




(86) Date de dépôt PCT/PCT Filing Date: 2005/11/08
(87) Date publication PCT/PCT Publication Date: 2006/05/18
(45) Date de délivrance/Issue Date: 2013/12/17
(85) Entrée phase nationale/National Entry: 2007/05/07
(86) N° demande PCT/PCT Application No.: US 2005/040669
(87) N° publication PCT/PCT Publication No.: 2006/053093
(30) Priorité/Priority: 2004/11/08 (US60/625,894)

(51) Cl.Int./Int.Cl. *G06F 15/16* (2006.01)
(72) Inventeur/Inventor:
JACKSON, DAVID BRIAN, US
(73) Propriétaire/Owner:
ADAPTIVE COMPUTING ENTERPRISES, INC., US
(74) Agent: KIRBY EADES GALE BAKER

(54) Titre : SYSTEME ET PROCEDE FOURNISSANT DES EXECUTIONS DE SYSTEME AU SEIN D'UN ENVIRONNEMENT INFORMATIQUE
(54) Title: SYSTEM AND METHOD OF PROVIDING SYSTEM JOBS WITHIN A COMPUTE ENVIRONMENT



(57) Abrégé/Abstract:

The invention relates to systems, methods and computer-readable media for using system jobs for performing actions outside the constraints of batch compute jobs submitted to a compute environment such as a cluster or a grid. The method for modifying a



(57) **Abrégé(suite)/Abstract(continued):**

compute environment from a system job comprises associating a system job to a queueable object, triggering the system job based on an event and performing arbitrary actions on resources outside of compute nodes in the compute environment. The queueable objects include objects such as batch compute jobs or job reservations. The events that trigger the system job may be time driven, such as ten minutes prior to completion of the batch compute job, or dependent on other actions associated with other system jobs. The system jobs may be utilized also to perform rolling maintenance on a node by node basis.

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

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
18 May 2006 (18.05.2006)

PCT

(10) International Publication Number
WO 2006/053093 A3

(51) International Patent Classification:

G06F 15/16 (2006.01)

(21) International Application Number:

PCT/US2005/040669

(22) International Filing Date:

8 November 2005 (08.11.2005)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

60/625,894 8 November 2004 (08.11.2004) US

(71) Applicant (for all designated States except US): **CLUSTER RESOURCES, INC.** [US/US]; 622 North 900 East, Spanish Fork, UT 84660 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **JACKSON, David, Brian** [US/US]; 10890 S. Conestoga Drive, Spanish Fork, UT 84660 (US).

(74) Agent: **ISAACSON, Thomas, M.**; Law Office of Thomas M. Isaacson, 850 Lindy Lane, Huntingtown, MD 20639 (US).

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

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

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

Declarations under Rule 4.17:


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

Published:

- with international search report

[Continued on next page]

(54) Title: SYSTEM AND METHOD OF PROVIDING SYSTEM JOBS WITHIN A COMPUTE ENVIRONMENT



(57) Abstract: The invention relates to systems, methods and computer-readable media for using system jobs for performing actions outside the constraints of batch compute jobs submitted to a compute environment such as a cluster or a grid. The method for modifying a compute environment from a system job comprises associating a system job to a queueable object, triggering the system job based on an event and performing arbitrary actions on resources outside of compute nodes in the compute environment. The queueable objects include objects such as batch compute jobs or job reservations. The events that trigger the system job may be time driven, such as ten minutes prior to completion of the batch compute job, or dependent on other actions associated with other system jobs. The system jobs may be utilized also to perform rolling maintenance on a node by node basis.

WO 2006/053093 A3

WO 2006/053093 A3



— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

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

(88) Date of publication of the international search report:

16 November 2006

SYSTEM AND METHOD OF PROVIDING SYSTEM JOBS WITHIN A COMPUTE ENVIRONMENT

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0002] The present invention relates to managing a compute environment or more specifically to a system and method of modifying or updating a compute environment using system jobs. One embodiment of the invention relates to rolling maintenance on a node-by-node basis within the compute environment.

2. Introduction

[0003] The present invention relates to a system and method of managing resources in the context of a compute environment which may be defined as a grid or cluster of computers. Grid computing may be defined as coordinated resource sharing and problem solving in dynamic, multi-institutional collaborations. Many computing projects require much more computational power and resources than a single computer or computer processor may provide. Networked computers with peripheral resources such as printers, scanners, I/O devices, storage disks, scientific devices and instruments, etc. may need to be coordinated and utilized to complete a task.

[0004] Grid/cluster resource management generally describes the process of identifying requirements, matching resources to applications, allocating those resources, and scheduling and monitoring compute resources over time in order to run applications or compute jobs as efficiently as possible. Each project will utilize a different set of

resources and thus is typically unique. In addition to the challenge of allocating resources for a particular job, administrators also have difficulty obtaining a clear understanding of the resources available, the current status of the environment and available resources, and real-time competing needs of various users. General background information on clusters and grids may be found in several publications. See, e.g., Grid Resource Management, State of the Art and Future Trends, Jarek Nabrzyski, Jennifer M. Schopf, and Jan Weglarz, Kluwer Academic Publishers, 2004; and Beowulf Cluster Computing with Linux, edited by William Gropp, Ewing Lusk, and Thomas Sterling, Massachusetts Institute of Technology, 2003.

[0005] It is generally understood herein that the terms grid and cluster are interchangeable in that there is no specific definition of either. In general, a grid will comprise a plurality of clusters as will be shown in FIG. 1. Several general challenges exist when attempting to maximize resources in a grid. First, there are typically multiple layers of grid and cluster schedulers. A grid 100 generally comprises a group of clusters or a group of networked computers. The definition of a grid is very flexible and may mean a number of different configurations of computers. The introduction here is meant to be general given the variety of configurations that are possible. A grid scheduler 102 communicates with a plurality of cluster schedulers 104A, 104B and 104C. Each of these cluster schedulers communicates with a plurality of resource managers 106A, 106B and 106C. Each resource manager communicates with a series of compute resources shown as nodes 108A, 108B, 108C, 108D, 108E, 108F, 108G, 108H, 108I.

[0006] Local schedulers (which may refer to the cluster schedulers 104A, 104B, 104C or the resource managers 106A, 106B, 106C) are closer to the specific resources 108 and may not allow grid schedulers 102 direct access to the resources. The resources are grouped into clusters 110, 112 and 114. Examples of cluster resources include data storage devices such as hard drives, compute resources such as computer processors,

network resources such as routers and transmission means, and so forth. The grid level scheduler 102 typically does not own or control the actual resources. Therefore, compute jobs are submitted from the high level grid-scheduler 102 to a local set of resources with no more permissions than then user would have. Compute jobs may also be submitted at the cluster scheduler layer of the grid or even directly at the resource managers. There are problems with the efficiency of the arrangement.

[0007] The heterogeneous nature of the shared resources causes a reduction in efficiency. Without dedicated access to a resource, the grid level scheduler 102 is challenged with the high degree of variance and unpredictability in the capacity of the resources available for use. Most resources are shared among users and projects and each project varies from the other. The difference in performance goals for various projects also reduces efficiencies. Grid resources are used to improve performance of an application but the resource owners and users have different performance goals: from optimizing the performance for a single application to getting the best system throughput or minimizing response time. Local policies may also play a role in performance.

[0008] FIG. 2 illustrates a current state of art that allows a scheduler/resource manager combination to submit and control standard batch compute jobs. An example of a batch job is a request from a weather service to process a hurricane analysis. The amount of computing resources are large and therefore the job is submitted to a cluster for processing. A batch job is submitted to the queue of a resource manager and is constrained to run within the cluster associated with that resource manager. A batch job 204, 206 or 208 within a queue 202 has the ability to have a number of steps in which each step may have dependencies on other steps, successful or failed completion of previous steps or similar relationships. The bounds of influence for the batch jobs are limited to running non-root applications or executables on that cluster or on compute nodes that are allocated to it.

[0009] The respective batch job is unable to do anything outside of the constrained space for the job. There are a number of deficiencies with this approach, particularly in that such a job is unable to modify the scheduling environment. The job is only able to operate within the scheduling environment and it is also constrained to only doing the specified actions. For example, the job may be constrained to run an executable within a compute node of the cluster (within its allocated space), but it is unable to run any other action within the cluster or within the other services of the cluster.

[0010] What is needed is a method by which a processing entity that can be queued can be submitted to the scheduler that can be more flexible and have a broader scope of impact on the compute environment.

SUMMARY OF THE INVENTION

[0011] Additional features and advantages of the invention will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by practice of the invention. The features and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other features of the present invention will become more fully apparent from the following description and appended claims, or may be learned by the practice of the invention as set forth herein.

[0012] The invention relates to systems, methods and computer-readable media for using system jobs for performing actions outside the constraints of batch compute jobs. System jobs may be conceptually thought of as intelligent agents. Typically, controlling and managing the resources within a compute environment such as a cluster or a grid are tasks performed by a scheduler or other management software. No actions, provisioning or reservations are made outside of the control of this software. The present invention

provides increased flexibility in managing and controlling the environment by using entities that are called system jobs that may include triggered events that are outside of events managed by a scheduler. Other concepts that apply to system jobs include state, retry capability, steps, time steps and dependencies.

[0012a] Certain exemplary embodiments can provide a method for managing a multi-node compute environment, the method comprising: associating a system job to a compute job that is scheduled according to a reservation to consume resources in a multi-node compute environment under common management, wherein the system job is initiated independent of a resource manager that establishes the reservation for the compute job in the multi-node compute environment; and performing an action on a resource within the multi-node compute environment by processing the system job.

[0012b] Certain exemplary embodiments can provide a system for managing a multi-node compute environment, the system comprising: a processor; and a computer-readable storage medium storing instructions, which, when performed by the processor, cause the processor to perform a method comprising: associating a system job to a compute job that is scheduled according to a reservation to consume resources in a multi-node compute environment, wherein the system job is initiated independent of a resource manager that establishes the reservation for the compute job in the multi-node compute environment; and performing an action on a resource within the multi-node compute environment by processing the system job.

[0012c] Certain exemplary embodiments can provide a non-transitory computer-readable storage medium storing instructions which, when executed by a computing device, cause the computing device to manage a multi-node compute environment, the instructions comprising: associating a system job to a compute job that is scheduled according to a reservation to consume resources in a multi-node compute environment, wherein the system job is initiated independent of a resource manager that establishes the reservation for the compute job in the multi-node compute environment; and performing an action on a resource within the multi-node compute environment by processing the system job.

[0013] As an example, the method embodiment of the invention comprises a method for modifying a compute environment from a system job by associating a system job to a queueable object, triggering the system job based on an event and performing arbitrary actions on resources in the compute environment. The queueable objects include objects such as batch compute jobs or job reservations. The events that trigger the system job may be time driven, such as ten minutes prior to completion of the batch compute job, or dependent on other actions associated with other system jobs.

[0014] Another embodiment of the invention relates to performing a rolling maintenance on a compute environment. A method of performing rolling maintenance on a node within a compute environment comprises receiving a submission of a system job associated with a node, performing a provisioning operation on the node, determining whether the provisioning was successful (health check) and if provisioning was successful, then terminating the system

job leaving the node available for use in the compute environment. If the provisioning was not successful, the system job reports an unsuccessful status via means such as an email to an administrator and creates a reservation for the node. Provisioning a node may involve updating software, provisioning an operating system or any other operation that may be performed on that node. The operation on each node is performed on a node by node basis independently and a time associated with the process may be an earliest possible time, a scheduled time or an earliest possible time after a predetermined period of time. System jobs may be submitted at both a grid level and a cluster level within a compute environment.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] In order to describe the manner in which the above-recited and other advantages and features of the invention can be obtained, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0016] FIG. 1 illustrates generally a grid scheduler, cluster scheduler, and resource managers interacting with compute nodes;

[0017] FIG. 2 illustrates a prior art submission of a batch job for processing in a compute environment;

[0018] FIG. 3 illustrates the use of system jobs in an architecture according to an aspect of the invention;

[0019] FIG. 4 illustrates the use of system jobs in a cluster of nodes;

[0020] FIG. 5 illustrates a method embodiment of the invention;

[0021] FIG. 6 is a flowchart illustrating an example embodiment of the invention; and

[0022] FIG. 7 illustrates a cluster with files system spanning multiple nodes.

DETAILED DESCRIPTION OF THE INVENTION

[0023] The present invention provides an improvement over the prior art by enabling system jobs or other processing entities that may be queued for processing in a compute environment to have the ability to perform arbitrary actions on resources outside the compute nodes in the environment. Furthermore, the invention enables actions to be taken associated with the submitted job outside the previously constrained space.

[0024] Embodiments of the invention relate to system jobs, and systems of creating and using system jobs, methods of creating and using system jobs, computer-readable media for controlling a computing device to manage system jobs and a compute environment operating according to the principles of the invention. As introduced above, one example of a job is a consume job that consumes resources for a particular project, such as a weather study. The present invention provides for a different type of job that is flexible and performs other operations and/ or modifications in the compute environment. System jobs may be created and/ or submitted remotely or internally within a compute environment and may spawn child operations into a resource manager but the master job resides strictly within the workload manager and/ or scheduler. System jobs will preferably contain one or more steps with dependencies.

[0025] Each step that is involved in processing a system job may consist of one or more tasks where each task modifies the internal or external environment of the compute environment or the job. Internal environment changes include, but are not limited to: creating reservations, setting variables, modifying credentials, policies, thresholds,

priorities, etc. External changes include modifying resources, database settings, peer interfaces, external credentials, launching arbitrary scripts, launching applications, provisioning resources, etc.

[0026] A system job may require several steps to complete its process and terminate.

Throughout this process, at various stages, a state of a particular task needs to be identified. Step state is based on success or failure of task execution. Steps may possess triggers. Steps may generate and consume job level and global level variables. Step dependencies may be based on internal or external factors including, but not limited to: job, step, trigger, time, or environment based dependencies. Time dependencies may be based on absolute time, or time relative to some job internal or external event.

Dependencies may include local or global variable settings. Dependencies may be based on return value of arbitrary configurable probes.

[0027] Steps may optionally allocate resources. Steps may optionally be associated with a walltime. There are several differentiators associated with system jobs. They allow at least one of: (1) integration of environmental data into job flow decisions; (2) creation of arbitrary probes, continuous task retry, etc.; (3) integration of environment data into task execution; (4) dynamic resource reallocation based results of previous tasks; (5) integration of compute tasks, tasks involving non-compute resources (i.e. data bases, provisioning systems, data managers, etc), and changes to compute environment meta data (such as policies, thresholds, priorities, credential configuration, etc); (6) access to live global cluster and job centric information, information; (7) envelopment of traditional compute tasks in higher layer wrappers; (8) allowing greater environment management; (8) synchronization of tasks managing unrelated resources and resource types; (9) co-allocation of resources and requirements, scheduling, reservation; (10) guarantees of completion for loose aggregations of request types application of tight and

loose time constraints on requests (including periodic window, timeframe proximity, and deadline based constraints); and (11) optimization of loose aggregations of requests.

[0028] System jobs are also referred to as workload management object event policies.

The purpose of a workload management object event policy is to allow actions to be associated with a workload management object such as a reservation, a compute/system job, a node, a cluster, a user, a resource manger and/or other queue-able workload units that trigger a given action either based on a time criteria or other measurable condition.

An example of this may be a system/compute job having an associated event policy that launches a script 10 minutes prior to a jobs completion. This script could send an e-mail to the user notifying them that the job is almost finished, or it can set in action the launch of another job that has a dependency on the results of the initial job being mostly complete. Another example is that of a reservation with an associated event policy that deletes temporary files and restarts all of the reserved nodes to purge them of sensitive data and to clear memory prior to usage by another entity.

[0029] An example of the method aspect of the invention comprises the steps of receiving a request for the creation of an entity to manage or perform at least one operation within a compute environment. The entity is preferably a system job as described herein. The method further comprises creating the entity, wherein the entity has arbitrary dependencies, associating the entity with a workload management object and using the entity to perform at least one operation and/or modification on the compute environment.

[0030] FIG. 3 illustrates an example of how a system job 326 can be used to set up a virtual private cluster or a job-specific virtual cluster. In FIG. 3, the user 312 submits a job 326 via a queue 302 to a resource manager 106. A queue 318 is also shown as having jobs submitted to the scheduler 104. The queue 310 illustrates in more detail the invention with a compute job and system jobs associated with it that will be processed on

the cluster 110. While the job 326 is submitted by the user 312, the associated system jobs may be selected by the user 312 or via an automatic process that receives some input from the user 312 and also may reference policy information or service level agreement information to generate system jobs to help to monitor and manager the compute environment for the submitted job 326.

[0031] The job steps discussed and the functions performed that are associated with the job may be arbitrary in the principles of the present invention. The concrete examples illustrate how the arbitrary capabilities may be applied. A queue 310 holds a system job 326 and a number of other job steps 320, 322, 324, 328. The first job step 320 involves contacting not the cluster but a provisioning manager 330 to set up a compute environment. The subsequent job step 322 arranges for storage management with a storage manager 332; the third job step 324 contacts a license manager 334 to make sure the applications that are needed are available. The fourth step 326 executes the actual job in the virtual environment 308 and the final step 328 involves staging the data out of this environment and destroying or collapsing the virtual cluster.

[0032] The above example illustrates the operation of system jobs where there could be any combination of the various tasks associated with a system job. System jobs have a number of distinct differences from standard consume jobs 326. The present invention allows for full support meaning that jobs allow arbitrary dependencies and combinations or relationships between job steps. They also allow arbitrary actions in which arbitrary things can be executed, arbitrary services can be driven, arbitrary data can be modified, arbitrary policies and configurations of the scheduler can be adjusted. They can be set to require resource allocation and can be set up so they only come live when those resources can be allocated and dedicated to the system job. They also have the ability to have arbitrary impact on the system.

[0033] Figure 4 shows an example of using a system job to perform a rolling maintenance. Rolling maintenance can include updating its software, performing rolling provisioning, patches and software upgrades as well as other functions. In a rolling maintenance, a site has a desire to either check or change current applications, operating systems or kernel versions in their compute nodes or other cluster resources. For example, assume that a compute node needs to have software reinstalled and updated. Previously, this process would be done by taking the entire node down after all the jobs assigned to that node are complete, making the system unavailable, installing by hand all the nodes with the new level of software and once checks are made turning all nodes back to the users to continue running jobs. This process is made more efficient by the application of system jobs.

[0034] Figure 4 illustrates a series of nodes 402 with the associated with resource manager 106, scheduler 104 and provisioning manager 330. Using system jobs, a system administrator, rather than performing all the above-mentioned steps, simply submits a system job which performs the update automatically. For example, the system job schedules at the earliest possible time on each node an independent node update, a software update and in addition to updating the node, it also performs a sanity and/or health check. In event of failure, the system job notifies the administrator so that he or she should take action as needed on the nodes that actually failed. This reduces the human administration time required in any update or modification.

[0035] Cluster 402 of FIG. 4 illustrates a series of jobs 1 - 6 running some of the nodes 1 - 5 with time along the X axis. As shown, node 1 is currently running job 1 and in some time in the future, job 1 will complete and a system job 1 will operate for some time, followed by job 5. Some of these nodes are currently empty, namely node 4 which is running system job 4. When the administrator actually schedules the system job, the system preferably identifies the earliest time that the job could occur on each node.

The system job may also be modified to identify any particular time to begin, i.e., it may be instructed to find the earliest time starting one week from today, an earliest possible time from any predetermined time or a scheduled time. For example, on node 4 the job can start immediately, which it does, and then update that node and turns it over to run job 4 which automatically happens as soon as it completes its health and sanity check.

[0036] On other nodes the system job is scheduled for immediate processing upon completion of existing workloads. The update is completed as soon as possible and the node is again automatically turned over to user access and jobs (shown as job 6) can begin or continue to run. The system jobs principle takes advantage of the fact that the system jobs are actually not running out on the compute host (the cluster). When a system job requires allocation of a resource such as node 1, as soon as node 1 is available, the job launches a request to the provisioning service 330. The provisioning service 330 then updates the node as necessary to handle the job. As soon as that step of the system job is complete, a health check trigger is launched verifying the node is operational. If the health check trigger is successful, the node is freed and the system job is canceled. If the health check is unsuccessful, an e-mail is sent out and the node is reserved indefinitely. The e-mail is sent to the administrator so he or she can correct whatever problems occurred. In a similar case, in all cases the system job is not actually run on the compute host even though the compute host is allocated and impacted by the system job.

[0037] FIG. 5 illustrates the method aspect of the invention related to the use of a system job required for maintenance. The method comprises a number of steps performed by the system job. The first step comprises the system job to transmit a communication to the provisioning manager to provision an allocated resource (502). Each system job will have a requirement for a specific node. For example, in the example shown in FIG. 5, the system job requires that the system job only runs with

regard to node 1 because it requires node 1. The job is not available to start until the node is allocated and dedicated to this job. Once that job runs, it uses the provisioning manager to provision a particular operating system (or for some other provisioning need) that has been requested.

[0038] Next, the method comprises running a script that communicates with the node to verify that the provisioning step was properly carried out and that the node is healthy (504). If step 504 reports success (506), then the system job sends an e-mail and terminates the job (508) thus allowing other compute jobs to immediately use the node within the cluster. If step (504) fails (506), then the system job reports the failure, and creates a system reservation for the node, and terminates the job (510) leaving the node in a reserve state until an administrator can respond to the failure and correct the operating system. This example was the application of a system job to allow for rolling maintenance.

[0039] Jobs associated with rolling maintenance that are scheduled are not a resource manager process. They are higher level jobs that perform arbitrary tasks outside processes handled by the resource manager. A trigger is a subset of a system job and has dependencies and can interface with web services, local processes, socket interfaces and can manage priorities. This allows an administrator to have the workload manager not being tied to a resource manager. The administrator can schedule a file system backup (e.g., job 1 and 2 will use the file system and job 3 will back up the file system). The scheduler typically has a locked model where the scheduler only knows about the resource manager.

[0040] Figure 6 shows another use of a system job, in particular for backing up a file system. In this particular situation, assume that a cluster has a number of file systems available and they are available across a parallel set of nodes. This scenario is illustrated in FIG. 7 in cluster 702 having a variety of sixteen nodes 704 with file system A (FSA),

file system B (FSB), file system C (FLC), and file system file system D (FSD). There are four nodes associated with each file system. Suppose the site has a goal of backing up each file system and in order to do that, it must quiesce each individual file system so that there is no activity when it is backed up. To quiesce each file system means to terminate activity thus allowing aspects of a parallel system to come to a completed state. When a system is quiesced, previously planned transmissions and signals are all delivered and activity is allowed to stop in a natural manner.

[0041] To accomplish this set of requirements, an object is created that submits a series of system jobs. The first system job requests allocation of all four nodes associated with file system A (602). This is performed using a feature requirement. Once it has all the nodes dedicated, the first step is that it issues a communication to the backup file system which backs up the file system (604). When that completes, the system job verifies the success of the process (606). In this case, regardless of whether the back was successful, the job reports the verification information and updates the database recording that information and then terminates allowing the nodes to be used by the user (608).

[0042] It is possible to modify the scenario slightly in which the file system must be quiesced. The file system can be quiesced for a period of time before everything synchronizes. Within a system job, it is possible to have the ability or step to force a duration, a step can either complete when its task is complete or when a duration has been reached. Therefore, this example could be modified so that step (602) simply to allocate the resources and quiesce them for a period of 10 minutes to allow full synchronization of the parallel aspects followed by the backup step (604) and step (606) which determines the success of the process, and wherein step (608) which updates the database with the success status.

[0043] To create a system job there are a number of different models. A system job can be automatically created by submitting a standard job to a particular quality of service

where the quality of service requires enablement of special services such as automatic provisioning or dedicated network bandwidth. In such a case, the user submits a standard job with a selected quality of service. An example of submitting a job to a quality of service that has a dedicated bandwidth. With such a request, the scheduler would take the job request and encapsulate it in a system job. The first step in a system job 1 is to identify the resources and then communicate with the network manager to dynamically partition the network so as to provide the guaranteed bandwidth. Once that is completed, the system job will proceed to allow the submitted job to process.

[0044] The same model is also used to allow data stage-in, data stage-out and have tightly coordinated resource usage after the environment is set up. The system jobs allow one to have a tight time frame control whereas the limitation of a job step simply says that one step will follow another step it does not constrain how tightly it must follow. Whereas a system job does allow one to tightly constrain that a subsequent job will run immediately allowing chaining of prerequisite job and post requisite steps. In the situation of a rolling maintenance, within the graphical user interface, a user does not even need to be aware that the system job exists. In most cases, system jobs are actually a method used under the covers to enable outlying functionality. In the case of a rolling maintenance the administrator simply goes to a graphical page and indicates that he or she would like this particular script to be run on all nodes which will automatically install the application. The administrator indicates that this application to be updated in all nodes using this particular provisioning manager. The rest of the steps are done automatically without his or her knowledge.

[0045] An important attribute of system jobs is that a system job is queueable. A system job can have dependency on types of resources, dependency on other system jobs or batch compute jobs. System jobs can incorporate dynamic content sensitive triggers, which allow them to customize the environment or customize the general local

scheduling environment. The steps in a system job may or may not have a duration, and they may or may not have a resource allocation or a resource co-allocation. They do have the ability to perform arbitrary execution or use arbitrary services. For example, system jobs can tap in and activate services such as a peer-to-peer service or a resource manager. Furthermore, system jobs can be reserved and can have relative or absolute priority.

[0046] Embodiments within the scope of the present invention may also include computer-readable media for carrying or having computer-executable instructions or data structures stored thereon. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer. By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to carry or store desired program code means in the form of computer-executable instructions or data structures. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or combination thereof) to a computer, the computer properly views the connection as a computer-readable medium. Thus, any such connection is properly termed a computer-readable medium. Combinations of the above should also be included within the scope of the computer-readable media.

[0047] Computer-executable instructions include, for example, instructions and data which cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. Computer-executable instructions also include program modules that are executed by computers in stand-alone or network environments. Generally, program modules include routines, programs, objects, components, and data structures, etc. that perform particular tasks or implement particular abstract data types. Computer-executable instructions, associated

data structures, and program modules represent examples of the program code means for executing steps of the methods disclosed herein. The particular sequence of such executable instructions or associated data structures represents examples of corresponding acts for implementing the functions described in such steps.

[0048] Those of skill in the art will appreciate that other embodiments of the invention may be practiced in network computing environments with many types of computer system configurations, including personal computers, hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. Embodiments may also be practiced in distributed computing environments where tasks are performed by local and remote processing devices that are linked (either by hardwired links, wireless links, or by a combination thereof) through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

CLAIMS:

1. A method for managing a multi-node compute environment, the method comprising:

 associating a system job to a compute job that is scheduled according to a reservation to consume resources in a multi-node compute environment under common management, wherein the system job is initiated independent of a resource manager that establishes the reservation for the compute job in the multi-node compute environment;

 and

 performing an action on a resource within the multi-node compute environment by processing the system job.
2. The method of claim 1, wherein the system job has at least one attribute selected from a list comprising event type, timeout, offset, action, dependency, argument list, state and threshold.
3. The method of claim 2, wherein the event type is at least one of a cancel, state and create event type.
4. The method of claim 1, wherein performing the action is based on a time offset associated with the beginning or completion of the compute job.

5. The method of claim 1, wherein performing the action comprises measuring a condition within the multi-node compute environment.
6. The method of claim 1, further comprising, after performing the action, to yield a performed action:

performing a health check of resources affected by the performed action.
7. The method of claim 6, wherein, if the health check indicates that the resources affected by the performed action are operational, releasing the resources for use in the multi-node compute environment for batch compute jobs.
8. The method of claim 6, wherein, if the health check indicates that the resources affected by the performed action is not operational, sending a message to an administrator.
9. The method of claim 1, wherein performing the action is triggered by an event associated with a second system job.
10. The method of claim 9, wherein the event associated with the second system job is completion of the second system job.

11. The method of claim 1, wherein performing the action is dependent on a second system job.

12. A system for managing a multi-node compute environment, the system comprising:

a processor; and

a computer-readable storage medium storing instructions, which, when performed by the processor, cause the processor to perform a method comprising:

associating a system job to a compute job that is scheduled according to a reservation to consume resources in a multi-node compute environment, wherein the system job is initiated independent of a resource manager that establishes the reservation for the compute job in the multi-node compute environment; and

performing an action on a resource within the multi-node compute environment by processing the system job.

13. A non-transitory computer-readable storage medium storing instructions which, when executed by a computing device, cause the computing device to manage a multi-node compute environment, the instructions comprising:

•
•
•
•


associating a system job to a compute job that is scheduled according to a reservation to consume resources in a multi-node compute environment, wherein the system job is initiated independent of a resource manager that establishes the reservation for the compute job in the multi-node compute environment; and

performing an action on a resource within the multi-node compute environment by processing the system job.

1/4

FIG. 1

PRIOR ART



2/4

FIG. 2

PRIOR ART





FIG. 3



3/4

FIG. 4


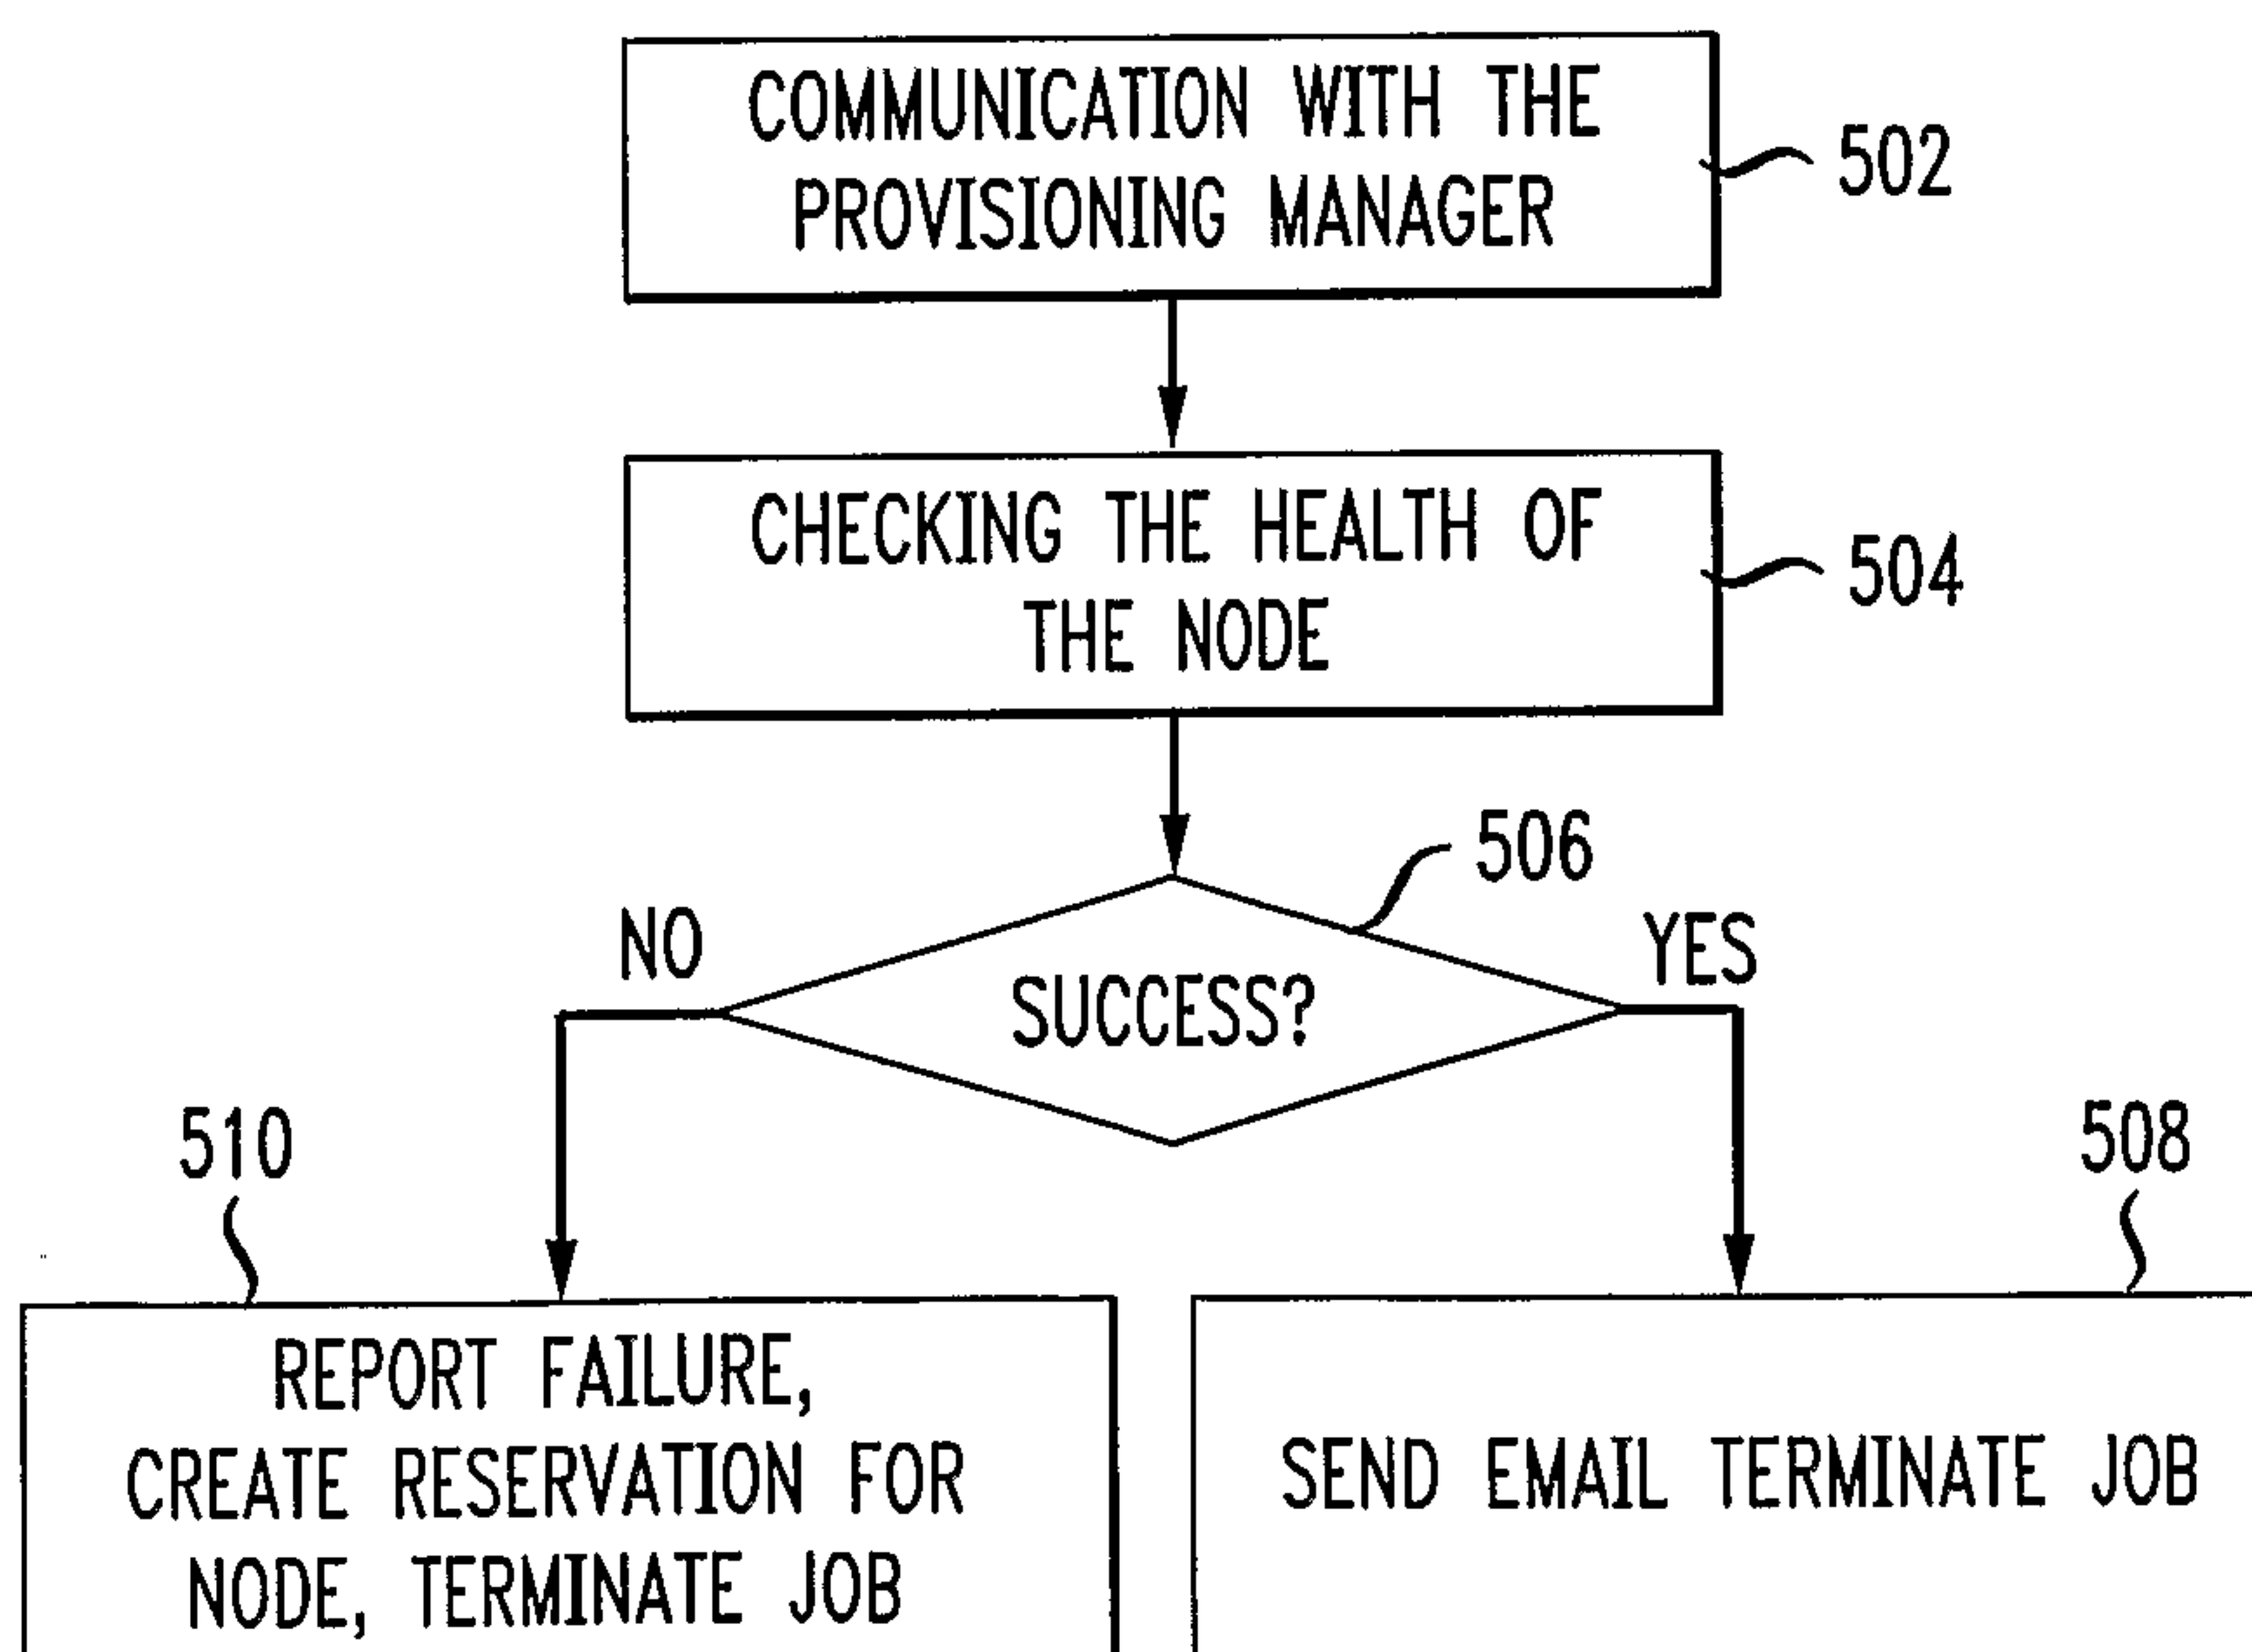



FIG. 5



4/4

FIG. 6*FIG. 7*