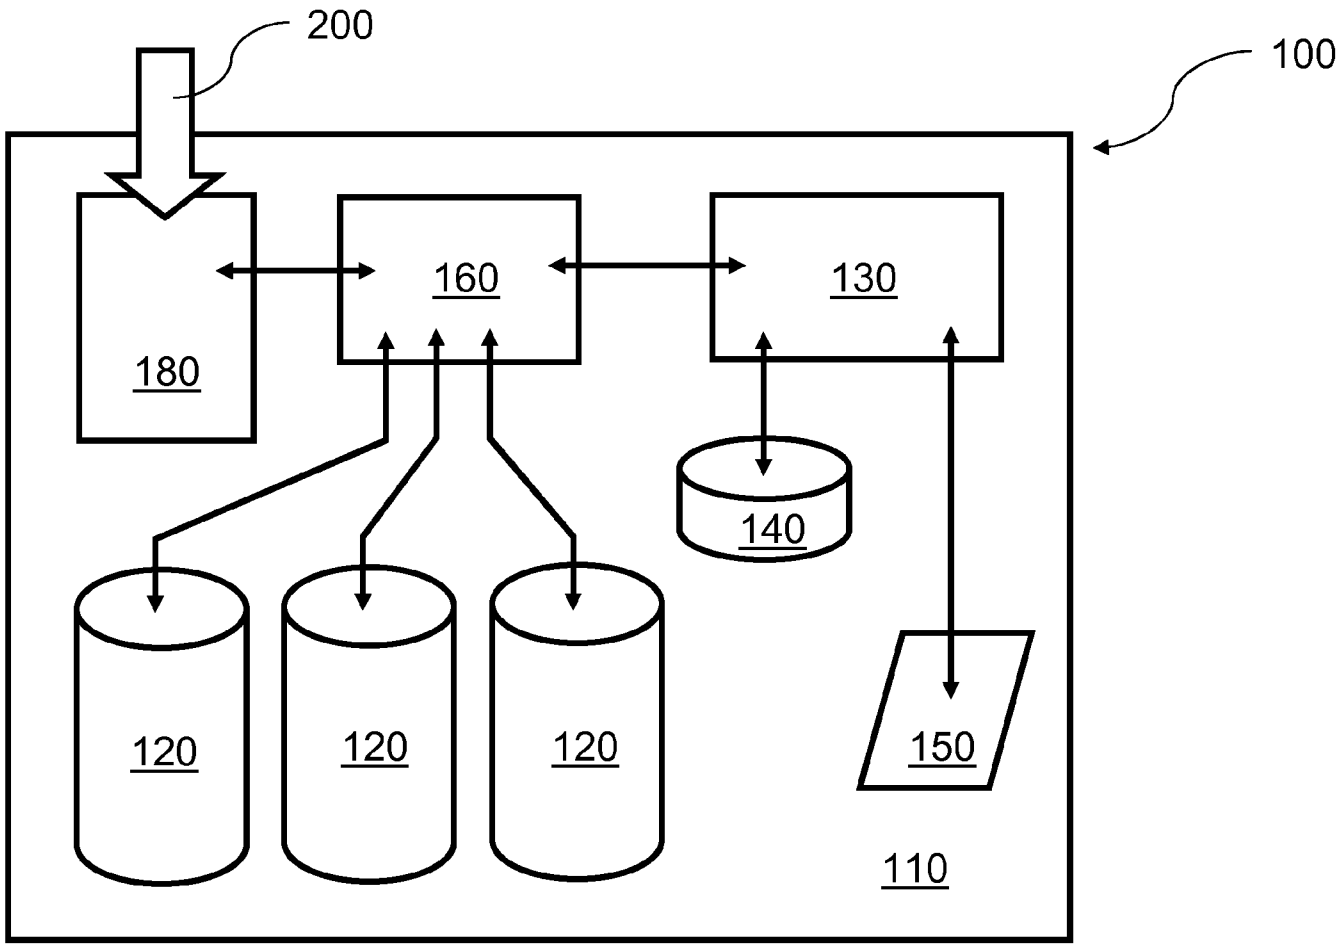

Fig. 2

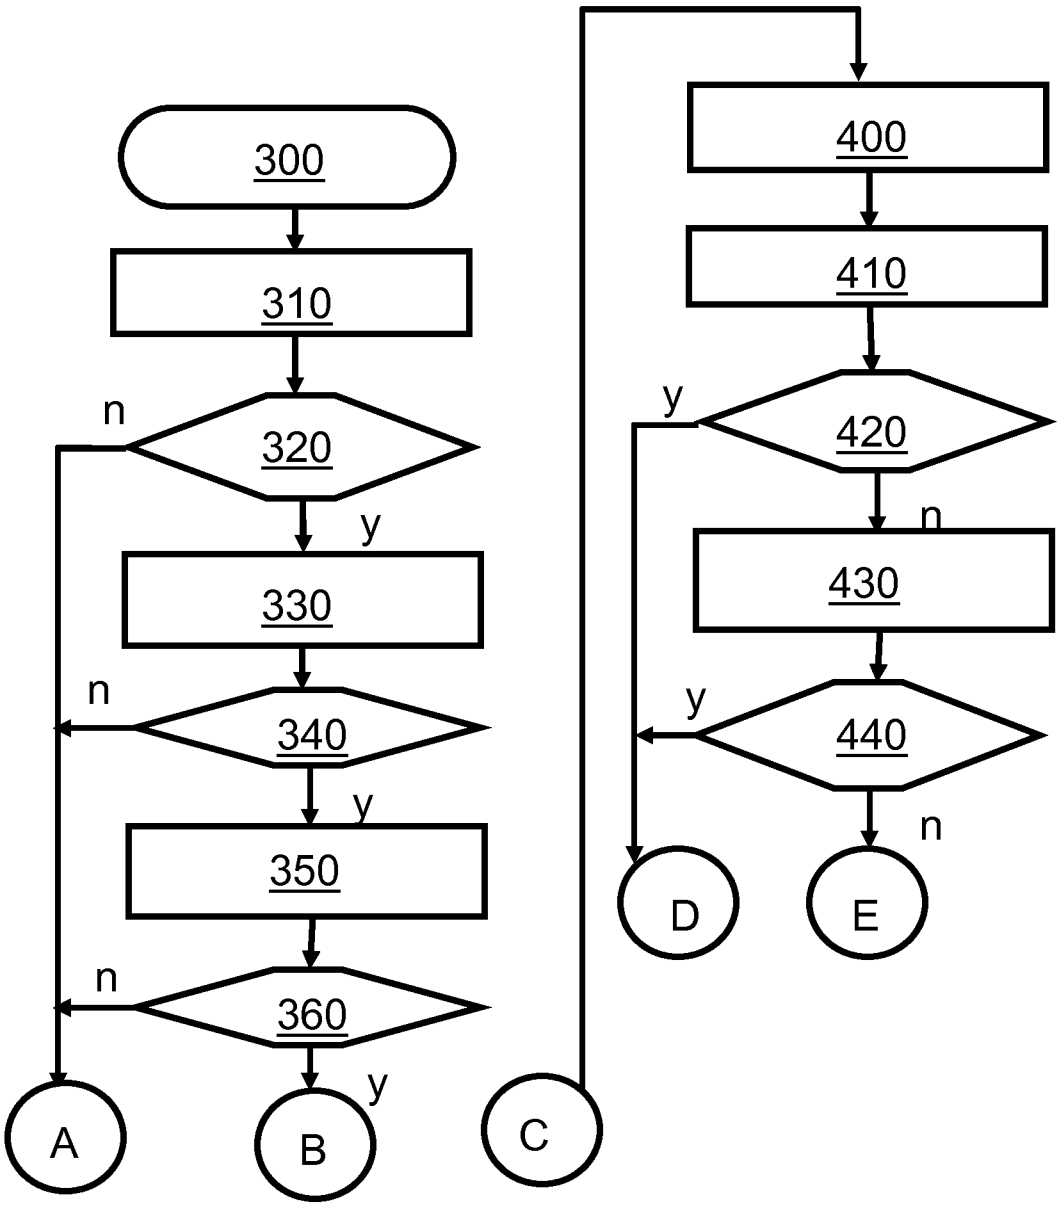
300
310
320
n
y
330
340
n
y
350
360
n
y
A
B
C
400
410
420
y
n
430
440
y
n
D
E

Fig. 3a

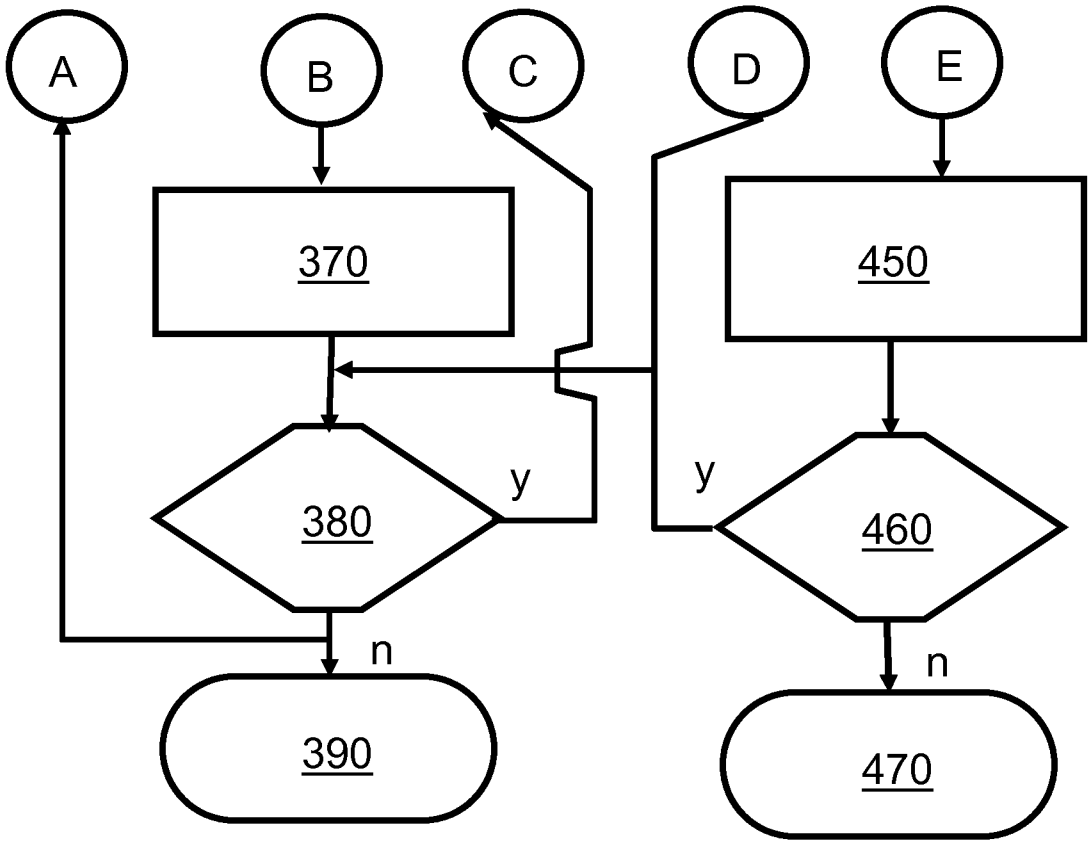

Fig. 3b

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2010/064959

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F3/06
ADD.

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)
G06F

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

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2007/233868 A1 (TYRRELL JOHN C [US] ET AL) 4 October 2007 (2007-10-04) paragraph [0041] - paragraph [0043]; claims 1,5; figures 1,4,16 paragraphs [0058], [0068], [0075], [0076], [0083], [0084], [0087], [0089]	1-11
Y	US 2007/239793 A1 (TYRRELL JOHN C [US] ET AL) 11 October 2007 (2007-10-11) paragraphs [0035], [0086] - paragraph [0094]; figures 1,16-19	1-11
A	US 2008/091748 A1 (BENIYAMA NOBUO [JP] ET AL) 17 April 2008 (2008-04-17) cited in the application paragraph [0033] - paragraph [0121]; figures 2-4,7,9,12	1-11

☐ 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 document 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 another 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	Date of mailing of the international search report
28 December 2010	05/01/2011

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31–70) 340–2040, Fax: (+31–70) 340–3016	Moens, Rudi

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/EP2010/064959

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007233868 A1	04-10-2007	EP 2002327 A2	17-12-2008
		WO 2007120524 A2	25-10-2007
US 2007239793 A1	11-10-2007	NONE	
US 2008091748 A1	17-04-2008	JP 2008097502 A	24-04-2008
		US 2010191906 A1	29-07-2010