

(51) International Patent Classification:
G06F 11/14 (2006.01)(21) International Application Number:
PCT/US2015/050867(22) International Filing Date:
18 September 2015 (18.09.2015)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
2885/CHE/2015 9 June 2015 (09.06.2015) IN(71) Applicant: **HEWLETT PACKARD ENTERPRISE DEVELOPMENT LP** [US/US]; 11445 Compaq Center Drive West, Houston, TX 77070 (US).(72) Inventors: **KAPANIPATHI, Sri Harshan**; Sy.No.192, Whitefield Road, Mahadevapura Post, Bangalore, Karnataka 560048 (IN). **VASUDEVAN, Deepak**; Sy.No.192, Whitefield Road, Mahadevapura Post, Bangalore, Karnataka 560048 (IN). **THAKER, Deven Mahendraray**; Sy.No.192, Whitefield Road, Mahadevapura Post, Bangalore, Karnataka 560048 (IN).(74) Agents: **ORTEGA, Arthur S.** et al.; Hewlett Packard Enterprise, 3404 E. Harmony Road, Mail Stop 79, Fort Collins, CO 80528 (US).

(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, BN, 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, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, 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, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, 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, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- 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))

[Continued on next page]

(54) Title: DATA REPLICATION MANAGEMENT




FIG. 1

(57) Abstract: In one example, initiating data replication process to copy source data set as target data sets to a plurality of target devices during a predefined time period, determining progress information regarding completion of the replication process based on historical and current performance information of the plurality of target devices, and if the progress information indicates that the replication process is not to be completed within a predefined time period, then selecting a target device from the plurality of target devices to complete the complete the replication process based on historical and current performance information of the plurality of target devices, and causing the selected target device to continue with the completion of the replication process, causing the non-selected target devices to discontinue with the replication process, and causing the replicated target data set to be copied to another target device of the plurality of target devices.



Published:

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

DATA REPLICATION MANAGEMENT

BACKGROUND

[0001] Computer systems include computing devices that communicate with other computing devices over a network. The computing devices may include storage devices to store data for later retrieval. The computing devices may copy data from one computing device to other computing devices.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Examples are described in the following detailed description and in reference to the drawings, in which:

[0003] Fig. 1 depicts an example system for data replication management in accordance with an example of the techniques of the present application;

[0004] Fig. 2 depicts an example flow chart of processes for data replication management in accordance with an example of the techniques of the present application;

[0005] Figs. 3A through 3B depict example diagrams of data replication management in accordance with an example of the techniques of the present application; and

[0006] Fig. 4 depicts an example block diagram showing a non-transitory, computer-readable medium that stores instructions for data replication management in accordance with an example of the techniques of the present application.

DETAILED DESCRIPTION

[0007] Computer systems include computing devices that communicate with other computing devices over a network. The computing devices may include storage devices to store data for later retrieval. The computing devices may copy data from one computing device to other computing devices.

[0008] Computer system environments may implement data replication techniques to copy data from one computing device to another computing device for data availability purposes. For example, a source computing device may implement a data replication process to copy or replicate data to a plurality of target computing devices to provide a backup for data availability purposes. For example, in the event that one of the computing devices encounter an error condition such that the data may no longer be available, then a copy of the data may be retrieved from one of the other target computing devices. Computer system environments may be required to manage increased amounts of data

including data volumes of unstructured data. In addition, computer systems may be configured as virtualization and cloud infrastructures to meet ever increasing resource demands. The computing devices may implement data replication techniques to provide high data availability. In one example, the computing devices may be configured to implement data replication techniques where copies are made to a plurality of target computing devices located across different locations or sites. However, data replication in such a configuration may encounter problems such as resource constraints or increases in data volume activity, such that the data replication process may not make copies in a timely manner which may result in some targets not having copies.

[0009] The present application discloses techniques to provide data replication management techniques which may improve data replication from a source computing device to across a plurality of target computing devices which may increase the performance of the system. For example, the data replication techniques provide at least one copy of data to a target computing device to help ensure data availability even during resource constraint conditions such as bandwidth or data throughput constraints.

[0010] In one example of the present techniques, disclosed is a data replication management module configured to initiate a data replication process to copy a source data set as target data sets to a plurality of target computing devices during a predefined time period. The management module may determine progress information regarding completion of the data replication process based on historical and current performance information of the plurality of target computing devices.

[0011] If the data replication management module determines that the progress information indicates that the data replication process is not to be completed within the predefined time period, then the management module may select a target computing device from the plurality of target computing devices to complete the data replication process based on historical and current performance information of the plurality of target computing devices. In one example, the data replication management module selects the target computing device that has highest performance compared to those of the other target computing devices. In this case, the management module may send a continue notification to the selected target computing device to cause the selected target computing device to continue with the completion of the data replication process. The management module may send a discontinue notification to non-selected target computing devices to cause the non-selected target computing devices to discontinue with the data replication process. The management module may cause the selected target computing

device to initiate a data replication process to copy the replicated target data set to another target computing device of the plurality of target computing devices.

[00012] On the other hand, if the data replication management module determines that the progress information indicates that the data replication process is to be completed within the predefined time period, then the module may send continue notifications to the target computing devices to cause the target computing device to continue with the completion of the data replication process.

[00013] In another example, the data replication management module may be configured to calculate the current performance information of a target computing device which relates to data throughput during the data replication process and calculate the historical performance information of the target computing device which relates to data throughput during previous data replication processes. In another example, the management module may be configured to release processing resources used by the computing device associated with the non-selected target computing devices for use by the computing device for the data replication process of the selected target computing device. In this manner, releasing resources may allow the computing device to use the resources to complete the data replication process with the selected target computing device. In another example, management module may be configured to receive from a host computer a request to initiate the data replication process.

[00014] In this manner, the present application discloses techniques to provide data replication management techniques which may help improve data replication from a source computing device to across a plurality of target computing devices which may increase the performance of the system. For example, the data replication techniques such that even during resource constraint at least one copy of data is made to help ensure data availability.

[00015] Fig. 1 depicts an example system 100 for data replication management in accordance with an example of the techniques of the present application. The system 100 includes a computing device 102 configured to communicate with target computing devices 112 over a network 110.

[00016] The computing device 102 includes a data replication management module 104 and a storage device 106 to store a source data set 108. The data replication management module 104 is configured to manage a data replication process to replicate or copy the source data set 108 to target data sets 114 (114-1 through 114-n) of target computing devices 112 (112-1 through 112-n), as explained below in further detail.

[00017] The data replication management module 104 may be configured to initiate a data replication process to copy the source data set 108 as target data sets 114 to a plurality of target computing devices 112 during a predefined time period. The predefined time period may be based on system 100 criteria such as on the characteristics of the source data set 108 such as size of the data and the like. In one example, the time period may be a predefined replication time window which may be configured after a backup and during off production time. The management module 104 may determine progress information regarding completion of the replication process based on historical and current performance information of the plurality of target computing devices 112. In one example, historical performance information may be defined as performance related information or data of the target computing devices over the time period of the previous replication processes. In one example, current performance information may be defined as performance related information or data of the target computing devices over or during the time period of the current data replication process. In one example, progress information may be defined as to what extent the current data replication process is complete such that a copy of the source data set 108 is made to the target data sets 114. In another example, the replication module 104 may determine the prorated progress of the data replication process toward completion. In one example, performance information may include data throughput (such as in Megabits per second (Mbps)) which include the amount of data moved or transferred successfully from computing device 102 to a target computing device 112 in a given time period.

[00018] If data replication management module 104 determines that the progress information indicates that the data replication process is not to be completed within the predefined time period, then management module may select a target computing device 112 from the plurality of target computing devices 112 to complete the complete the data replication process based on historical and current performance information of the plurality of target computing devices. In one example, the data replication management module selects the target computing device that has highest performance compared to those of the other target computing devices. It should be understood that the data replication management module may select the target computing device based on other techniques. For example, selection techniques may be based a target computing device site preference, which may be set by a user, such as administrator of the system, based on a parameter such as backup consolidation. In other examples, selection techniques may be based on storage capacity of the target computing devices, target computing device capabilities such as high availability, target computing device uptime and the like. In this

case, management module 104 may send a continue notification to the selected target computing device 112 to cause the selected target computing device to continue with the completion of the data replication process. The management module 104 may send a discontinue notification to non-selected target computing devices 112 to cause the non-selected target computing devices to discontinue with the data replication process. The management module 104 may cause the selected target computing device 112 to initiate a data replication process to copy the replicated target data set to another target computing device of the plurality of target computing devices.

[00019] The data replication management module 104 may check the progress of the data replication process using techniques such as polling the status, interrupt or any other like technique. The “continue” and “discontinue” notifications may comprise any means of notifications such as commands, messages according to particular protocols and the like. The notifications may include notifications to start the data replication process, to stop the process, to pause the process, to restart the process and the like. The source data set 108 may include a group of data blocks associated with particular applications such as file systems, databases and the like. The data replication process may include a process of copying from computing device 102 to target computing devices 112 data blocks in a successive manner one block after another until a complete copy is made as target data sets 114 to the target computing devices. In another example, the replication process may be based on deduplication based replication techniques (also known as low bandwidth replication) such that data that is not at the target is not transferred.

[00020] On the other hand, if data replication management module 104 determines that the progress information indicates that the data replication process is to be completed within the predefined time period, then the module may send continue notifications to the target computing devices 112 to cause the target computing device to continue with the completion of the data replication process.

[00021] In another example, data replication management module 104 may be configured to calculate the current performance information of a target computing device 112 which relates to data throughput during the data replication process and calculate the historical performance information of the target computing device which relates to data throughput during previous data replication processes. The performance information of computing devices 112 may be any performance metric, such as data throughput, related to processing data to transfer data across network 110. In another example, management module 104 may be configured to release processing resources used by the computing device associated with the non-selected target computing devices for use by the

computing device for the data replication process of the selected target computing device. In one example, processing resources may any computing or processing resources for processing data associated with hardware, software or a combination thereof of the computing device. In another example, management module 104 may be configured to receive from a host computer a request to initiate the data replication process.

[00022] In another example, computing device 102 may be configured to provide storage management functionality as described herein. In one example, computing device 102 may provide functionality to process Input/Output (IO) requests from host computers to access source data set 108 on storage device 106.

[00023] The storage device 106 may be defined as any electronic means to store data for later retrieval. The source data set 108 may be stored as storage volumes which may by logical units of data that can be defined across multiple storage devices. The IO requests may include requests to read data from storage device 106 as volumes and requests to write data to the storage devices as volumes. The source data set 108 may be defined as a group of data or volumes to be replicated or copied from one computing device to another computing device. The computing device 102 and target computing devices 112 may include any electronic device or means of processing data. The storage device 106 may refer to a physical storage element, such as a disk-based storage element (e.g., hard disk drive, optical disk drive, etc.) or other type of storage element (e.g., semiconductor storage element). In one example, multiple storage devices within a storage subsystem can be arranged as an array configuration. The description of source data set 106 applies to the description of target data sets 114.

[00024] The network 110 may include any means of electronic or data communication between computing device 102 and target computing devices 112. The network 110 may include a local area network, Internet and the like. The network 110 may include Fibre Channel network, SCSI (Small Computer System Interface) link, Serial Attached SCSI (SAS) link and the like. The network 110 may include expanders, concentrators, routers, and other communications devices.

[00025] The system 100 of Fig. 1 shows example computing device 102 to communicate with target computing devices 112 and should be understood that other configurations may be employed to practice the techniques of the present application. For example, system 100 may be configured to include a plurality of computing devices 102 to communicate with a plurality of target computing devices 112. In another example, storage device 106 is shown as a single component but it should be understood that the storage device may be plurality of storage devices distributed across a plurality of

computing devices 102. In another example, source data set 108 is shown as a single component but it should be understood that the source data set may be plurality of source data sets distributed across a plurality of computing devices 102. In another example, data replication management module 104 is shown as a single component but it should be understood that the module may be plurality of modules distributed across a plurality of computing devices 102. The components of system 100 may be implemented in hardware, software or a combination thereof. In one example, data replication management module 104 may be implemented in hardware, software or a combination thereof. In another example, the functionality of the components of system 100 may be implemented using technology related to Personal Computers (PCs), server computers, tablet computers, mobile computers and the like.

[00026] Fig. 1 shows a system 100 to provide data replication management. The storage system 100 may include computer-readable storage medium comprising (e.g., encoded with) instructions executable by a processor to implement functionalities described herein in relation to Fig. 1. In some examples, the functionalities described herein in relation to instructions to implement data replication management module 104 functions, and any additional instructions described herein in relation to storage medium, may be implemented as engines or modules comprising any combination of hardware and programming to implement the functionalities of the modules or engines, as described below. The functions of data replication management module 104 may be implemented by a computing device which may be a server, blade enclosure, desktop computer, laptop (or notebook) computer, workstation, tablet computer, mobile phone, smart device, or any other processing device or equipment including a processing resource. In examples described herein, a processor may include, for example, one processor or multiple processors included in a single computing device or distributed across multiple computing devices.

[00027] Fig. 2 depicts an example flow chart 200 of a process for data replication management in accordance with an example of the techniques of the present application. To illustrate operation, it may be assumed that system 100 includes computing device 102 configured to communicate with target computing devices 112 over network 110. The data replication management module 104 may be configured to implement the techniques of the present application and functionality described herein.

[00028] It should be understood the process depicted in Fig. 2 represents generalized illustrations, and that other processes may be added or existing processes may be removed, modified, or rearranged without departing from the scope and spirit of

the present application. In addition, it should be understood that the processes may represent instructions stored on a computer-readable storage medium that, when executed, may cause a processor to respond, to perform actions, to change states, and/or to make decisions. Alternatively, the processes may represent functions and/or actions performed by functionally equivalent circuits like analog circuits, digital signal processing circuits, Application Specific Integrated Circuits (ASICs), or other hardware components associated with the system. Furthermore, the flow charts are not intended to limit the implementation of the present application, but rather the flow charts illustrate functional information to design/fabricate circuits, generate software, or use a combination of hardware and software to perform the illustrated processes.

[00029] The process 200 may begin at block 202, where data replication management module 104 initiates a data replication process to copy source data set 108 as target data sets 114 to a plurality of target computing devices 112 during a predefined time period. In one example, it may be assumed that a host computing device sent a notification such as a command to computing device 102 to initiate the data replication process. Processing proceeds to block 204.

[00030] At block 204, data replication management module 104 determines progress information regarding completion of the data replication process based on historical and current performance information of the plurality of target computing devices. Processing proceeds to block 206.

[00031] At block 206, data replication management module 104 checks progress information to determine whether the information indicates that the data replication process is not to be completed within a predefined time period. In one example, if data replication management module 104 determines that the data replication process is not to be completed within a predefined time period, then processing proceeds to block 208. On the other hand, in another example, if data replication management module 104 determines that the data replication process is to be completed within a predefined time period, then processing proceeds to End block. In this case, in one example, data replication management module 104 sends continue notifications to target computing devices 112 to cause the target computing devices to continue with the completion of the data replication process.

[00032] At block 208, data replication management module 104 selects a target computing device 112 from the plurality of target computing devices to complete the complete the data replication process based on historical and current performance information of the plurality of target computing devices. In one example, the data

replication management module selects the target computing device that has highest performance compared to those of the other target computing devices. It should be understood that the data replication management module may select the target computing device based on other techniques. Processing proceeds to block 210.

[00033] At block 210, data replication management module 104 causes the selected target computing device 112 to continue with the completion of the data replication process. In one example, data replication management module 104 sends a continue notification to the selected target computing device to cause the selected target computing device to continue with the completion of the data replication process. Processing proceeds to block 212.

[00034] At block 212, data replication management module 104 causes the non-selected target computing devices to discontinue with the data replication process. . In one example, data replication management module 104 sends a discontinue notification to non-selected target computing devices 112 to cause the non-selected target computing devices to discontinue with the data replication process. Processing proceeds to block 214.

[00035] At block 214, data replication management module 104 initiates a data replication process to copy the replicated target data set to another target computing device of the plurality of target computing devices. Processing proceeds to the End block for other processing.

[00036] In this manner, the present application discloses techniques to provide data replication management techniques which may help improve data replication from a source computing device to across a plurality of target computing devices which may increase the performance of the system. For example, the data replication techniques such that even during resource constraint at least one copy of data is made to help ensure data availability.

[00037] The process 200 of Fig. 2 shows an example process and it should be understood that other configurations can be employed to practice the techniques of the present application. For example, process 200 may be configured to communicate with a plurality of computing devices 102 and the like.

[00038] Figs. 3A through 3B depicts example diagrams of processes for data replication management in accordance with an example of the techniques of the present application. To illustrate operation, it may be assumed that system 100 includes computing device 102 configured to communicate with target computing devices 112 over

network 110. The data replication management module 104 may be configured to implement the techniques of the present application and functionality described herein.

[00039] Fig. 3A shows a table 300 to illustrate an example process of the data replication management techniques of the present application. Fig. 3B shows a graph 320 to illustrate an example process of the data replication management techniques of the present application.

[00040] Turning to Fig. 3A, to illustrate operation, in one example, it may be assumed that system 100 includes computing device 102 configured to communicate with four target computing devices 112-1, 112-2, 112-3, 112-4 over network 110. The data replication management module 104 may be configured to implement the techniques of the present application and functionality described herein.

[00041] In this example, data replication management module 104 of computing devices 102 may be configured to initiate a data replication process to copy source data set 108 as target data sets 114-1, 114-2, 114-3, 114-4 associated with respective target computing devices 112-1, 112-2, 112-3, 112-4 during a predefined time period of 200 minutes. In one example, management module 104 may be configured to receive from a host computer a request to initiate the data replication process. The management module 104 determines progress information regarding completion of the data replication process based on historical and current performance information of target computing devices 112-1, 112-2, 112-3, 112-4. In this case, management module 104 determines progress by polling the progress in 5 minutes intervals during the time period or window starting from time = 5 to time =165 minutes, as shown by column 302 of table 300.

[00042] The data replication management module 104 determines the performance information of target computing devices 112-1, 112-2, 112-3, 112-4. For example, management module 104 determines current performance information of target computing devices 112-1, 112-2, 112-3, 112-4 as shown by respective current performance information columns 304-1, 304-2, 304-3, 304-4 of table 300. In this case, current performance information relates to data throughput between computing device 102 and each of target computing devices 112-1, 112-2, 112-3, 112-4.

[00043] The data replication management module 104 determines the overall progress information. In this case, management module 104 determines the current prorated performance information for all of the target computing devices, as shown by column 306 of table 300. In this case, the current performance information relates to the current data replication process.

[00044] The data replication management module 104 determines the historical performance information as shown by column 308 of table 300. In this case, the historical performance information relates to the average historical data throughput of past or previous data replication processes.

[00045] The data management module 104 determines progress information based on current and historical performance information of target computing devices 112-1, 112-2, 112-3, 112-4. In this case, in one example, management module 104 determines or calculates progress information based on current performance information of target computing devices 112-1, 112-2, 112-3, 112-4 as shown by respective columns 304-1, 304-2, 304-3, 304-4 of table 300. In this case, in one example, management module 104 determines or calculates progress information based on average historical information of target computing devices 112-1, 112-2, 112-3, 112-4 as shown by column 308 of table 300.

[00046] The data management module 104 determines whether the progress information indicates that the data replication process is not to be completed within the predefined time period.

[00047] If data management module 104 determines that the progress information indicates that the data replication process is not to be completed within the predefined time period, then management module may select a target computing device 112 from plurality of target computing devices 112-1, 112-2, 112-3, 112-4 to complete the complete the data replication process based on historical and current performance information of the plurality of target computing devices. In one example, the data replication management module selects the target computing device that has highest performance compared to those of the other target computing devices. It should be understood that the data replication management module may select the target computing device based on other techniques

[00048] On the other hand, if data replication management module 104 determines that the progress information indicates that the data replication process is to be completed within the predefined time period, then the module may send continue notifications to target computing devices 112-1, 112-2, 112-3, 112-4 to cause the target computing device to continue with the completion of the data replication process.

[00049] In this case, at time = 70 minutes (column 302), data replication management module 104 determines that the overall prorated progress information (column 306) is decreasing to a value that is below the value of the historical performance information (column 308), also as shown by reference 312. For example, the historical performance information (column 308) indicates a value of 23 Mbps while the overall

performance (column 306) decreased to 22 Mbps (at time = 70). If the current configuration were to continue, then it may be that the replication process may not complete successfully resulting in no copies of the source data being made to target computing devices. In this case, data replication management module 104 determines that it is time to select a target computing device 112 from plurality of target computing devices 112-1, 112-2, 112-3, 112-4 to complete the replication process.

[00050] In this case, data replication management module 104 determines that the progress information indicates that the data replication process is not to be completed within the predefined time period, then the management module selects a target computing device 112 from plurality of target computing devices 112-1, 112-2, 112-3, 112-4 to complete the complete the data replication process based on historical and current performance information of the plurality of target computing devices. In this case, management module 104 selects the third target computing device 112-3 since it has the highest performance as shown by data throughput of column 304-3 (value of 39 Mbps) of table 300 compared to those of the other target computing devices 112-1, 112-2, 112-4. The management module 104 sends a continue notification to the selected target computing device 112-3 to cause the selected target computing device to continue with the completion of the data replication process, as shown by reference 310 of table 300. In one example, management module 104 may be configured to release processing resources used by the computing device 102 associated with the non-selected target computing devices 112-1, 112-2 and 112-4 for use by the computing device for the data replication process of the selected target computing device 112-3. In this manner, releasing resources may allow the computing device to use the resources to complete replication process with the selected target computing device. In this case, the release of resources may allow computing device 102 to focus on the replication process directed to the selected target computing device 112-3 to increase performance (data throughput) from 39 Mbps (at time =70) to 60 Mbps (at time = 165) as shown by reference 310 at table 300. In addition, the system takes some time to adjust the overall performance (data throughput) shown by reference 312 such that thereafter performance increases from 45 Mbps (at time =105) to 60 Mbps (at time = 165) as shown by column 306 at table 300.

[00051] In this case, management module 104 send a discontinue notification to non-selected target computing devices 112-1, 112-2, 112-4 to cause the non-selected target computing devices to discontinue with the data replication process. This is illustrated by the "pause" indication in respective columns 304-1, 304-2, 304-4 of table 300 associated with respective computing devices 112-1, 112-2, 112-4.

[00052] In one example, data replication management module 104 may cause the selected target computing device 112-3 to initiate a data replication process to copy the replicated target data set to another target computing device of the plurality of target computing devices. For example, after target computing device 112-3 completes the data replication process, it may initiate a data replication process with target computing devices 11-2.

[00053] In this manner, the present application discloses techniques to provide data replication management techniques which may help improve data replication from a source computing device to across a plurality of target computing devices which may increase the performance of the system. For example, the data replication techniques such that even during resource constraint at least one copy of data is made to help ensure data availability.

[00054] Turning to Fig. 3B, graph 320 illustrates an example process of the data replication process of Fig. 3A. The graph 320 shows time on the x axis (horizontal) in time intervals and data throughput on the y axis (vertical) in Mbps. The graph 320 shows that during the initial time period, data replication management module 104 is performing a data replication process that includes all of the target computing devices 112-1, 112-2, 112-3, 112-4 as shown by reference 322. In this case, initially, the system may be said to be operating in a multiple data replication mode such that all targeting computing devices are configured to perform the data replication process, as shown by period 322 of graph 320.

[00055] However, during the next time interval, data replication management module 104 determines that the current performance is decreasing below the historical performance information during the interval (data replication window time) as shown by reference 324 of graph 320. In particular, the current performance is shown by reference 328 while the historical performance is shown as a value of 23 Mbps by reference 330 at graph 320.

[00056] In this case, data replication management module 104 selects target computing device 112-3 since it has the highest performance as represented by the highest data throughput. The management module 104 sends a continue notification to the selected target computing device 112-3 to cause the selected target computing device to continue with the completion of the data replication process, as shown by reference 326 of graph 320. In one example, the data replication management module selects the target computing device that has highest performance compared to those of the other target computing devices. It should be understood that the data replication management module

may select the target computing device based on other techniques. In this case, the system is said to be operating in a single data replication mode such that only one target computing device is performing the replication process. The management module 104 sends a discontinue notification to non-selected target computing devices 112-1, 112-2, 112-4 to cause the non-selected target computing devices to discontinue with the data replication process. This allows the current performance or data throughput of the selected target computing device 112-3 for the data replication process to increase after the selection process as shown by reference 328 of graph 320. The current performance shown by reference 328 is higher and increases over the historical performance after the selection process as shown by reference 330 of graph 320.

[00057] In this manner, the present application discloses techniques to provide data replication management techniques which may help improve data replication from a source computing device to across a plurality of target computing devices which may increase the performance of the system. For example, the data replication techniques such that even during resource constraint at least one copy of data is made to help ensure data availability.

[00058] The process of Figs. 3A and 3B shows example processes and it should be understood that other configurations can be employed to practice the techniques of the present application. For example, processes above may be configured to communicate with a plurality of computing devices communicating with target computing devices.

[00059] Fig. 4 is an example block diagram showing a non-transitory, computer-readable medium that stores code for operation in accordance with an example of the techniques of the present application. The non-transitory, computer-readable medium is generally referred to by the reference number 400 and may be included in the system in relation to Fig. 1. The non-transitory, computer-readable medium 400 may correspond to any typical storage device that stores computer-implemented instructions, such as programming code or the like. For example, the non-transitory, computer-readable medium 400 may include one or more of a non-volatile memory, a volatile memory, and/or one or more storage devices. Examples of non-volatile memory include, but are not limited to, electrically erasable programmable Read Only Memory (EEPROM) and Read Only Memory (ROM). Examples of volatile memory include, but are not limited to, Static Random Access Memory (SRAM), and dynamic Random Access Memory (DRAM). Examples of storage devices include, but are not limited to, hard disk drives, compact disc drives, digital versatile disc drives, optical drives, and flash memory devices.

[00060] A processor 402 generally retrieves and executes the instructions stored in the non-transitory, computer-readable medium 400 to operate the present techniques in accordance with an example. In one example, the tangible, computer-readable medium 400 can be accessed by the processor 402 over a bus 404. A first region 406 of the non-transitory, computer-readable medium 400 may include data replication management module 104 functionality as described herein. The data replication management module 104 functionality may be implemented in hardware, software or a combination thereof.

[00061] For example, block 408 provides instructions which may include instructions to initiate a data replication process, as described herein. In one example, the instructions may include instructions to initiate a data replication process to copy the source data set as target data sets to a plurality of target computing devices during a predefined time period, as described herein.

[00062] For example, block 410 provides instructions which may include instructions to determine progress information, as described herein. In one example, the instructions may include instructions to determine progress information regarding completion of the data replication process based on historical and current performance information of the plurality of target computing devices, as described herein.

[00063] For example, block 412 provides instructions which may include instructions to select a target computing, as described herein. In one example, the instructions may include instructions to select a target computing device from the plurality of target computing devices to complete the data replication process based on historical and current performance information of the plurality of target computing devices, as described herein.

[00064] For example, block 414 provides instructions which may include instructions to cause the selected target computing device to continue, as described herein. In one example, the instructions may include instructions to cause the selected target computing device to continue with the completion of the data replication process, as described herein.

[00065] For example, block 416 provides instructions which may include instructions to cause the non-selected target computing devices to discontinue, as described herein. In one example, the instructions may include instructions to cause the non-selected target computing devices to discontinue with the data replication process, as described herein.

[00066] For example, block 418 provides instructions which may include instructions to cause the selected device to copy data, as described herein. In one

example, the instructions may include instructions to cause the replicated target data set to be copied to another target computing device of the plurality of target computing devices, as described herein.

[00067] Although shown as contiguous blocks, the software components can be stored in any order or configuration. For example, if the non-transitory, computer-readable medium 400 is a hard drive, the software components can be stored in non-contiguous, or even overlapping, sectors.

[00068] As used herein, a “processor” may include processor resources such as at least one of a Central Processing Unit (CPU), a semiconductor-based microprocessor, a Graphics Processing Unit (GPU), a Field-Programmable Gate Array (FPGA) configured to retrieve and execute instructions, other electronic circuitry suitable for the retrieval and execution instructions stored on a computer-readable medium, or a combination thereof. The processor fetches, decodes, and executes instructions stored on medium 400 to perform the functionalities described below. In other examples, the functionalities of any of the instructions of medium 400 may be implemented in the form of electronic circuitry, in the form of executable instructions encoded on a computer-readable storage medium, or a combination thereof.

[00069] As used herein, a “computer-readable medium” may be any electronic, magnetic, optical, or other physical storage apparatus to contain or store information such as executable instructions, data, and the like. For example, any computer-readable storage medium described herein may be any of Random Access Memory (RAM), volatile memory, non-volatile memory, flash memory, a storage drive (e.g., a hard drive), a solid state drive, any type of storage disc (e.g., a compact disc, a DVD, etc.), and the like, or a combination thereof. Further, any computer-readable medium described herein may be non-transitory. In examples described herein, a computer-readable medium or media is part of an article (or article of manufacture). An article or article of manufacture may refer to any manufactured single component or multiple components. The medium may be located either in the system executing the computer-readable instructions, or remote from but accessible to the system (e.g., via a computer network) for execution. In the example of Fig. 4, medium 400 may be implemented by one computer-readable medium, or multiple computer-readable media.

[00070] In examples described herein, computing device 102 may communicate with components implemented on separate devices or system(s) via a network interface device of the computing device. For example, computing device 102 may communicate with storage device 106 via a network interface device of the computing device 102. In

another example, computing device 102 may communicate with target computing devices 112 via a network interface device of the computing device 102. In examples described herein, a "network interface device" may be a hardware device to communicate over at least one computer network. In some examples, a network interface may be a Network Interface Card (NIC) or the like. As used herein, a computer network may include, for example, a Local Area Network (LAN), a Wireless Local Area Network (WLAN), a Virtual Private Network (VPN), the Internet, or the like, or a combination thereof. In some examples, a computer network may include a telephone network (e.g., a cellular telephone network).

[00071] In some examples, instructions 408-418 may be part of an installation package that, when installed, may be executed by processor 402 to implement the functionalities described herein in relation to instructions 408-418. In such examples, medium 400 may be a portable medium, such as a CD, DVD, or flash drive, or a memory maintained by a server from which the installation package can be downloaded and installed. In other examples, instructions 408-418 may be part of an application, applications, or component(s) already installed on computing device 102 including processor 402. In such examples, the medium 400 may include memory such as a hard drive, solid state drive, or the like. In some examples, functionalities described herein in relation to Figs. 1 through 4 may be provided in combination with functionalities described herein in relation to any of Figs. 1 through 4.

[00072] The foregoing describes a novel and previously unforeseen approach for storage management. While the above application has been shown and described with reference to the foregoing examples, it should be understood that other forms, details, and implementations may be made without departing from the spirit and scope of this application.

WHAT IS CLAIMED IS:

1. A computing device comprising:
 - a storage device to store a source data set; and
 - a data replication management module to:
 - initiate a data replication process to copy the source data set as target data sets to a plurality of target computing devices during a predefined time period,
 - determine progress information regarding completion of the data replication process based on historical and current performance information of the plurality of target computing devices,
 - if the progress information indicates that the data replication process is not to be completed within the predefined time period, then
 - select a target computing device from the plurality of target computing devices to complete the complete the data replication process based on historical and current performance information of the plurality of target computing devices, and
 - send a continue notification to the selected target computing device to cause the selected target computing device to continue with the completion of the data replication process,
 - send a discontinue notification to non-selected target computing devices to cause the non-selected target computing devices to discontinue with the data replication process, and
 - cause the selected target computing device to initiate a data replication process to copy the replicated target data set to another target computing device of the plurality of target computing devices.
2. The computing device of claim 1, wherein if the progress information indicates that the data replication process is to be completed within the predefined time period, then send continue notifications to the target computing devices to cause the target computing device to continue with the completion of the data replication process.
3. The computing device of 1, wherein the data replication management module to calculate the current performance information of a target computing device which relates to data throughput during the data replication process and calculate the historical performance information of the target computing device which relates to data throughput during previous data replication processes.

4. The computing device of claim 1, wherein the data replication management module to release processing resources used by the computing device associated with the non-selected target computing devices for use by the computing device for the data replication process of the selected target computing device.

5. The computing device of claim 1, wherein the data replication management module to receive from a host computer a request to initiate the data replication process.

6. A method of data replication management, the method comprising:

initiating a data replication process to copy source data set as target data sets to a plurality of target computing devices during a predefined time period;

determining progress information regarding completion of the data replication process based on historical and current performance information of the plurality of target computing devices; and

if the progress information indicates that the data replication process is not to be completed within a predefined time period, then

selecting a target computing device from the plurality of target computing devices to complete the complete the data replication process based on historical and current performance information of the plurality of target computing devices, and

causing the selected target computing device to continue with the completion of the data replication process, causing the non-selected target computing devices to discontinue with the data replication process, and causing the replicated target data set to be copied to another target computing device of the plurality of target computing devices.

7. The method of claim 6, wherein if the progress information indicates that the data replication process is to be completed within the predefined time period, then sending continue notifications to the target computing devices to cause the target computing device to continue with the completion of the data replication process.

8. The method of claim 6, further comprising calculating the current performance information of a target computing device which relates to data throughput during the data replication process and calculating the historical performance information of the target

computing device which relates to data throughput during previous data replication processes.

9. The method of claim 6, further comprising releasing processing resources used by the computing device associated with the non-selected target computing devices for use by the computing device for the data replication process of the selected target computing device.

10. The method of claim 6, further comprising receiving from a host computer a request to initiate the data replication process.

11. A non-transitory computer-readable medium having computer executable instructions stored thereon for data replication management, the instructions are executable by a processor to:

- initiate a data replication process to copy source data set as target data sets to a plurality of target computing devices during a predefined time period;

- determine progress information regarding completion of the data replication process based on historical and current performance information of the plurality of target computing devices; and

- if the progress information indicates that the data replication process is not to be completed within a predefined time period, then

- select a target computing device from the plurality of target computing devices to complete the complete the data replication process based on historical and current performance information of the plurality of target computing devices, and

- cause the selected target computing device to continue with the completion of the data replication process, cause the replicated target data set to be copied to another target computing device of the plurality of target computing devices, and cause the non-selected target computing devices to discontinue with the data replication process.

12. The non-transitory computer-readable medium of claim 11, further comprising instructions that if executed cause a processor to: send a discontinue notification to non-selected target computing devices to cause the non-selected target computing devices to discontinue with the data replication process.

13. The non-transitory computer-readable medium of claim 11, further comprising instructions that if executed cause a processor to: calculate the current performance information of a target computing device which relates to data throughput during the data replication process and calculate the historical performance information of the target computing device which relates to data throughput during previous data replication processes.

14. The non-transitory computer-readable medium of claim 11 further comprising instructions that if executed cause a processor to: release processing resources used by the computing device associated with the non-selected target computing devices for use by the computing device for the data replication process of the selected target computing device.

15. The non-transitory computer-readable medium of claim 11 further comprising instructions that if executed cause a processor to: receive from a host computer a request to initiate the data replication process.





FIG. 1



3/5

302	304-1	304-2	304-3	304-4	306	308
POLL DURING PERIOD	TARGET 1	TARGET 2	TARGET 3	TARGET 4	OVERALL THROUGHPUT FOR ALL REPLICATION (MB/SEC)	HISTORICAL AVERAGE
5	20	25	40	18	26	23
10	21	22	43	19	26	23
15	20	35	45	33	23	23
20	18	32	43	21	29	23
25	18	22	42	18	25	23
30	17.3	28.4	42	22.4	28	23
35	16.6	28.8	39	22.6	27	23
40	15.9	29.2	40	22.8	27	23
45	15.2	29.6	43	23	28	23
50	14.5	30	45	23.2	28	23
55	13.8	30.4	43	23.4	28	23
60	13.1	30.8	42	23.6	27	23
65	12.4	31.2	42	23.8	27	23
70	11.7	11.7	39	24	22	23
75	11	10	42	18	20	23
80	7.7	11	39	19	19	23
85	9.6	17	38	20	21	23
90	8.9	12	40	21	20	23
95	8.2	13	41	24	22	23
100	7.5	13	42	21	21	23
105	pause	pause	45	pause	45	23
110	pause	pause	49	pause	49	23
115	pause	pause	50	pause	50	23
120	pause	pause	53	pause	53	23
125	pause	pause	55	pause	55	23
130	pause	pause	56	pause	56	23
135	pause	pause	60	pause	60	23
140	pause	pause	55	pause	55	23
145	pause	pause	50	pause	50	23
150	pause	pause	53	pause	53	23
155	pause	pause	55	pause	55	23
160	pause	pause	56	pause	56	23
165	pause	pause	60	pause	60	23

FIG. 3A




FIG. 3B

5/5




FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2015/050867**A. CLASSIFICATION OF SUBJECT MATTER****G06F 11/14(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)

G06F 11/14; G06F 12/16; G06F 17/00; G06F 17/30

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

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: data replication, performance, completion, network, target, instruction, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2012-0284229 A1 (MI-JEOM KIM et al.) 08 November 2012 See paragraphs [0042] and [0049]; claim 1; and figure 2.	1-15
A	US 8,458,418 B1 (PRATAP SINGH et al.) 04 June 2013 See column 2, lines 54-58; claim 1; and figure 1.	1-15
A	US 8,126,844 B2 (ROBERT K.KAST) 28 February 2012 See column 3, lines 9-16; claim 1; and figure 2.	1-15
A	US 2012-0233419 A1 (TATSURO YAMASHITA) 13 September 2012 See paragraphs [0084]-[0086]; claim 1; and figure 1.	1-15
A	EP 2474919 A1 (CHINA UNIONPAY CO., LTD.) 11 July 2012 See paragraphs [0021]-[0023]; claim 1; and figure 1.	1-15



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 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

29 February 2016 (29.02.2016)

Date of mailing of the international search report

29 February 2016 (29.02.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/050867

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012-0284229 A1	08/11/2012	KR 10-1544483 B1 KR 10-2012-0116774 A US 8849756 B2 WO 2012-141404 A1	17/08/2015 23/10/2012 30/09/2014 18/10/2012
US 8458418 B1	04/06/2013	None	
US 8126844 B2	28/02/2012	US 2011-0093855 A1 US 7882061 B1	21/04/2011 01/02/2011
US 2012-0233419 A1	13/09/2012	WO 2012-120667 A1	13/09/2012
EP 2474919 A1	11/07/2012	CN 102004743 A CN 102004743 B EP 2474919 A4 US 2012-0246116 A1 US 8527459 B2 WO 2011-026313 A1	06/04/2011 14/08/2013 30/04/2014 27/09/2012 03/09/2013 10/03/2011