



(51) International Patent Classification:
G06F 15/80 (2006.01)

(21) International Application Number:
PCT/US2013/072600

(22) International Filing Date:
2 December 2013 (02.12.2013)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).

(72) Inventors: **CALHOUN, Michael Bozich**; 8000 Foothills Blvd., Roseville, California 95747 (US). **JIZRAWI, Alexander V.**; 8000 Foothills Blvd., Roseville, California 95747 (US).

(74) Agents: **PAGAR, Preetam B.** et al.; Hewlett-Packard Company, Intellectual Property Administration, 3404 E. Harmony Rd. MS:35, Fort Collins, Colorado 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, LT, 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, 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))

Published:

- with international search report (Art. 21(3))

(54) Title: SYSTEM WIDE MANAGEABILITY



(57) Abstract: Techniques for providing a system wide manageability interface are provided. In one aspect, manageability data may be received from a plurality of manageability subsystems. Each manageability subsystem may be contained in a domain. The received manageability data may be coalesced into a system wide view. The system wide view may span all domains.

SYSTEM WIDE MANAGEABILITY

BACKGROUND

[0001] Modern server computer systems may take many forms. A simple server may include a single processor, with a single core, on a single motherboard, in a single enclosure. A more complex system may include multiple processors and/or cores on a single motherboard. Even more complex systems may include multiple motherboards, containing varying numbers of processors and cores, each within their own enclosure. The enclosures may be coupled together via external cabling. In yet other complex systems, often referred to as blade systems, the motherboards (blades) with associated processors / cores are contained within a blade enclosure. The enclosure itself may provide connectivity paths between the blades, removing the need for additional cabling between the blades. Supporting functions, such as power and cooling, may be shared between the blades.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] FIG. 1 is an example of a system that may utilize a virtualized manageability data interface described herein.

[0003] FIG. 2 is another example of a system that may utilize a virtualized manageability data interface described herein.

[0004] FIG. 3 is an example of a high level flow diagram for receiving and coalescing manageability data according to the techniques described herein.

[0005] FIG. 4 is another example of a high level flow diagram for receiving and coalescing manageability data, as well as operating on the manageability data, according to the techniques described herein.

[0006] FIG. 5 is an example of a high level flow diagram for receiving and coalescing manageability data according to the techniques described herein.

[0007] FIG. 6 is another example of a high level flow diagram for receiving, coalescing, and local node configuration, based on manageability data according to the techniques described herein.

DETAILED DESCRIPTION

[0008] The availability of server systems in many different forms is useful to an end user. The end user may be able to select a system that best meets the needs of the workloads that are being run on the system. Unfortunately, what is beneficial for the end user may create problems for server system vendors.

[0009] System firmware allows the various components of a server system to communicate with each other and form a complete system. In a blade computer system, for example, the system firmware may allow multiple blades to be aggregated together, such that they operate as a single server system. There are several possible ways a server vendor may provide firmware with such capabilities. Each of these ways may have its own problems.

[0010] For example, a server system vendor may create entirely custom firmware for each possible configuration of server system. As should be clear, custom firmware for each possible system would result in an increase in costs for the vendor, resulting in increased costs for the end user. In addition, custom system firmware for each possible configuration would be very difficult for a vendor to maintain. Another example solution may involve designing the system firmware to accommodate the largest possible system that may ever be deployed. As should be clear, this would result in the firmware being overly complex for smaller systems, because the single firmware would have to support both large and small systems.

[0011] Another possible solution is to provide firmware on each individual component (e.g. each blade) that is capable of communicating with each other component to form a self discovering system firmware. For example this solution may involve the various components operating in a master/slave

relationship, wherein one component is the master and is responsible for configuration of the overall system. Again, such a solution increases complexity of the system firmware, as each component may need to be able to act as master or slave at any given time. Furthermore, fault handling becomes overly complex, as any recovery action is dependent on if component is acting as a master or slave.

[0012] Techniques described herein overcome the problems described above and others by providing a Virtualized Manageability Data Interface (VMDI) in system firmware. Modern servers include manageability subsystems, which are often referred to as baseboard management controllers (BMC). A BMC may be coupled to all manageable elements, such as processors and sensors, within a domain. The BMC, which can also be referred to as the manageability subsystem, may be coupled to the system firmware through the VMDI.

[0013] Upon system startup, the manageability subsystem of each node may send manageability data, which may also be referred to as manageability information, to the VMDI. Manageability data may include information such as the number and types of processors installed in the domain, sensor readings, installed memory, or any other information about the domain. It should be understood that each manageability subsystem only reports manageability information about its own domain.

[0014] The system firmware may then take the manageability data received through the VMDI from all of the domains and coalesce that information into a single, system wide view. This information may then be used by the system firmware to configure the system as a whole. For example, the system may include three domains, each with two processors. Each domain may report, through the VMDI, that there are two processors in the domain. The system firmware may aggregate this information to indicate that the overall system has six processors.

[0015] In addition, the system firmware may then translate system wide configuration commands into domain specific configuration commands, which may also be referred to as management commands. Continuing with the

processor example, the system firmware may need to assign processor numbers to each of the six processors. The system firmware may assign processor numbers one and two to the first domain, three and four to the second domain, and five and six to the third domain. These commands could be sent to the respective domains. However, there is no need for each domain to have any information about the other domains. Each domain is able to operate as if it were operating in isolation. As should be clear, this reduces the complexity of the management subsystem for each domain, because each domain need only be concerned with itself. The remainder of the overall system is hidden from the individual domains.

[0016] Finally, there may be some domain specific events that may need system wide information to be properly processed. For example, if all domains are sharing a limited power source, and one domain requests additional power, the system firmware may arbitrate this request. The system firmware may first determine how much power is being used by all domains combined. When sufficient power is available, the system firmware may grant the request. The requesting domain need not have any knowledge that other domains are also sharing the power source. Likewise, in some cases a single domain may cause an event that would affect the system as a whole. For example, a processor failure in a single domain may need to shut the entire system down. The system firmware may receive such events and determine if system wide action is needed.

[0017] Because each domain essentially operates in isolation for purposes of configuration and fault handling, server systems can be created with an arbitrary number of domains. If an end user's workload necessitates a smaller, single domain system, such a need can be satisfied. If an end user needs a larger, multi domain system, the need can be satisfied by combining multiple domains. Because the multi domain nature of the larger system is hidden by the system firmware, there is no need for specialized development within each of the domains. In other words, as system size needs to increase, additional domains are added in a modular fashion.

[0018] FIG. 1 is an example of a system 100 that may utilize a virtualized manageability data interface described herein. The system may include a plurality of domains 110-a-c. Although only three domains are shown, it should be understood that the system may have any number of domains. The system may also include a non-transitory processor readable medium 120. The medium, may contain thereon instructions which form the system firmware. For example, the instructions may include instructions to implement a VMDI 130.

[0019] For purposes of simplicity of explanation, domain 110-a will be described. However, it should be understood that all domains may have a similar structure. Domain 110-a may include a node 111-a. A node, for example, may be a modular component that may be combined with other nodes to form larger systems. For example, a node may be an individual blade in a blade computing system. A node may be a standalone server contained in an enclosure that may be coupled to other nodes to form larger systems. Regardless of the specific form factor, what should be understood is that nodes may be the building blocks through which larger systems are created.

[0020] Each node may include one or more processors 112-a. The specific number of processors per node is relatively unimportant. Likewise, each processor may have one or more cores (not shown). Again, the specific number of cores per processor is relatively unimportant. The processor 112-a may be responsible for executing the workloads assigned to the node. The node may also include sensors 113-a. These sensors may measure the environmental conditions of the node. For example, there may be sensors that measure the temperature of the processor, of the incoming cooling air, and the outgoing exhaust air. In addition, there may be sensors that measure the amount of current or voltage being generated and/or drawn by various components in the node. The sensors may be used to detect fault conditions within the node.

[0021] The node may also include a management subsystem 114-a. The management subsystem may include a management processor 115-a. The management processor may often be referred to as a baseboard management processor (BMC). The management processor may be coupled to the

processor 112-a and sensors 113-a and is able to read and communicate the status of those components to the system firmware, which is described in further detail below. The management processor may also be coupled to a management firmware (not shown). The management firmware may provide the instructions necessary to allow the management subsystem 114-a to communicate with the system firmware 120.

[0022] It should be noted that in some cases, the management subsystem 114-a and its associated firmware may not be provided by the vendor of the node. For example, the node vendor may simply purchase the management subsystem as an off-the-shelf component to be integrated into the node. This is very similar to how many different node vendors all utilize off-the-shelf processors provided by processor companies such as Intel™. As such, the solution describe herein is operable without requiring customization of the management subsystem firmware, as such customization may not even be possible given that the node vendor may be using an off-the-shelf management subsystem.

[0023] As shown in FIG. 1, there appears to be a one to one correspondence between a domain and a node. However, this depiction is for purposes of simplicity of explanation, not by way of limitation. A domain is intended to include all components that are managed by a single management subsystem. A node is intended to include all resources that may be available to allow the system to operate on workloads. Thus, it is possible that multiple nodes may be contained in the same domain, with only one node including a management subsystem. Although in many cases there will be a one to one relationship between domains and nodes.

[0024] The system firmware 120 may include VMDI instructions 130. The VMDI instructions may be executed by the processors 112-a-c in the domains 110-a-c as part of the system firmware. System firmware are the instructions that allow the processors from the various domains to communicate with each other and to behave as a single system. System firmware may be defined as software that is tightly coupled to the system and runs on the same hardware that is used by the system to process workloads. In other words, the system

firmware operates in the same address space as the operating system and applications running on the system. The VMDI instructions provide an interface for each of the management subsystems of the domains to communicate domain specific information to the system firmware. The system firmware may coalesce this information into a system wide view by aggregating the management information from each of the domains. The system firmware may then make system wide configuration decisions. Finally, the system firmware may translate the system wide configuration into domain specific management commands. Thus, each domain need not be aware of the system wide configuration, but rather only needs to respond to commands sent to that domain alone

[0025] In operation, at system startup, each domain 110, through its respective management subsystem 114, may send manageability information to the system firmware through the VMDI 130. Each domain may be unaware that it is communicating with a system firmware that is also communicating with other domains. Thus, each domain has no knowledge of the other domains. The system firmware may then proceed to make system wide configuration decisions. These system wide configuration decisions may include any of the types of configuration that are needed for a large system. Specific parameters are described in further detail below.

[0026] Once the system firmware has made configuration decisions, these decisions may be segmented and translated into management commands, also referred to as events, that are specific to a particular node. These types of commands are referred to as node local commands or events, because they are only applicable to the node that is receiving the command. In other words, system wide configuration is broken down into the configuration needed by each node individually, without requiring each node to have complete knowledge of the overall larger system. Upon receipt of the node local commands, each node may configure itself accordingly. After each node has configured itself according to the node local commands, the entire system is then configured as one larger system.

[0027] During operation of the system, an individual node may need to request a resource that is shared between all of the nodes. In such cases, the node may send a request to the system firmware through the VMDI. The node itself is only aware that it is sending the resource request to the system firmware, as the VMDI is completely transparent to the node. The system firmware then determines if the resource request may be granted, based on the resource usage of the system as a whole. If the request is granted, the node may be sent a management command indicating that the resource request was granted.

[0028] A node may also request a shared resource setting that may have impacts to the overall system. The system firmware may receive the request through the VMDI and make a determination as to the proper system wide setting. A node may also request use of a resource that may only be used by one node at a time. The system firmware may determine which node should be granted use of the resource. In yet another example, a node may experience a fault that may or may not impact the entire system. The fault indication may be sent to the system firmware. The system firmware may decide if system wide action (such as shutdown) is needed, or if the fault is isolated to a particular node, such that only node local corrective action is needed.

[0029] FIG. 2 is another example of a system that may utilize a virtualized manageability data interface described herein. FIG. 2 generally contains the same elements as FIG. 1, and for simplicity of explanation, the description of those elements is not repeated. FIG. 2 describes additional elements of the VMDI. For example, the VMDI may include system configuration tables 240, system sensor tables 245, system data tables 250, and system resource tables 255. In addition, FIG. 2 depicts certain components and resources that may be shared 260 by all nodes within the system. For example, FIG. 2 depicts power resources 265, cooling resources 270, and notification resources 275. Although specific examples, are presented, it should be understood that these examples are for purposes of explanation, and not limitation.

[0030] System configuration 240 may include information about the current system configuration and status. For example, the system configuration can

include information about which resources have been assigned to which domains or which resources are out of service. When node management information is received from a node, the system configuration may be examined to determine the appropriate action to take. For example, if a node is requesting additional power 265, the system configuration may be reviewed to determine how much power remains to be allocated. Depending on this result, the request for additional power may or may not be granted.

[0031] The system firmware may also include system sensor tables 245. The system sensor tables may store the current status of all the sensors in each of the domains. The system firmware may then make decisions based on the data contained in the system sensor table. For example, in one implementation, if all the processors in a single domain are reporting a high temperature, but none of the processors in any other domain are reporting a high temperature, the system firmware may determine there is a problem with only the node reporting high temperature. As such, corrective action may be taken on that node only. The system firmware may generate node local commands to shut down the offending node. However, if all nodes are reporting out of range temperature, the problem may be system wide. The system firmware may then generate node local commands for all nodes, instructing the nodes to shut down, which essentially shuts down the entire system. In other implementations, the system firmware may treat the entire system as a whole, and an error condition from any node may result in system wide corrective action. Thus, depending on the particular implementation and error conditions, the system firmware may take node specific or system wide corrective action.

[0032] The system firmware may also include system data tables 250. The system data tables may include information that may be consumed by the system firmware when configuring the system. As mentioned above, each domain may send node local management information to the system firmware through the VMDI. For example, each node may send information related to the hardware configured on the node. One example may be the number of processors configured on the node. The system firmware may then coalesce all of the information received from the various nodes into the system data tables.

These tables may reflect all of the hardware in the system, regardless of which node the hardware resides on. The table data may then be consumed by configuration routines within the system firmware or the operating system to configure the system as a whole.

[0033] The system firmware may also include a system resource list 255. The system resource list may include information about resources that are shared by all nodes within the system. For example, cooling resources 270, such as fans, may be shared by all nodes within the system. As another example, notification resources 275 may be shared by all nodes in the system. An example of a notification resource may be a display panel on the front of an enclosure, such as a blade enclosure. The display panel may only be usable by a single node at a time.

[0034] The system firmware may make use of the system resource table 255 when determining how to configure system wide resources. For example, an individual node may be experiencing below normal operating temperatures, and may request that the speed of the system cooling fans be decreased. However, the node itself is not aware that the overall system contains other nodes, all of which may not be experiencing the same operating temperature conditions. The system resource list may be used to determine the proper settings for resources that are shared between all of the nodes.

[0035] As another example, the system may include a notification resources 275, such as a display panel on the front of an enclosure. If multiple nodes have information that is to be displayed on the notification panel, the system firmware may need to make a decision as to which node is to be granted use of the notification panel. The system resource list may include information about which resources have been granted to which nodes. The system resource list may be used when determining how shared resources are to be granted to nodes.

[0036] FIG. 3 is an example of a high level flow diagram for receiving and coalescing manageability data according to the techniques described herein. In block 310, manageability data may be received from a plurality of manageability subsystems. Each of the manageability subsystems may be contained in a

domain. As described above, a domain may include all of the hardware resources that are managed by a single manageability subsystem.

[0037] In block 320, the received manageability data may be coalesced into a system wide view spanning all domains. As described above, multiple independent domains may be aggregated to form a larger system. Each individual domain may maintain its independent nature. However, the virtualized manageability data interface provides a mechanism for system firmware to receive manageability data from each of the domains and combine that data into a larger system wide view.

[0038] FIG. 4 is another example of a high level flow diagram for receiving and coalescing manageability data, as well as operating on the manageability data, according to the techniques described herein. In block 405, just as above, manageability data may be received from a plurality of manageability subsystems. Each of the manageability subsystems may be contained in a domain. In block 410 the received manageability data may be coalesced into a system wide view spanning all domains.

[0039] In block 415, system wide configuration may be determined based on the system wide view. As explained above, the system firmware is made aware of all of the hardware resources present in each of the domains. The system firmware may then determine how to configure the entire system, as it would be configured if all of the resources of the plurality of domains were all contained in the same domain. In block 420, the configuration may be translated into domain specific configuration. As explained above, each of the domains are independent from each other. Thus, when a particular domain is to be configured, the configuration must be specific to that domain. In block 420, the domain specific configuration may be sent to the manageability subsystem in each domain. The manageability subsystems may then configure the resources of each domain in accordance with the domain specific configurations.

[0040] In block 430, a resource request from the manageability subsystem of a single domain may be received. A domain may be requesting a resource for which there is a limited system wide supply. For example, one of the domains may be requesting additional power in order to allow a processor within that

domain to operate at a higher speed. There may be a system wide limitation on the amount of power available.

[0041] In block 435, it may be determined that sufficient resources are available based on the system wide view. As explained above, the system firmware maintains a system wide view of all available resources. When sufficient resources are available to grant a request from a single domain, the resources may be allocated to that domain. In block 440, the resource request may be granted when sufficient resources are available.

[0042] In block 445, a configuration request may be received from the manageability subsystem of a single domain. A configuration request may be different than a resource request, in that a configuration request may be requesting configuration changes that affect the entire system, as opposed to a single domain. For example, a system fan may only spin at one speed at a time. A domain may be experiencing higher than normal operating temperatures and might request a change in configuration to increase the speed of the system fan. However, changing the speed of the system fan affects all domains, not just the requesting domain.

[0043] In block 450, the proper system wide configuration may be determined based on the configuration request and the system wide view. The system firmware may evaluate the request from a single domain in view of the system wide configuration to determine the correct system wide configuration. For example, in the case of a system fan, if one domain requests additional cooling, the proper response may be to increase the speed of the fan, even if no other domain requests additional cooling. As another example, a domain may request use of the notification facilities of the system. However, a different domain may already be using the notification facilities. The system firmware may then determine the proper configuration of the notification facilities, given a system wide view. For example, the system firmware may determine which domain's notification information is of higher priority, and may grant the use of the notification facilities to that domain. In block 455, the system may be configured based on the determination. In other words, the system may be configured after taking into account the configuration requests from all domains.

[0044] In block 460, the fault indication may be received from the manageability subsystem of a single domain. As explained above, each manageability subsystem is coupled to the hardware resources of the domain. If a fault occurs, the manageability subsystem may send an indication of the fault to the system firmware. In block 465, a determination of when system wide corrective action is needed may be made. The determination may be made based on the system wide view. Some faults that occur within a domain may not have any impact on any other domains. For example, an indication of a corrected memory error would have no impact on the system as a whole. However, a fault such as a processor failure due to overheating, may necessitate shutting down all domains that make up the system. The system firmware may decide the proper corrective action based on the system wide view. In block 470, the determined system wide corrective action may be executed. In some cases, it may be determined that no corrective action is required. In other cases, system shutdown may be needed. The corrective action may lie somewhere in between.

[0045] FIG. 5 is an example of a high level flow diagram for receiving and coalescing manageability data according to the techniques described herein. In block 510, manageability information may be received from a plurality of nodes through a virtualized manageability data interface. Each node may comprise a manageability subsystem. As described above, each node may operate independently of all other nodes and include a manageability subsystem. The manageability subsystem may send manageability information to a system firmware.

[0046] In block 520, the manageability information may be coalesced into a system wide view. As explained above, system firmware may provide a virtualized manageability data interface to receive manageability information from all of the nodes. The system firmware may then coalesce the received information into a system wide view that includes the information from all of the nodes that make up the system.

[0047] FIG. 6 is another example of a high level flow diagram for receiving, coalescing, and local node configuration, based on manageability data

according to the techniques described herein. In block 610, just as above, manageability information may be received from a plurality of nodes through a virtualized manageability data interface. Each node may comprise a manageability subsystem. In block 620, the manageability information may be coalesced into a system wide view.

[0048] In block 630, the proper system wide configuration may be determined. For example, after the manageability information from the plurality of nodes is received, the system firmware has a system wide view of all of the resources that are being provided by each node. The system firmware may then be able to determine how those resources are to be configured for the system as a whole.

[0049] In block 640, the proper system wide configuration may be translated into node local management commands. As explained above, each node may be unaware of the other nodes that, when combined, form the overall system. Thus, when the system has determined the proper configuration overall, those configuration commands may be translated into commands that address each node individually. These are the node local commands. In block 650, the node local management commands may be sent to the manageability subsystem of each node of the plurality of nodes. In other words, each node may receive configurations commands that affect that node only. Once all nodes are configured individually, the resultant combination is a system with operational parameters as determined by the system firmware.

We Claim:

1. A system comprising:

a plurality of nodes, each node comprising a manageability subsystem to perform node management functions; and

a system firmware running on the nodes, the system firmware providing a virtualized manageability data interface coupled to the manageability subsystems, to provide system wide management functions.

2. The system of claim 1 wherein each node of the plurality of nodes further comprises:

a management firmware to send node local management events to the system firmware, wherein each nodes management events are independent of other nodes.

3. The system of claim 1 wherein each node of the plurality of nodes further comprises:

a management firmware to receive node local management commands, wherein the node local management commands are generated by the system firmware, the system firmware generating the commands based on the plurality of nodes.

4. The system of claim 1 wherein the system firmware further:

receives management events from the plurality of nodes through the virtualized manageability data interface; and

consolidates the events into a system wide view.

5. The system of claim 1 wherein the system firmware further:

determines system wide management actions;

segments the system wide management actions into node local commands; and

sends the node local commands to the plurality of nodes.

6. The system of claim 1 wherein the system firmware further:
 - receives node configuration data from each of the plurality of nodes; and
 - aggregates the node configuration data into system wide configuration data, wherein an operating system access the system wide configuration data for system management.
7. The system of claim 1 wherein the system firmware further:
 - allocates shared system resources based on information received from the management subsystems through the virtualized manageability data interface.
8. The system of claim 1 wherein the system firmware further:
 - initiates system wide actions based on a node local fault event.
9. A method comprising:
 - receiving manageability data from a plurality of manageability subsystems, each manageability subsystem contained in a domain; and
 - coalescing the received manageability data into a system wide view spanning all domains.
10. The method of claim 9, further comprising:
 - determining system wide configuration based on the system wide view;
 - translating the system wide configuration into domain specific configuration; and
 - sending the domain specific configuration to the manageability subsystem in each domain.
11. The method of claim 9 further comprising:
 - receiving a resource request from the manageability subsystem of a single domain;
 - determining that sufficient resources are available based on the system wide view; and

granting the resource request when sufficient resources are available.

12. The method of claim 9 further comprising:

receiving a configuration request from the manageability subsystem of a single domain;

determining a proper system wide configuration based on the configuration request and the system wide view; and

configuring the system based on the determination.

13. The method of claim 9 further comprising:

receiving a fault indication from the manageability subsystem of a single domain;

determining when system wide corrective action is needed based on the system wide view; and

executing the determined system wide corrective action.

14. A non-transitory processor readable medium containing thereon a set of instructions, which when executed by the processor cause the processor to:

receive, through a virtualized manageability data interface, manageability information from a plurality of nodes, wherein each node comprises a manageability subsystem; and

coalesce the manageability information into a system wide view.

15. The medium of claim 15 further comprising instructions to:

determine proper system wide configuration;

translate the proper system wide configuration into node local management commands; and

send the node local management commands to the manageability subsystem of each node of the plurality of nodes.




FIG. 1




FIG. 2

3 / 6




FIG. 3

4 / 6




FIG. 4

5 / 6




FIG. 5

6 / 6




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/072600**A. CLASSIFICATION OF SUBJECT MATTER****G06F 15/80(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 15/80; G06F 15/16; G06F 12/16; G06F 1/00

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: manageability, node, subsystem, firmware

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2006-0095702 A1 (JOHN HICKMAN et al.) 04 May 2006 See abstract; claims 1-9; and figures 3-6.	1-15
A	US 2011-0295944 A1 (BUONOMO MICHAEL ANTHONY et al.) 01 December 2011 See abstract; claims 1, 8; and figure 1.	1-15
A	US 2009-0216910 A1 (DUCHESNEAU DAVID D.) 27 August 2009 See abstract; claim 1; and figure 7.	1-15
A	US 2007-0101167 A1 (CRAIG A. LINDLEY et al.) 03 May 2007 See abstract; claims 1-5; and figures 1-4.	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

31 July 2014 (31.07.2014)

Date of mailing of the international search report

01 August 2014 (01.08.2014)

Name and mailing address of the ISA/KR


International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

LEE, Seok Hyung

Telephone No. +82-42-481-5983



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/072600

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2006-0095702 A1	04/05/2006	CN 1760851 A CN 1760851 C0 TW I354207 B US 7734753 B2	19/04/2006 30/04/2008 11/12/2011 08/06/2010
US 2011-0295944 A1	01/12/2011	US 8676940 B2	18/03/2014
US 2009-0216910 A1	27/08/2009	CA 2718136 A1 EP 2145362 A2 US 2013-306276 A1 US 8706914 B2 US 8706915 B2 WO 2008-131446 A2 WO 2008-131446 A3	30/10/2008 20/01/2010 21/11/2013 22/04/2014 22/04/2014 30/10/2008 11/02/2010
US 2007-0101167 A1	03/05/2007	US 8745124 B2	03/06/2014