



- (51) International Patent Classification:
G06F 21/50 (2013.01) H04L 9/08 (2006.01)
G06F 21/60 (2013.01)
- (21) International Application Number:
PCT/US2017/043343
- (22) International Filing Date:
21 July 2017 (21.07.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:

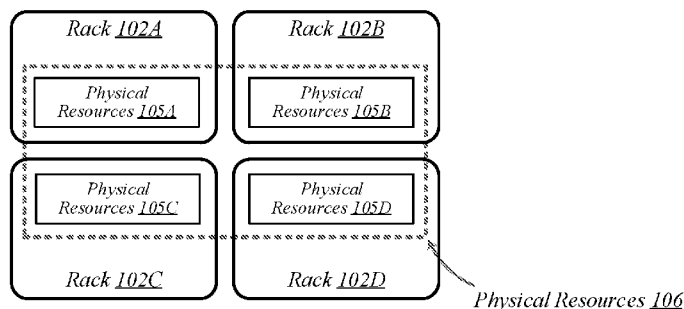
| | | |
|------------|-------------------------------|----|
| 62/365,969 | 22 July 2016 (22.07.2016) | US |
| 62/376,859 | 18 August 2016 (18.08.2016) | US |
| 62/427,268 | 29 November 2016 (29.11.2016) | US |
| 15/656,798 | 21 July 2017 (21.07.2017) | US |

- (71) Applicant: INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).
- (72) Inventors: MUNOZ, Alberto J.; 1739 Hawkins Drive, Los Altos, California 94024 (US). NACHIMUTHU, Murugasamy K.; 8441 SW 186th Avenue, Beaverton, Oregon 97007 (US). KUMAR, Mohan J.; 18680 SW Marko Lane, Aloha, Oregon 97007 (US). POWIERTOWSKI, Wojciech; 1906 SW Elliott Avenue, Portland, Oregon 97214 (US). GHETIE, Sergiu D.; 897 NE Brennan Lane, Hillsboro, Oregon 97124 (US). UPASANI, Neeraj S.; 16803 NW Oakdridge Drive, Portland, Oregon 97229 (US). DALVI, Sagar V.; 2111 NE 25th Avenue, Hillsboro, Oregon 97214 (US). NNEBE, Chukwunye S.; 965 Palmer Circle, Folsom, California 95630 (US). GULLORY, Jeanne; 2111 NE 25th Avenue, Hillsboro, Oregon 97124 (US).

(54) Title: TECHNIQUES TO VERIFY AND AUTHENTICATE RESOURCES IN A DATA CENTER COMPUTER ENVIRONMENT

FIG. 1

Data Center 100



(57) Abstract: Embodiments are generally directed apparatuses, methods, techniques and so forth to receive a sled manifest comprising identifiers for physical resources of a sled, receive results of an authentication and validation operations performed to authenticate and validate the physical resources of the sled, determine whether the results of the authentication and validation operations indicate the physical resources are authenticate or not authenticate. Further and in response to the determination that the results indicate the physical resources are authenticated, permit the physical resources to process a workload, and in response to the determination that the results indicate the physical resources are not authenticated, prevent the physical resources from processing the workload.



(74) **Agent:** RUSSELL, Matthew S.; KACVINSKY DAISAK
BLUNI PLLC, 430 Davis Drive, Suite 150, Morrisville,
North Carolina 27560 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, 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).

Published:

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

**TECHNIQUES TO VERIFY AND AUTHENTICATE RESOURCES IN A DATA
CENTER COMPUTER ENVIRONMENT**

RELATED CASES

5 This application claims priority to United States Patent Application Serial Number
15/656,798 filed July 21, 2017, United States Provisional Patent Application Number
62/365,969, filed July 22, 2016, United States Provisional Patent Application Number
62/376,859, filed August 18, 2016, and United Provisional Patent Application Number
62/427,268, filed November 29, 2016, each of which are hereby incorporated by reference in
10 their entirety.

TECHNICAL FIELD

Embodiments described herein generally include verifying hardware and software
resources in a data center computer environment.

BACKGROUND

15 A computing data center may include one or more computing systems including a
plurality of compute nodes that may include various compute structures (e.g., servers or sleds)
and may be physically located on multiple racks. The sleds may include a number of physical
resources interconnected via one or more compute structures and buses. In some instances, a
computing data center including the sleds and physical resources may process sensitive,
20 confidential and/or valuable information such as medical or financial records, proprietary
business data, and licensed multimedia content. Thus, customers of processing this information
may require strict control of the hardware and software used to process this valuable information
to prevent attacks and ensure the integrity of the data center's physical resources. Thus,
embodiments may be directed to preventing changes to hardware and software and solving other
25 problems to ensure the integrity of a data center.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the invention are illustrated by way of example, and not by way of
limitation, in the figures of the accompanying drawings in which like reference numerals refer to
similar elements.

30 **FIG. 1** illustrates an example of a data center.

FIG. 2 illustrates an example of a rack.

FIG. 3 illustrates an example of a data center.

FIG. 4 illustrates an example of a data center.

FIG. 5 illustrates an example of a switching infrastructure.

FIG. 6 illustrates an example of a data center.

FIG. 7 illustrates an example of a sled.

FIG. 8 illustrates an example of a data center.

FIG. 9 illustrates an example of a data center.

5 **FIG. 10** illustrates an example of a sled.

FIG. 11 illustrates an example of a data center.

FIG. 12 illustrates an example of a data center.

FIG. 13A illustrates an example of a sled.

FIG. 13B illustrates an example of a distribution and attestation flow.

10 **FIG. 14** illustrates an example of a data center.

FIG. 15 illustrates an example of a first logic flow.

FIG. 16 illustrates an example of a second logic flow.

FIG. 17 illustrate a second example of a third logic flow.

DETAILED DESCRIPTION

15 Various embodiments may generally be directed to verifying and authenticating resources in a data center computer environment. More specifically, embodiments include authenticating and verifying each physical resource of a sled via a verification controller of the sled. The verification controller may be a secure hardware component capable of processing instructions stored in a secure hardware location, for example. The verification controller may
20 generate results of the verification and authentication of the resources indicating whether they are successfully verified and authenticated or not.

In embodiments, the results and a sled manifest indicating each of the physical resources may be provided to a pod management controller. The pod management controller may verify and authenticate the results and the sled manifest and based on the verification and
25 authentication of the results and the sled manifest; the pod management controller may permit or prevent physical resources to be used in composed nodes. These and other details will become more apparent in the following description.

Reference is now made to the drawings, wherein like reference numerals are used to refer to like elements throughout. In the following description, for purposes of explanation, numerous
30 specific details are set forth in order to provide a thorough understanding thereof. It may be evident, however, that the novel embodiments can be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to facilitate a description thereof. The intention is to cover all modifications, equivalents, and alternatives consistent with the claimed subject matter.

FIG. 1 illustrates a conceptual overview of a data center 100 that may be representative of a data center or other type of computing network in/for which one or more techniques described herein may be implemented per various embodiments. As shown in FIG. 1, data center 100 may contain a plurality of racks, each of which may house computing equipment comprising a respective set of physical resources. In the non-limiting example depicted in FIG. 1, data center 100 contains four racks 102A to 102D, which house computing equipment comprising respective sets of physical resources (PCRs) 105A to 105D. According to this example, a collective set of physical resources 106 of data center 100 includes the various sets of physical resources 105A to 105D that are distributed among racks 102A to 102D. Physical resources 106 may include resources of multiple types, such as – for example – processors, co-processors, accelerators, field-programmable gate arrays (FPGAs), memory, and storage. The embodiments are not limited to these examples.

The illustrative data center 100 differs from typical data centers in many ways. For example, in the illustrative embodiment, the circuit boards (“sleds”) on which components such as CPUs, memory, and other components are placed are designed for increased thermal performance. In particular, in the illustrative embodiment, the sleds are shallower than typical boards. In other words, the sleds are shorter from the front to the back, where cooling fans are located. This decreases the length of the path that air must travel across the components on the board. Further, the components on the sled are spaced further apart than in typical circuit boards, and the components are arranged to reduce or eliminate shadowing (i.e., one component in the air flow path of another component). In the illustrative embodiment, processing components such as the processors are located on a top side of a sled while near memory, such as DIMMs, are located on a bottom side of the sled. As a result of the enhanced airflow provided by this design, the components may operate at higher frequencies and power levels than in typical systems, thereby increasing performance. Furthermore, the sleds are configured to blindly mate with power and data communication cables in each rack 102A, 102B, 102C, 102D, enhancing their ability to be quickly removed, upgraded, reinstalled, and/or replaced. Similarly, individual components located on the sleds, such as processors, accelerators, memory, and data storage drives, are configured to be easily upgraded due to their increased spacing from each other. In the illustrative embodiment, the components additionally include hardware attestation features to prove their authenticity.

Furthermore, in the illustrative embodiment, the data center 100 utilizes a single network architecture (“fabric”) that supports multiple other network architectures including Ethernet and Omni-Path. The sleds, in the illustrative embodiment, are coupled to switches via optical fibers,

which provide higher bandwidth and lower latency than typical twisted pair cabling (e.g., Category 5, Category 5e, Category 6, etc.). Due to the high bandwidth, low latency interconnections and network architecture, the data center 100 may, in use, pool resources, such as memory, accelerators (e.g., graphics accelerators, FPGAs, ASICs, etc.), and data storage drives that are physically disaggregated, and provide them to compute resources (e.g., processors) on an as needed basis, enabling the compute resources to access the pooled resources as if they were local. The illustrative data center 100 additionally receives usage information for the various resources, predicts resource usage for different types of workloads based on past resource usage, and dynamically reallocates the resources based on this information.

10 The racks 102A, 102B, 102C, 102D of the data center 100 may include physical design features that facilitate the automation of a variety of types of maintenance tasks. For example, data center 100 may be implemented using racks that are designed to be robotically-accessed, and to accept and house robotically-manipulability resource sleds. Furthermore, in the illustrative embodiment, the racks 102A, 102B, 102C, 102D include integrated power sources that receive higher current than typical for power sources. The increased current enables the power sources to provide additional power to the components on each sled, enabling the components to operate at higher than typical frequencies. **FIG. 2** illustrates an exemplary logical configuration of a rack 202 of the data center 100. As shown in FIG. 2, rack 202 may generally house a plurality of sleds, each of which may comprise a respective set of physical resources. In the particular non-limiting example depicted in FIG. 2, rack 202 houses sleds 204-1 to 204-4 comprising respective sets of physical resources 205-1 to 205-4, each of which constitutes a portion of the collective set of physical resources 206 comprised in rack 202. With respect to FIG. 1, if rack 202 is representative of – for example – rack 102A, then physical resources 206 may correspond to the physical resources 105A comprised in rack 102A. In the context of this example, physical resources 105A may thus be made up of the respective sets of physical resources, including physical storage resources 205-1, physical accelerator resources 205-2, physical memory resources 204-3, and physical compute resources 205-5 comprised in the sleds 204-1 to 204-4 of rack 202. The embodiments are not limited to this example. Each sled may contain a pool of each of the various types of physical resources (e.g., compute, memory, accelerator, storage). By having robotically accessible and robotically manipulability sleds comprising disaggregated resources, each type of resource can be upgraded independently of each other and at their own optimized refresh rate.

FIG. 3 illustrates an example of a data center 300 that may generally be representative of one in/for which one or more techniques described herein may be implemented according to

various embodiments. In the particular non-limiting example depicted in FIG. 3, data center 300 comprises racks 302-1 to 302-32. In various embodiments, the racks of data center 300 may be arranged in such fashion as to define and/or accommodate various access pathways. For example, as shown in FIG. 3, the racks of data center 300 may be arranged in such fashion as to define and/or accommodate access pathways 311A, 311B, 311C, and 311D. In some embodiments, the presence of such access pathways may generally enable automated maintenance equipment, such as robotic maintenance equipment, to physically access the computing equipment housed in the various racks of data center 300 and perform automated maintenance tasks (e.g., replace a failed sled, upgrade a sled). In various embodiments, the dimensions of access pathways 311A, 311B, 311C, and 311D, the dimensions of racks 302-1 to 302-32, and/or one or more other aspects of the physical layout of data center 300 may be selected to facilitate such automated operations. The embodiments are not limited in this context.

FIG. 4 illustrates an example of a data center 400 that may generally be representative of one in/for which one or more techniques described herein may be implemented according to various embodiments. As shown in FIG. 4, data center 400 may feature an optical fabric 412. Optical fabric 412 may generally comprise a combination of optical signaling media (such as optical cabling) and optical switching infrastructure via which any particular sled in data center 400 can send signals to (and receive signals from) each of the other sleds in data center 400. The signaling connectivity that optical fabric 412 provides to any given sled may include connectivity both to other sleds in a same rack and sleds in other racks. In the particular non-limiting example depicted in FIG. 4, data center 400 includes four racks 402A to 402D. Racks 402A to 402D house respective pairs of sleds 404A-1 and 404A-2, 404B-1 and 404B-2, 404C-1 and 404C-2, and 404D-1 and 404D-2. Thus, in this example, data center 400 comprises a total of eight sleds. Via optical fabric 412, each such sled may possess signaling connectivity with each of the seven other sleds in data center 400. For example, via optical fabric 412, sled 404A-1 in rack 402A may possess signaling connectivity with sled 404A-2 in rack 402A, as well as the six other sleds 404B-1, 404B-2, 404C-1, 404C-2, 404D-1, and 404D-2 that are distributed among the other racks 402B, 402C, and 402D of data center 400. The embodiments are not limited to this example.

FIG. 5 illustrates an overview of a connectivity scheme 500 that may generally be representative of link-layer connectivity that may be established in some embodiments among the various sleds of a data center, such as any of example data centers 100, 300, and 400 of FIGs. 1, 3, and 4. Connectivity scheme 500 may be implemented using an optical fabric that features a

dual-mode optical switching infrastructure 514. Dual-mode optical switching infrastructure 514 may generally comprise a switching infrastructure that is capable of receiving communications according to multiple link-layer protocols via a same unified set of optical signaling media, and properly switching such communications. In various embodiments, dual-mode optical switching infrastructure 514 may be implemented using one or more dual-mode optical switches 515. In various embodiments, dual-mode optical switches 515 may generally comprise high-radix switches. In some embodiments, dual-mode optical switches 515 may comprise multi-ply switches, such as four-ply switches. In various embodiments, dual-mode optical switches 515 may feature integrated silicon photonics that enable them to switch communications with significantly reduced latency in comparison to conventional switching devices. In embodiments, the dual-mode switch may be a single physical network wire that may be capable of carrying Ethernet or Omni-Path communication, which may be auto-detected by the dual-mode optical switch 515 or configured by the Pod management controller. This allows for the same network to be used for Cloud traffic (Ethernet) or High Performance Computing (HPC), typically Omni-Path or Infiniband. Moreover, and in some instances, an Omni-Path protocol may carry Omni-Path communication and Ethernet communication. In some embodiments, dual-mode optical switches 515 may constitute leaf switches 530 in a leaf-spine architecture additionally including one or more dual-mode optical spine switches 520. Note that in some embodiments, the architecture may not be a leaf-spine architecture, but may be a two-ply switch architecture to connect directly to the sleds.

In various embodiments, dual-mode optical switches may be capable of receiving both Ethernet protocol communications carrying Internet Protocol (IP packets) and communications according to a second, high-performance computing (HPC) link-layer protocol (e.g., Intel's Omni-Path Architecture's, Infiniband) via optical signaling media of an optical fabric. As reflected in FIG. 5, with respect to any particular pair of sleds 504A and 504B possessing optical signaling connectivity to the optical fabric, connectivity scheme 500 may thus provide support for link-layer connectivity via both Ethernet links and HPC links. Thus, both Ethernet and HPC communications can be supported by a single high-bandwidth, low-latency switch fabric. The embodiments are not limited to this example.

FIG. 6 illustrates a general overview of a rack architecture 600 that may be representative of an architecture of any particular one of the racks depicted in FIGs. 1 to 4 according to some embodiments. As reflected in FIG. 6, rack architecture 600 may generally feature a plurality of sled spaces into which sleds may be inserted, each of which may be robotically-accessible via a rack access region 601. In the particular non-limiting example

depicted in FIG. 6, rack architecture 600 features five sled spaces 603-1 to 603-5. Sled spaces 603-1 to 603-5 feature respective multi-purpose connector modules (MPCMs) 616-1 to 616-5. In some instances, when a sled is inserted into any given one of sled spaces 603-1 to 603-5, the corresponding MPCM may couple with a counterpart MPCM of the inserted sled. This coupling
5 may provide the inserted sled with connectivity to both signaling infrastructure and power infrastructure of the rack in which it is housed.

Included among the types of sleds to be accommodated by rack architecture 600 may be one or more types of sleds that feature expansion capabilities. **FIG. 7** illustrates an example of a sled 704 that may be representative of a sled of such a type. As shown in FIG. 7, sled 704 may
10 comprise a set of physical resources 705, as well as an MPCM 716 designed to couple with a counterpart MPCM when sled 704 is inserted into a sled space such as any of sled spaces 603-1 to 603-5 of FIG. 6. Sled 704 may also feature an expansion connector 717. Expansion connector 717 may generally comprise a socket, slot, or other type of connection element that is capable of accepting one or more types of expansion modules, such as an expansion sled 718.
15 By coupling with a counterpart connector on expansion sled 718, expansion connector 717 may provide physical resources 705 with access to supplemental computing resources 705B residing on expansion sled 718. The embodiments are not limited in this context.

FIG. 8 illustrates an example of a rack architecture 800 that may be representative of a rack architecture that may be implemented in order to provide support for sleds featuring
20 expansion capabilities, such as sled 704 of FIG. 7. In the particular non-limiting example depicted in FIG. 8, rack architecture 800 includes seven sled spaces 803-1 to 803-7, which feature respective MPCMs 816-1 to 816-7. Sled spaces 803-1 to 803-7 include respective primary regions 803-1A to 803-7A and respective expansion regions 803-1B to 803-7B. With respect to each such sled space, when the corresponding MPCM is coupled with a counterpart
25 MPCM of an inserted sled, the primary region may generally constitute a region of the sled space that physically accommodates the inserted sled. The expansion region may generally constitute a region of the sled space that can physically accommodate an expansion module, such as expansion sled 718 of FIG. 7, in the event that the inserted sled is configured with such a module.

FIG. 9 illustrates an example of a rack 902 that may be representative of a rack
30 implemented according to rack architecture 800 of FIG. 8 according to some embodiments. In the particular non-limiting example depicted in FIG. 9, rack 902 features seven sled spaces 903-1 to 903-7, which include respective primary regions 903-1A to 903-7A and respective expansion regions 903-1B to 903-7B. In various embodiments, temperature control in rack 902 may be

implemented using an air cooling system. For example, as reflected in FIG. 9, rack 902 may feature a plurality of fans 919 that are generally arranged to provide air cooling within the various sled spaces 903-1 to 903-7. In some embodiments, the height of the sled space is greater than the conventional “1U” server height. In such embodiments, fans 919 may generally
5 comprise relatively slow, large diameter cooling fans as compared to fans used in conventional rack configurations. Running larger diameter cooling fans at lower speeds may increase fan lifetime relative to smaller diameter cooling fans running at higher speeds while still providing the same amount of cooling. The sleds are physically shallower than conventional rack dimensions. Further, components are arranged on each sled to reduce thermal shadowing (i.e.,
10 not arranged serially in the direction of air flow). As a result, the wider, shallower sleds allow for an increase in device performance because the devices can be operated at a higher thermal envelope (e.g., 250W) due to improved cooling (i.e., no thermal shadowing, more space between devices, more room for larger heat sinks, etc.).

MPCMs 916-1 to 916-7 may be configured to provide inserted sleds with access to power
15 sourced by respective power modules 920-1 to 920-7, each of which may draw power from an external power source 921. In various embodiments, external power source 921 may deliver alternating current (AC) power to rack 902, and power modules 920-1 to 920-7 may be configured to convert such AC power to direct current (DC) power to be sourced to inserted sleds. In some embodiments, for example, power modules 920-1 to 920-7 may be configured to
20 convert 277-volt AC power into 12-volt DC power for provision to inserted sleds via respective MPCMs 916-1 to 916-7. The embodiments are not limited to this example.

MPCMs 916-1 to 916-7 may also be arranged to provide inserted sleds with optical signaling connectivity to a dual-mode optical switching infrastructure 914, which may be the same as – or similar to – dual-mode optical switching infrastructure 514 of FIG. 5. In various
25 embodiments, optical connectors contained in MPCMs 916-1 to 916-7 may be designed to couple with counterpart optical connectors contained in MPCMs of inserted sleds to provide such sleds with optical signaling connectivity to dual-mode optical switching infrastructure 914 via respective lengths of optical cabling 922-1 to 922-7. In some embodiments, each such length of optical cabling may extend from its corresponding MPCM to an optical interconnect loom 923
30 that is external to the sled spaces of rack 902. In various embodiments, optical interconnect loom 923 may be arranged to pass through a support post or other type of load-bearing element of rack 902. The embodiments are not limited in this context. Because inserted sleds connect to an optical switching infrastructure via MPCMs, the resources typically spent in manually configuring the rack cabling to accommodate a newly inserted sled can be saved.

FIG. 10 illustrates an example of a sled 1004 that may be representative of a sled designed for use in conjunction with rack 902 of FIG. 9 according to some embodiments. Sled 1004 may feature an MPCM 1016 that comprises an optical connector 1016A and a power connector 1016B, and that is designed to couple with a counterpart MPCM of a sled space in conjunction with insertion of MPCM 1016 into that sled space. Coupling MPCM 1016 with such a counterpart MPCM may cause power connector 1016 to couple with a power connector comprised in the counterpart MPCM. This may generally enable physical resources 1005 of sled 1004 to source power from an external source, via power connector 1016 and power transmission media 1024 that conductively couples power connector 1016 to physical resources 1005.

Sled 1004 may also include dual-mode optical network interface circuitry 1026. Dual-mode optical network interface circuitry 1026 may generally comprise circuitry that is capable of communicating over optical signaling media according to each of multiple link-layer protocols supported by dual-mode optical switching infrastructure 914 of FIG. 9. In some embodiments, dual-mode optical network interface circuitry 1026 may be capable both of Ethernet protocol communications and of communications according to a second, high-performance protocol. In various embodiments, dual-mode optical network interface circuitry 1026 may include one or more optical transceiver modules 1027, each of which may be capable of transmitting and receiving optical signals over each of one or more optical channels. The embodiments are not limited in this context.

Coupling MPCM 1016 with a counterpart MPCM of a sled space in a given rack may cause optical connector 1016A to couple with an optical connector comprised in the counterpart MPCM. This may generally establish optical connectivity between optical cabling of the sled and dual-mode optical network interface circuitry 1026, via each of a set of optical channels 1025. Dual-mode optical network interface circuitry 1026 may communicate with the physical resources 1005 of sled 1004 via electrical signaling media 1028. In addition to the dimensions of the sleds and arrangement of components on the sleds to provide improved cooling and enable operation at a relatively higher thermal envelope (e.g., 250W), as described above with reference to FIG. 9, in some embodiments, a sled may include one or more additional features to facilitate air cooling, such as a heatpipe and/or heat sinks arranged to dissipate heat generated by physical resources 1005. It is worthy of note that although the example sled 1004 depicted in FIG. 10 does not feature an expansion connector, any given sled that features the design elements of sled 1004 may also feature an expansion connector according to some embodiments. The embodiments are not limited in this context.

FIG. 11 illustrates an example of a data center 1100 that may generally be representative of one in/for which one or more techniques described herein may be implemented according to various embodiments. As reflected in FIG. 11, a physical infrastructure management framework 1150A may be implemented to facilitate management of a physical infrastructure 1100A of data center 1100. In various embodiments, one function of physical infrastructure management framework 1150A may be to manage automated maintenance functions within data center 1100, such as the use of robotic maintenance equipment to service computing equipment within physical infrastructure 1100A. In some embodiments, physical infrastructure 1100A may feature an advanced telemetry system that performs telemetry reporting that is sufficiently robust to support remote automated management of physical infrastructure 1100A. In various embodiments, telemetry information provided by such an advanced telemetry system may support features such as failure prediction/prevention capabilities and capacity planning capabilities. In some embodiments, physical infrastructure management framework 1150A may also be configured to manage authentication of physical infrastructure components using hardware attestation techniques. For example, robots may verify the authenticity of components before installation by analyzing information collected from a radio frequency identification (RFID) tag associated with each component to be installed. The embodiments are not limited in this context.

As shown in FIG. 11, the physical infrastructure 1100A of data center 1100 may comprise an optical fabric 1112, which may include a dual-mode optical switching infrastructure 1114. Optical fabric 1112 and dual-mode optical switching infrastructure 1114 may be the same as – or similar to – optical fabric 412 of FIG. 4 and dual-mode optical switching infrastructure 514 of FIG. 5, respectively, and may provide high-bandwidth, low-latency, multi-protocol connectivity among sleds of data center 1100. As discussed above, with reference to FIG. 1, in various embodiments, the availability of such connectivity may make it feasible to disaggregate and dynamically pool resources such as accelerators, memory, and storage. In some embodiments, for example, one or more pooled accelerator sleds 1130 may be included among the physical infrastructure 1100A of data center 1100, each of which may comprise a pool of accelerator resources – such as co-processors and/or FPGAs, for example – that is available globally accessible to other sleds via optical fabric 1112 and dual-mode optical switching infrastructure 1114.

In another example, in various embodiments, one or more pooled storage sleds 1132 may be included among the physical infrastructure 1100A of data center 1100, each of which may comprise a pool of storage resources that is available globally accessible to other sleds via

optical fabric 1112 and dual-mode optical switching infrastructure 1114. In some embodiments, such pooled storage sleds 1132 may comprise pools of solid-state storage devices such as solid-state drives (SSDs). In various embodiments, one or more high-performance processing sleds 1134 may be included among the physical infrastructure 1100A of data center 1100. In some
5 embodiments, high-performance processing sleds 1134 may comprise pools of high-performance processors, as well as cooling features that enhance air cooling to yield a higher thermal envelope of up to 250W or more. In various embodiments, any given high-performance processing sled 1134 may feature an expansion connector 1117 that can accept a far memory expansion sled, such that the far memory that is locally available to that high-performance
10 processing sled 1134 is disaggregated from the processors and near memory comprised on that sled. In some embodiments, such a high-performance processing sled 1134 may be configured with far memory using an expansion sled that comprises low-latency SSD storage. The optical infrastructure allows for compute resources on one sled to utilize remote accelerator/FPGA, memory, and/or SSD resources that are disaggregated on a sled located on the same rack or any
15 other rack in the data center. The remote resources can be located one switch jump away or two-switch jumps away in the spine-leaf network architecture described above with reference to FIG. 5. The embodiments are not limited in this context.

In various embodiments, one or more layers of abstraction may be applied to the physical resources of physical infrastructure 1100A in order to define a virtual infrastructure, such as a
20 software-defined infrastructure 1100B. In some embodiments, virtual computing resources 1136 of software-defined infrastructure 1100B may be allocated to support the provision of cloud services 1140. In various embodiments, particular sets of virtual computing resources 1136 may be grouped for provision to cloud services 1140 in the form of SDI services 1138. Examples of cloud services 1140 may include – without limitation – software as a service (SaaS) services
25 1142, platform as a service (PaaS) services 1144, and infrastructure as a service (IaaS) services 1146.

In some embodiments, management of software-defined infrastructure 1100B may be conducted using a virtual infrastructure management framework 1150B. In various
embodiments, virtual infrastructure management framework 1150B may be designed to
30 implement workload fingerprinting techniques and/or machine-learning techniques in conjunction with managing allocation of virtual computing resources 1136 and/or SDI services 1138 to cloud services 1140. In some embodiments, virtual infrastructure management framework 1150B may use/consult telemetry data in conjunction with performing such resource allocation. In various embodiments, an application/service management framework 1150C may

be implemented in order to provide QoS management capabilities for cloud services 1140. The embodiments are not limited in this context.

FIG. 12 illustrates an example of a data center 1200 that may generally be representative of a data center or other type of computing network in/for which one or more techniques described herein may be implemented according to various embodiments. As shown in FIG. 12, the data center 1200 may be similar to and include features and components previously discussed. For example, the data center 1200 may generally contain a plurality of racks 1202A to 1202D, each of which may house computing equipment including a respective set of physical resources 1205A- x to 1205D- x , where x may be any positive integer from 1 to 4. The physical resources 1205 may be contained within a number of sleds 1204A through 1204D. As mentioned, the physical resources 1205 may include resources of multiple types, such as – for example – processors, co-processors, fully-programmable gate arrays (FPGAs), memory, accelerators, and storage. Moreover, the physical resources 1205 may be physical memory resources; physical compute resources, physical storage resources, physical accelerator resources, etc.

In embodiments, the physical resources 1205 may be pooled within racks and between racks and provided as a composed node to process workloads. For example, physical resources 1205A-1 of sled 1204A-1 may be pooled with physical resources 1205A-3 of sled 1204A-3 to provide combined processing capabilities for workloads across sleds within the same rack, *e.g.* rack 1202A. Similarly, physical resources 1205 of one or more racks may be combined with physical resources of one or more other racks to create a pool of physical resources to process a workload as a composed node. In one example, the physical resources 1205A-3 may be combined and pooled with physical resources of 1205B-1, which are located within rack 1202A and rack 1202B, respectively. Any combination of physical resources 1205 may be pooled together as a composed node to process workloads and embodiments are not limited in this manner. Moreover, some embodiments may include more or less physical resources 1205, sleds 1204, and racks 1202 and the illustrated example should not be construed in a limiting manner.

In the illustrated example of FIG. 12, the data center 1200 may provide verification and authentication capabilities for the physical resources 1205. The verification and authentication capabilities include determining that each of the physical resources 1205 are the original equipment manufacturer (OEM) components that are intended to be in a sled 1204 based on a sled manifest distributed with the sled 1204 and operating according to their intended function, for example. The sled manifest may include an identifier for the sled 1204, inventory (identifiers) of the physical resources 1205 of the sled 1204, hardware authentication information

for each physical resource 1205 of a sled 1204, and software/firmware validation information for each of the physical resource 1205 of a sled 1204.

In embodiments, each of the sleds 1204 may perform verification and authentication operations using the sled manifest and generate results indicating success or failure of the operations. Each sled 1204 may provide the results and sled manifest to a pod management controller 1231 via one or more secure links. As will be discussed in more detail below, each sled 1204 may include a verification controller, to ensure that each physical resource 1205 is authenticated and is operating as intended based on hardware and software authentication and verification. Further, the verification controller may first verify the sled manifest. The verification controller may use and access a hash value of a sled manifest used as part of the results from a secure location. The verification controller may compare the hash value of the sled manifest stored in the secure location with a generated hash value sled manifest to verify the sled manifest. Once verified, the verification controller may authenticate and verify the physical resources 1205 and generate results. These results, which may include a hash value of the sled manifest used for authentication, may be signed by the verification controller using a private key and provided to the pod management controller 1231 through one or more secure links. The secure links may utilize one or more secure protocols, *e.g.* a transport layer security (TLS) protocol, a secure sockets layer (SSL) protocol, a private communications transport (PCT) protocol, and so forth.

The pod management controller 1231 may receive the results and sled manifest and validate them to ensure that the results are the actual results of the operations performed and the sled manifest is the actual sled manifest from an authentic sled 1204. In an example, the pod management controller 1231 may ensure the sled manifest received from the authenticated sled 1204 is authentic by utilizing a digital signature. The manufacturer of sled 1204 (or any other authorized party) may sign the sled manifest with a private key and the pod management controller 1231 may validate the digital signature using a public key that may be obtained from a trusted third party or from a public key certificate signed by a trusted CA.

The pod management controller 1231 may also verify the integrity of the sled manifest contents using a hash value. For example, the sled manifest signature may contain a hash value of the sled manifest. The pod management controller 1231 can itself calculate a hash value of the sled manifest and compare it to the hash value included with the signature. If the values match, the sled manifest has not been modified in any way from the time the signature was created. This hash value of the sled manifest can also be compared with the received hash value from the sled 1204, which is part of the results from the sled verification. If the values match, then the pod

manager controller 1231 has an assurance that this same sled manifest was used for sled verification by the verification controller. Moreover, using hashing techniques may ensure that the sled manifest is not corrupted during communication to the pod management controller 1231.

Similarly, the pod management controller 1231 may authenticate the results of the operations performed on the sled 1204 utilizing a signature of a private key of a verification controller. The results may be signed by the verification controller. The pod management controller 1231 may authenticate that the results were generated by the verification controller using a public key, which may be obtained from the sled manifest and a trusted third party.

In some instances, the pod management controller 1231 may verify the integrity of the content of the results based on a signature. As similarly discussed above, a hash value of the results may be generated by the verification controller (or a sled) and may be included as part of the signature of those results. The pod management controller 1231 may generate its hash value of the received results and perform a comparison with the hash included in the signature of the results. If the hash values match the pod management controller 1231 can trust that the results are authentic (i.e., generated by a legitimate verification controller in the target sled).

The pod management controller 1231 may enable the physical resources 1205 of a sled 1204 to process a workload and data based on the outcome of the authentication and verification of the sled manifest and the results performed on the sled 1204. For example, if the sled manifest and the results are authenticated, and verified by the pod management controller 1231, the pod management controller 1231 may enable the physical resources of the sled to be used to compose a node. The pod management controller 1231 may add the inventory of the physical resources 1205 to a database used to generate a composed node, for example. However, the pod management controller 1231 may prevent physical resources 1205 from being utilized if at least one of the results and the sled manifest are not authenticated and verified.

FIG. 13A illustrates an example of a sled 1304 that may be representative of a sled designed for use in conjunction with the racks discussed herein, for example. In embodiments, sled 1304 may have similar components and functionality as sled 1004 discussed in FIG. 10. Sled 1304 may feature physical resources 1305 including physical memory resources 1305-1, physical compute resources 1305-2, physical storage resources 1305-3, and physical accelerator resources 1305-4. Moreover, these physical resources 1305 may be processors, co-processors, fully-programmable gate arrays (FPGAs), memory, accelerators, and storage. The sled 1304 may also include a non-volatile memory 1354 having a sled manifest 1356, a verification controller 1352, a management controller 1352, and an interface 1364 coupled with a secure link.

The non-volatile memory 1354 having the sled manifest 1356 may be a secure memory such that non-volatile memory 1354 cannot be changed or only be changed with an appropriate access (write) privileges. In embodiments, the non-volatile memory 1354 may be incorporated in or be part of a processing unit, such as a CPU and a trusted platform component such as an
5 FPGA, CPLD, and an ASIC, or may be a separate non-volatile memory. The sled manifest 1356 may also be updated when valid and authenticate resources are used to replace other resources, *e.g.* due to a failure or performance upgrade.

In embodiments, the sled manifest 1356 may include information about the physical resources 1305 to ensure that they are authentic and operate as designed/intended. For example,
10 the sled manifest 1356 may include hardware authentication information for each of the physical resources 1305 and software/firmware validation information for the physical resources 1305. The hardware authentication information may include a public key that may be used to verify the physical resource 1305. As will be discussed in more detail below, the public key may be used to authenticate a signed nonce or certificate signed by the hardware of the physical
15 resources 1305 using a secure private key. Similarly, the sled manifest 1356 may include public keys of the firmware/software associated with the physical resources 1305. In some instances, the sled manifest 1356 may include hash values previously generated based on the firmware/software that may be used to authenticate and validate the firmware and software. The sled manifest 1356 may also include other information such as identification information to
20 uniquely identify the sled 1304, other identification information to identify each of the physical resources 1305, details about minimal firmware and security revisions, and datacenter rules and policies.

In embodiments, the verification controller 1352 of the sled 1304 may utilize the sled manifest 1356 to authenticate and verify the physical resources 1305. The verification controller
25 1352 may include secure circuitry and process one or more instructions stored in a secure memory, such as non-volatile memory, to validate/authenticate the sled manifest 1356 and the physical resources 1305. The one or more instructions may cause the verification controller 1352 to authenticate and verify the physical resources 1305 when a sled is first inserted in a rack, for example. The verification controller 1352 may also process one or more instructions to verify
30 and authenticate the physical resources 1305 when a change occurs to a physical resource 1305, *e.g.* a resource is replaced.

In some embodiments, the verification controller 1352 may operate when the sled 1304 is first plugged/inserted into a rack and before other hardware/elements, *e.g.* when the management controller 1362 begins processing information and data. Moreover, the verification controller

1352 may authenticate and verify the physical resources 1305 before they are enabled and presented to a pod management controller as possible resources for use in composed nodes by the management controller 1362.

In embodiments, the verification controller 1352 may retrieve/receive the sled manifest 5 1356 from the non-volatile memory 1354 authenticate the sled manifest and perform the authentication and verification operations. To authenticate, the sled manifest 1356, the verification controller 1352 may compare a hash value originally generated by a manufacturer and stored in a secure location of non-volatile memory with a hash value the verification controller 1352 generates based on the sled manifest. The non-volatile memory storing the hash 10 value of the sled manifest 1356 may be different or the same as non-volatile memory 1254. If the original hash value matches the generated hash value, the verification controller 1352 may verify the sled manifest 1356 and perform authentication and verification operations using the sled manifest 1356.

In embodiments, the authentication and verification operations include the verification 15 controller 1352 generating a nonce and sending the nonce to a physical resource 1305. The physical resource 1305 may sign the nonce with a private key which it stores securely and is provided by the manufacturer of the resource 1305, for example. The verification controller 1352 may receive the nonce signed with the private key of the physical resource 1305 from the physical resource 1305 and determine whether the physical resource 1305 is authentic. More 20 specifically, the verification controller 1352 may use the public key based on a certificate in the sled manifest 1356 associated with the physical resource 1305 to authenticate the signed nonce from the physical resource 1305. The public key may be used to verify the signed nonce and if the verification succeeds the verification controller 1352 may authenticate the physical resource 1305. If it fails, the verification controller 1352 may indicate that the physical resource is not 25 authenticated. The verification controller 1352 may authenticate each of the physical resources 1305 of the sled in this manner. Embodiments are not limited in this manner. The physical resources 1305 In some embodiments, the verification controller 1352 may utilize a cryptographic query using Intel's® Enhanced Privacy Identification (EPID).

The verification controller 1352 may also verify the firmware/software for each of the 30 physical resources 1305. For example, the verification controller 1352 may receive a hash value generated based on the firmware/software from a physical resource 1305. In some instances, the hash value received from the physical resource 1305 may be generated based on a request from the verification controller 1352. The verification controller 1352 may compare the hash value received from the physical resource 1305 with a hash value stored in the sled manifest 1356. The

stored hash value may have previously been generated based on the firmware/software, at the time of manufacturer, for example. If the hash values match, the verification controller 1352 may verify the firmware/software associated with the physical resource 1305. Alternatively, if the hash values do not match, the verification controller 1352 may not verify the firmware/software for the physical resource 1305. The verification controller 1352 may perform this operation for each of the physical resources having firmware/software.

In embodiments, the verification controller 1352 may perform authentication and verification operations for each of the physical resources 1305 of the sled 1304 and generate results of the operations. The results may indicate whether each of the authentication and verification of the physical resources 1305 was successful or not successful. In some embodiments, the verification controller 1352 may sign the results with a private key such that they may be authenticated by a pod management controller, as previously discussed. Once the verification controller 1352 has completed the authentication and verification operations, the sled 1304 may continue its initialization process, which may include one or more operations performed by the management controller 1362. The verification controller 1352 may also detect firmware corruption that may lead to malicious attacks trying to control of the platform. Once detected, the verification controller 1352 may recover verified firmware to prevent these types of attacks.

The management controller 1362, which may be a specialized microcontroller, such as a baseboard management controller or sled management controller, may be embedded in a motherboard of the sled 1304 to interface between system management software and platform hardware, *e.g.* physical resources 1305. Further, the management controller 1362 may also interface with other components of a data center, such as a pod management controller and a rack management controller. For example, the management controller 1362 may receive/retrieve the results of the authentication and verification operations performed by the verification controller 1352 and the sled manifest 156 to communicate to pod or rack management controller via the interface 1364.

The management controller 1362 may also establish a secure link with a rack management controller and a pod management controller to communicate information, such as the results and the sled manifest 1356, via the interface 1364. In embodiments, the management controller 1362 may establish the secure link with a pod management controller using one or more protocols, such as a transport layer security (TLS) protocol, a secure sockets layer (SSL) protocol, a private communications transport (PCT) protocol, and so forth. Once the secure link is established, the management controller 1362 may communicate the results which may be

signed by the verification controller 1352 to a pod management controller via the interface 1364. The management controller 1262 may also send the sled manifest 1356, which may also have been signed by the manufacturer of the sled or an authorized third party, for example. As previously discussed, a pod management controller may receive the signed results and sled
5 manifest 1356 and perform its authentication and verification before enabling the physical resources 1305 to be used in a composed node.

FIG. 13B illustrates an example of a distribution and attestation flow 1320 to ensure that devices, such as physical resources 1305 to detect counterfeit and improper devices.

Embodiments include support for securely maintaining a unique private/public key pair and
10 perform cryptographic functions (message signing), as previously discussed. Embodiments also include maintaining platform or sled support to maintain a set of root certificates and platform policies in a manifest, such as a sled manifest 1356.

These features may be provided by one or more certificate authorities 1322 for devices and resources, such as a device vendor or OEM, and an integrator 1234 to integrate physical
15 resources 1305 into a platform or sled 1304 to provide storage and processing capabilities. Embodiments are not limited in this manner, and other trusted entities may be involved to enable sure and authenticate resources.

In embodiments, the certificate authorities 1322 may generate and maintain a root key pair and root certificate for their devices or physical resources 1305, which may be provided to
20 the integrator 1324. Each certificate authority 1322 or vendor may maintain and provide their own unique root key pair and root certificate for the resources they manufacture and are implemented in a platform or sled 1304, for example. The integrator 1324 may collect all of the root key pairs and root certificates for resources 1305 to be implemented in a sled 1304, for example. The integrator 1324 may generate the sled manifest 1356 based on these root key pairs
25 and root certificates. The root key pairs and root certificates may include the public keys used to verify and authenticate the signed nonce(s) received by the physical resources 1305, as previously discussed. The sled manifest 1356 may also be signed by the integrator 1324 and provisioned secure non-volatile memory, such as non-volatile memory 1354, or in trusted component, such as a trusted platform module, a platform resilience technology device, and so
30 forth. The integrator 1324 may also provision a unique platform key pair in the secure non-volatile memory or trusted component.

The certificate authorities 1322 or vendors may also manufacture devices and generate unique per device private/public key pairs and device certificates signed by the root private key. The certificate authorities 1322 may provision a physical resource 1305, securely, with the

private/public key pair (device key pair). For example, the physical resource 1305 may maintain the device key pair in a secure, non-volatile memory, hardware fuses, and so forth. A device certificate associated with a physical resource 1305 may also be provisioned with the physical resource 1305. The physical resources 1305 including the device key pair and the device
5 certificate may be provided to the integrator 1324 to be implemented in a sled 1304.

In embodiments, a sled 1304 may utilize the sled manifest 1356 to authenticate and validate each of its' physical resources 1305 during a boot or reboot cycle, as previously mentioned. For example, a sled 1304 including circuitry, such as the verification controller 1352 or the like, may retrieve the sled manifest 1356 and may verify the sled manifest's signature.
10 Moreover, verification controller 1352 may discover all of the physical resources 1305 attached and coupled to the sled 1304. The verification controller 1352 may retrieve the device certificates for each of the physical resources 1305, for example. The verification controller 1352 may verify the device certificates utilizing the associated root keys and root certificates.

In embodiments, the verification controller 1352 generates a random nonce for every
15 physical resource 1305 and sends the nonce and a request for device details (hardware stepping, version, firmware revision, firmware measurement, etc.) to each of the physical resources 1305. The physical resources 1305 may collect the requested details and attaches the random nonce. Each of physical resources 1305 may sign the response, including the requested details and nonce, with the physical resources 1305 device private key and sends the response to the
20 verification controller 1352. The verification controller 1352 receives the response confirms the random nonce, verifies the signature using the public key, and performs a compliance check against policies store in the sled manifest 1356. The platform policies stored in the sled manifest 1356 may control the sled's actions to non-compliant devices, *e.g.* physical resources 1305 that failed verification, physical resources 1305 that were verified, but have a deficient firmware
25 version, etc.. The sled 1304 may take an action, such as reporting on the physical resources 1305, disabling the link between the sled 1304 and the physical resource 1305, disable the physical resource 1305 itself, preventing the sled 1304 from completing its boot cycle, and preventing the physical resource 1305 to be used as part of a composed node, for example. The sled 1304 may send the results and the sled manifest to a datacenter management entity, such as a pod
30 management controller 1331, as discussed herein.

FIG. 14 illustrates another example of a data center 1400 that may be representative of a data center or other type of computing network in/for which one or more techniques described herein may be implemented according to various embodiments. As shown in FIG. 14, the data center 1400 may be similar to and include features and components previously discussed. For

example, the data center 1400 may generally contain elements and component similar as those discussed in FIG. 12.

In embodiments, the physical resources 1405 may be pooled within racks and between racks and provided as a composed node to process workloads. In embodiments, the pod management controller 1431 may determine physical resources 1405 that are available to
5 generate composed nodes based on whether they are authenticated and verified as previously discussed. For example, the pod management controller 1431 may authenticate each of the results and the sled manifest using public keys that it may have received from a trusted third party or an OEM. The pod management controller 1431 may use the public key associated with a
10 private key used to sign the results to authenticate the results. Similarly, the pod management controller 1431 may use a public key associated with a private key used to sign the sled manifest to authenticate the sled manifest.

The pod management controller 1431 may also validate the results and the sled manifest using hash values. For example, the pod management controller 1431 may receive a hash value
15 for the results from a sled 1404, generate a hash value for the results and compare the received hash value and the generated hash value to determine whether they match. If they match, the pod management controller 1431 may determine that the content of the results have not been corrupted or changed.

~~The~~ pod management controller 1431 may receive a hash value as part of the signature
20 and with the sled manifest. The pod management controller 1431 may generate a hash value of the sled manifest, compare the hash value in the manifest signature with the generated hash value to determine whether they match or not. If they match, the pod management controller 1431 may determine that the sled manifest has not changed or been corrupted. Embodiments are not limited in this manner.

The pod management controller 1431 may enable the physical resources 1405 that are
25 successfully authenticated and validated by the sled 1404 and confirmed by the pod management controller 1431 to be used in a composed node, *e.g.* a clustering of physical resources to process a workload based on various requirements. Thus, the pod management controller 1431 ensures that the physical resources 1405 for use as a composed node is an OEM provided resource and is
30 operating according to its intended function.

In embodiments, the pod management controller 1431 may generate a composed node based on one or more requirements stipulated by a user or client 1433 to process a workload. These requirements may include performance requirements, *e.g.* processing requirements, memory requirements, network requirements, and so forth. Other requirements may also be

stipulated. For example, a user or client 1433 may specify that a workload is processed via manufacturer(s) of the physical resource 1405. For example, a user may require that a workload is processed on Intel® Corp. processing devices. Thus, the pod management controller 1431 may use these requirements to generate one or more composed nodes to process the workload. A
5 composed node may include one or more physical resources that may or may not be within the same sled 1404 and rack 1402. As previously discussed, a composed node may include physical resources 1405 within the same rack 1402 and from different racks 1402.

The pod management controller 1431 may also generate a platform certificate request to have a platform certificate generated and signed for a composed node by the certificate authority
10 1435. The platform certificate may include information from each of the sleds manifests associated with physical resources used to generate a composed node. Moreover, the platform certificate may include a listing of the physical resources that make up the composed node, for example.

The certificate authority 1435 may receive the platform certificate request and generate
15 the platform certificate. The certificate authority 1435 may sign the platform certificate with a private key, which may be used to authenticate the platform certificate. The platform certificate may be communicated to the pod management controller 1431 and the sleds 1404 having the physical resources 1405 used in the composed node. In some instances, a client 1433 may request the platform certificate from a data center to ensure that the physical resources 1405 meet
20 the requirements to process a workload. The pod management controller 1431 may send the signed platform certificate to a client 1433. In embodiments, the client 1433 may authenticate the platform certificate using a public key of the certificate authority 1435, for example.

FIG. 15 illustrates an embodiment of logic flow 1500. The logic flow 1500 may be representative of some or all the operations executed by one or more embodiments described
25 herein. For example, the logic flow 1500 may illustrate operations performed by a verification controller to authenticate and verify physical resources of a sled. However, embodiments are not limited in this, and one or more operations may be performed by other components or systems discussed herein.

At block 1502, the logic flow 1500 includes detecting a system initialization. The system
30 initialization may be generated when a sled is plugged or inserted into a sled slot of a rack. As part of the system initialization, power may be applied to various elements of a sled including a management controller, a verification controller, non-volatile memory, and physical resources. In embodiments, the verification controller may authenticate and verify the physical resources of a system before the system completely initializes. If one or more of the physical resources cannot

be authenticated or validated, the verification controller may prevent a sled from completing the initialization process.

At block 1504, the logic flow 1500 include retrieve (or receiving) a sled manifest from a non-volatile (secure) memory. The sled manifest includes identification, authentication, and verification information for a sled and the physical resources of the sled. At block 1506, the sled manifest may be verified to ensure the integrity of the sled manifest and its contents. For example, a verification controller may calculate a hash value based on the sled manifest and compare it with a previously generated and verified hash value stored in a secure location or hardware. In some instances, the verified hash value may have been generated by a manufacturer of the sled at the time of manufacturer or when a change is made to the sled. Embodiments are not limited to this example, and in some instances, the verified hash value may be generated when a change occurs on the sled, for example.

At block 1508, the logic flow 1500 includes generating a nonce, which may be a randomly generated number used once to authenticate a physical resource. Each physical resource may be authenticated with a different nonce, for example. At block 1510, the verification controller may send the nonce to a physical resource. The physical resource may store a private key in a secure hardware location to sign the nonce Further, and at block 1512, the verification controller may receive the signed nonce from the physical resource. The verification controller may authenticate and verify the physical resource using a public key in the sled manifest. More specifically, the verification controller may determine whether the nonce communicated to the physical resource is the same as the signed nonce received from the physical resource. If they match, the verification controller may authenticate the physical resource at block 1514. If the physical resource is not successfully authenticated, the verification controller may prevent a sled from completing its initialization process at block 1518.

If a physical resource is successfully authenticated, the verification controller may determine whether additional physical resources exist to authenticate at decision block 1516. If no additional physical resources exist, the verification may generate results that may be sent to a pod management controller at block 1520. The results may indicate each of the physical resources and an indication as to whether they were authenticated or not authenticated. In some embodiments, the verification controller may sign the results with a private key that it stores in a secure hardware location. If additional physical resources exist, the logic flow 1500 may repeat to authenticate each of the physical resources.

FIG. 16 illustrates an embodiment of logic flow 1600. The logic flow 1600 may be representative of some or all the operations executed by one or more embodiments described

herein. For example, the logic flow 1600 may illustrate operations performed a pod management controller to authenticate an verify results and a sled manifest. However, embodiments are not limited in this, and one or more operations may be performed by other components or systems discussed herein.

5 At block 1602, the logic flow 1600 includes receiving a sled manifest and results of an authentication and validation operations from a sled. The sled manifest may include information to identify the physical resources of the sled and used to authenticate/validate the physical resources. The results may be generated by a verification controller, as previously discussed, and include whether the authentication and the verification of the physical resources were successful
10 or not successful. The results also include a hash value of the sled manifest. Note that in embodiments the results and the sled manifest may not be communicated in the same communication, but may be communicated in different communications in one or more frames, for example. Embodiments are not limited in this manner.

 At block 1604, the logic flow 1600 may include authenticating and validating the sled
15 manifest received from a sled. More specifically, the pod management controller may authenticate the sled manifest using a public key corresponding to a private key used to sign the sled manifest by the sled, *e.g.* from a manufacturer or trusted third party. In some embodiments, the signature used to sign the sled manifest may include a hash value of the sled manifest. The pod management controller may validate the sled manifest via comparing the hash value in the
20 signature with a hash value generated by the pod management controller. If the hash values match, the sled manifest may be authenticated. However, if they do not match, the physical resources associated with the sled manifest may be prevented from be used in a composed node at block 1610.

 At block 1606, the logic flow 1600 includes determining whether the results are authentic
25 and are generated by the appropriate verification controller. For example, the pod management controller may authenticate the results via a private/public key combination. More specifically, the pod management controller may authenticate the results with a public key in the manifest and may be associated with a private key of a verification controller. If the results are successfully verified with the public key, they are authenticated.

30 The pod management controller may also validate the results using hash values. The pod management controller may generate a hash value based on the results and compare the hash value with a hash that is included in the signature of the results. If the hash values match, the pod management controller may validate that the results have not been changed.

The pod management controller may use the authenticated and validated results to ensure that the authentication and validation operations performed on the sled were successful. If the results cannot be authenticated, validated, or do not indicate the physical resources were successfully authenticated/validated, the pod management controller may not permit the physical resources of the sled to be used in composed nodes at block 1610.

If the pod management controller cannot authenticate or validate the sled manifest, the pod management controller may not permit the physical resources of the sled to be used in composed nodes at block 1610. If the pod management controller can authenticate and validate the sled manifest, the pod management controller may enable the physical resources to be used as part a composed node. The pod management controller may add the physical resources to a database used to generate composed nodes, for example.

FIG. 17 illustrates an embodiment of logic flow 1700. The logic flow 1700 may be representative of some or all the operations executed by one or more embodiments described herein. For example, the logic flow 1700 may illustrate operations performed a pod management controller. However, embodiments are not limited in this, and one or more operations may be performed by other components or systems discussed herein.

At block 1702, the logic flow 1700 may include determining a composed node to process a workload. The composed node may include one or more physical resources to process the workload. In some embodiments, the pod management controller may generate the composed node based on requirements for the workload, as previously discussed.

At block 1704, the logic flow 1700 includes determining a platform certificate for the composed node. The platform certificate may be a listing of the physical resources used to generate the composed node. In some embodiments, the pod management controller may send a platform certificate request to a certificate authority to sign the platform certificate.

At block 1706, embodiments include receiving the signed platform certificate from a certificate authority. Further and at block 1708 embodiments including causing the platform certificate to be sent to a client. In other instances, the pod management controller may receive the signed platform certificate from the certificate authority and send it to the client, for example. At block 1710, the logic flow includes causing the workload to be processed by the composed node.

The detailed disclosure now turns to providing examples that pertain to further embodiments. Examples one through thirty provided below are intended to be exemplary and non-limiting.

In a first example, a system, a device, an apparatus, and so forth may include a memory; and a pod management controller coupled to the memory. The pod management controller may determine whether a sled manifest and results are authentic and valid, the results generated by an authentication and validation operation performed to authenticate and validate physical resources
5 of a sled, determine whether the results of the authentication and validation operation indicate the physical resources are authentic and valid. The pod management controller may permit the physical resources to process a workload in response to a determination that the sled manifest and results are authentic and valid, and the results indicate the physical resources are authentic and valid and prevent the physical resources from processing the workload in response to a
10 determination that at least one of the sled manifest is not authentic and valid, the results are not authentic and valid, and the results indicate the physical resources are not authentic and valid.

In a second example and in furtherance of the first example, a system, a device, an apparatus, and so forth including the pod management controller to receive the sled manifest comprising identifiers for the physical resources from the sled, and receive the results of the
15 authentication and validation operation performed to authenticate and validate the physical resources from the sled.

In a third example and in furtherance of any previous example, a system, a device, an apparatus, and so forth including the pod management controller to authenticate the sled
20 manifest using a public key, the public key obtained from an original manufacturer of the sled or a trusted third party.

In a fourth example and in furtherance of any previous example, a system, a device, an apparatus, and so forth including the pod management controller to generate a hash value based on the sled manifest, compare the hash value with another hash value communicated with the
25 results, validate the sled manifest if the hash value and the another hash value match, and invalidate the sled manifest if the hash value and the another hash value do not match.

In a fifth example and in furtherance of any previous example, a system, a device, an apparatus, and so forth including the pod management controller to authenticate the results using a
30 public key, the public key obtained from the sled manifest.

In a sixth example and in furtherance of any previous example, a system, a device, an apparatus, and so forth including the pod management controller to add an indication of the
physical resources in a database to permit the physical resources to process the workload.

In a seventh example and in furtherance of any previous example, a system, a device, an apparatus, and so forth including the pod management controller to receive the sled manifest and
35 the results of the authentication and verification operations via a secure link with the sled.

In an eighth example and in furtherance of any previous example, a system, a device, an apparatus, and so forth including the pod management controller to generate a composed node comprising at least one of the physical resources of the sled.

5 In a ninth example and in furtherance of any previous example, a system, a device, an apparatus, and so forth including the pod management controller to send a platform certificate request to generate a platform certificate for the composed node to a certificate authority, the platform certificate to identify physical resources including the at least one of the physical resources used for the composed node, and receive a signed platform certificate.

10 In a tenth example and in furtherance of any previous example, a system, a device, an apparatus, and so forth including the pod management controller to send the signed platform certificate for the composed node to a client device.

In an eleventh example and in furtherance of any previous example, a non-transitory computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to determine whether a sled manifest and results are authentic and valid, the results generated by an authentication and validation operation performed to
15 authenticate and validate physical resources of a sled, determine whether the results of the authentication and validation operation indicate the physical resources are authentic and valid, permit the physical resources to process a workload in response to a determination that the sled manifest and results are authentic and valid, and the results indicate the physical resources are
20 authentic and valid, and prevent the physical resources from processing the workload in response to a determination that at least one of the sled manifest is not authentic and valid, the results are not authentic and valid, and the results indicate the physical resources are not authentic and valid.

In a twelfth example and in furtherance of any previous example, a non-transitory
25 computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to receive the sled manifest comprising identifiers for the physical resources from the sled, and receive the results of the authentication and validation operation performed to authenticate and validate the physical resources from the sled.

In a thirteenth example and in furtherance of any previous example, a non-transitory
30 computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to authenticate the sled manifest using a public key, the public key obtained from an original manufacturer of the sled or a trusted third party.

In a fourteenth example and in furtherance of any previous example, a non-transitory
35 computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to generate a hash value based on the sled manifest, compare the hash

value with another hash value communicated with the results, validate the sled manifest if the hash value and the another hash value match, and invalidate the sled manifest if the hash value and the another hash value do not match.

5 In a fifteenth example and in furtherance of any previous example, a non-transitory computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to authenticate the results using a public key, the public key obtained from the sled manifest.

10 In a sixteenth example and in furtherance of any previous example, a non-transitory computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to add an indication of the physical resources in a database to permit the physical resources to process the workload.

15 In a seventeenth example and in furtherance of any previous example, a non-transitory computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to receive the sled manifest and the results of the authentication and verification operations via a secure link with the sled.

In an eighteenth example and in furtherance of any previous example, a non-transitory computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to generate a composed node comprising at least one of the physical resources of the sled.

20 In a nineteenth example and in furtherance of any previous example, a non-transitory computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to send a platform certificate request to generate a platform certificate for the composed node to a certificate authority, the platform certificate to identify physical resources including the at least one of the physical resources used for the composed node, and
25 receive a signed platform certificate.

In a twentieth example and in furtherance of any previous example, a non-transitory computer-readable storage medium, comprising a plurality of instructions, that when executed, enable processing circuitry to send the signed platform certificate for the composed node to a
30 client device.

In a twenty-first example and in furtherance of any previous example, a computer-implemented method may include determining whether a sled manifest and results are authentic and valid, the results generated by an authentication and validation operation performed to authenticate and validate physical resources of a sled, determining whether the results of the

authentication and validation operation indicate the physical resources are authentic and valid, permitting the physical resources to process a workload in response to a determination that the sled manifest and results are authentic and valid, and the results indicate the physical resources are authentic and valid, and preventing the physical resources from processing the workload in response to a determination that at least one of the sled manifest is not authentic and valid, the results are not authentic and valid, and the results indicate the physical resources are not authentic and valid.

In a twenty-second example and in furtherance of any previous example, a computer-implemented method may include receiving the sled manifest comprising identifiers for the physical resources from the sled, and receiving the results of the authentication and validation operation performed to authenticate and validate the physical resources from the sled.

In a twenty-third example and in furtherance of any previous example, a computer-implemented method may include authenticating the sled manifest using a public key, the public key obtained from an original manufacturer of the sled or a trusted third party.

In a twenty-fourth example and in furtherance of any previous example, a computer-implemented method may include generating a hash value based on the sled manifest, comparing the hash value with another hash value communicated with the results, validating the sled manifest if the hash value and the another hash value match, and invalidating the sled manifest if the hash value and the another hash value do not match.

In a twenty-fifth example and in furtherance of any previous example, a computer-implemented method may include authenticating the results using a public key, the public key obtained from the sled manifest.

In a twenty-sixth example and in furtherance of any previous example, a computer-implemented method may include adding an indication of the physical resources in a database to permit the physical resources to process the workload.

In a twenty-seventh example and in furtherance of any previous example, a computer-implemented method may include receiving the sled manifest and the results of the authentication and verification operations via a secure link with the sled.

In a twenty-eighth example and in furtherance of any previous example, a computer-implemented method may include generating a composed node comprising at least one of the physical resources of the sled.

In a twenty-ninth example and in furtherance of any previous example, a computer-implemented method may include sending a platform certificate request to generate a platform certificate for the composed node to a certificate authority, the platform certificate to identify

physical resources including the at least one of the physical resources used for the composed node, and receiving a signed platform certificate.

In a thirtieth example and in furtherance of any previous example, a computer-implemented method may include sending the signed platform certificate for the composed node
5 to a client device.

Some embodiments may be described using the expression “one embodiment” or “an embodiment” along with their derivatives. These terms mean that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one
10 embodiment. The appearances of the phrase “in one embodiment” in various places in the specification are not necessarily all referring to the same embodiment. Further, some embodiments may be described using the expression “coupled” and “connected” along with their derivatives. These terms are not necessarily intended as synonyms for each other. For example, some embodiments may be described using the terms “connected” and “coupled” to indicate that
15 two or more elements are in direct physical or electrical contact with each other. The term “coupled,” however, may also mean that two or more elements are not in direct contact with each other, but yet still co-operate or interact with each other.

It is emphasized that the Abstract of the Disclosure is provided to allow a reader to quickly ascertain the nature of the technical disclosure. It is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. Also, in the
20 preceding Detailed Description, it can be seen that various features are grouped together in a single embodiment for the purpose of streamlining the disclosure. This method of disclosure is not to be interpreted as reflecting an intention that the claimed embodiments require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter lies in less than all features of a single disclosed embodiment. Thus the
25 following claims are at this moment incorporated into the Detailed Description, with each claim standing on its own as a separate embodiment. In the appended claims, the terms “including” and “in which” are used as the plain-English equivalents of the respective terms “comprising” and “wherein,” respectively. Moreover, the terms “first,” “second,” “third,” and so forth, are used merely as labels, and are not intended to impose numerical requirements on their objects.

30 What has been described above includes examples of the disclosed architecture. It is, of course, not possible to describe every conceivable combination of components and/or methodologies, but one of ordinary skill in the art may recognize that many further combinations and permutations are possible. Accordingly, the novel architecture is intended to embrace all

such alterations, modifications and variations that fall within the spirit and scope of the appended claims.

LISTING OF CLAIMS

What is claimed is:

1. An apparatus, comprising:
a memory; and
5 a pod management controller coupled with the memory, the pod management controller
to:
determine whether a sled manifest and results are authentic and valid, the results
generated by an authentication and validation operation performed to authenticate and
validate physical resources of a sled;
10 determine whether the results of the authentication and validation operation
indicate the physical resources are authentic and valid;
permit the physical resources to process a workload in response to a
determination that the sled manifest and results are authentic and valid, and the results
indicate the physical resources are authentic and valid; and
15 prevent the physical resources from processing the workload in response to a
determination that at least one of the sled manifest is not authentic and valid, the results
are not authentic and valid, and the results indicate the physical resources are not
authentic and valid.
2. The apparatus of claim 1, the pod management controller to:
20 receive the sled manifest comprising identifiers for the physical resources from the sled;
and
receive the results of the authentication and validation operation performed to
authenticate and validate the physical resources from the sled.
3. The apparatus of claim 1, the pod management controller to authenticate the sled
25 manifest using a public key, the public key obtained from an original manufacturer of the sled or
a trusted third party.
4. The apparatus of claim 1, the pod management controller to:
generate a hash value based on the sled manifest;
compare the hash value with another hash value communicated with the results;
30 validate the sled manifest if the hash value and the another hash value match; and
invalidate the sled manifest if the hash value and the another hash value do not match.
5. The apparatus of claim 1, the pod management controller to authenticate the results using
a public key, the public key obtained from the sled manifest.

6. The apparatus of claim 1, the pod management controller to add an indication of the physical resources in a database to permit the physical resources to process the workload.
7. The apparatus of claim 1, the pod management controller to receive the sled manifest and the results of the authentication and verification operations via a secure link with the sled.
- 5 8. The apparatus of claim 1, the pod management controller to generate a composed node comprising at least one of the physical resources of the sled.
9. The apparatus of claim 1, the pod management controller to:
- send a platform certificate request to generate a platform certificate for a composed node to a certificate authority, the platform certificate to identify physical resources including the at
- 10 least one of the physical resources used for the composed node; and
- receive a signed platform certificate.
10. The apparatus of claim 9, the pod management controller to send the signed platform certificate for the composed node to a client device.
11. A non-transitory computer-readable storage medium, comprising a plurality of
- 15 instructions, that when executed, enable processing circuitry to:
- determine whether a sled manifest and results are authentic and valid, the results generated by an authentication and validation operation performed to authenticate and validate physical resources of a sled;
- determine whether the results of the authentication and validation operation indicate the
- 20 physical resources are authentic and valid;
- permit the physical resources to process a workload in response to a determination that the sled manifest and results are authentic and valid, and the results indicate the physical resources are authentic and valid; and
- prevent the physical resources from processing the workload in response to a
- 25 determination that at least one of the sled manifest is not authentic and valid, the results are not authentic and valid, and the results indicate the physical resources are not authentic and valid.
12. The non-transitory computer-readable storage medium of claim 11, comprising a plurality of instructions, that when executed, enable processing circuitry to:
- receive the sled manifest comprising identifiers for the physical resources from the sled;
- 30 and
- receive the results of the authentication and validation operation performed to authenticate and validate the physical resources from the sled.

13. The non-transitory computer-readable storage medium of claim 11, comprising a plurality of instructions, that when executed, enable processing circuitry to authenticate the sled manifest using a public key, the public key obtained from an original manufacturer of the sled or a trusted third party.
- 5 14. The non-transitory computer-readable storage medium of claim 11, comprising a plurality of instructions, that when executed, enable processing circuitry to:
- generate a hash value based on the sled manifest;
 - compare the hash value with another hash value communicated with the results;
 - validate the sled manifest if the hash value and the another hash value match; and
 - 10 invalidate the sled manifest if the hash value and the another hash value do not match.
15. The non-transitory computer-readable storage medium of claim 11, comprising a plurality of instructions, that when executed, enable processing circuitry to authenticate the results using a public key, the public key obtained from the sled manifest.
16. The non-transitory computer-readable storage medium of claim 11, comprising a
15 plurality of instructions, that when executed, enable processing circuitry to add an indication of the physical resources in a database to permit the physical resources to process the workload.
17. The non-transitory computer-readable storage medium of claim 11, comprising a plurality of instructions, that when executed, enable processing circuitry to receive the sled manifest and the results of the authentication and verification operations via a secure link with
20 the sled.
18. The non-transitory computer-readable storage medium of claim 11, comprising a plurality of instructions, that when executed, enable processing circuitry to generate a composed node comprising at least one of the physical resources of the sled.
19. The non-transitory computer-readable storage medium of claim 11, comprising a
25 plurality of instructions, that when executed, enable processing circuitry to:
- send a platform certificate request to generate a platform certificate for a composed node to a certificate authority, the platform certificate to identify physical resources including the at least one of the physical resources used for the composed node; and
 - receive a signed platform certificate.
- 30 20. The non-transitory computer-readable storage medium of claim 19, comprising a plurality of instructions, that when executed, enable processing circuitry to send the signed platform certificate for the composed node to a client device.
21. A computer-implemented method to verify and authentic physical resources, comprising:

determining whether a sled manifest and results are authentic and valid, the results generated by an authentication and validation operation performed to authenticate and validate physical resources of a sled;

5 determining whether the results of the authentication and validation operation indicate the physical resources are authentic and valid;

permitting the physical resources to process a workload in response to a determination that the sled manifest and results are authentic and valid, and the results indicate the physical resources are authentic and valid; and

10 preventing the physical resources from processing the workload in response to a determination that at least one of the sled manifest is not authentic and valid, the results are not authentic and valid, and the results indicate the physical resources are not authentic and valid.

22. The computer-implemented method of claim 21, comprising:

receiving the sled manifest comprising identifiers for the physical resources from the sled;

15 authenticating the sled manifest using a public key, the public key obtained from an original manufacturer of the sled or a trusted third party; and

receiving the results of the authentication and validation operation performed to authenticate; and

20 authenticating the results using a public key, the public key obtained from the sled manifest and validating the physical resources from the sled.

23. The computer-implemented method of claim 21, comprising:

generating a hash value based on the sled manifest;

comparing the hash value with another hash value communicated with the results;

validating the sled manifest if the hash value and the another hash value match; and

25 invalidating the sled manifest if the hash value and the another hash value do not match.

24. The computer-implemented method of claim 21, comprising:

adding an indication of the physical resources in a database to permit the physical resources to process the workload.

25. The computer-implemented method of claim 21, comprising:

30 generating a composed node comprising at least one of the physical resources of the sled.

FIG. 1

Data Center 100

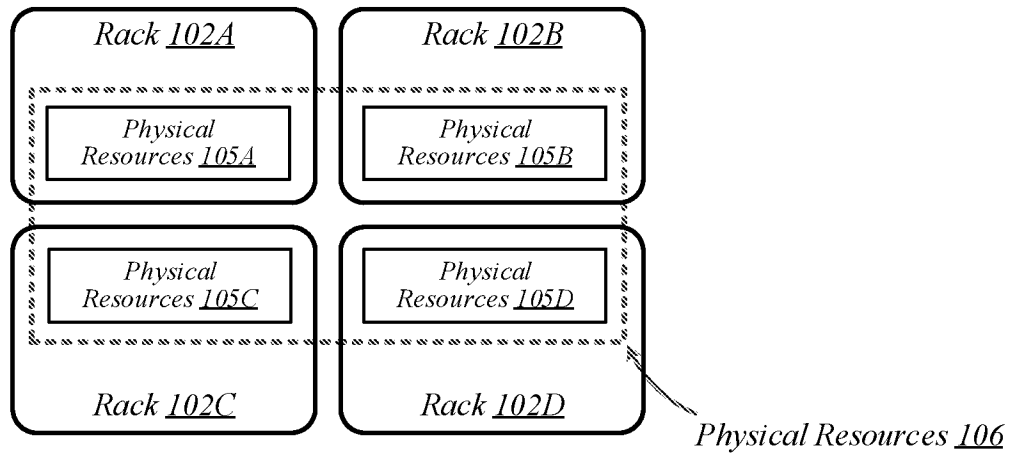
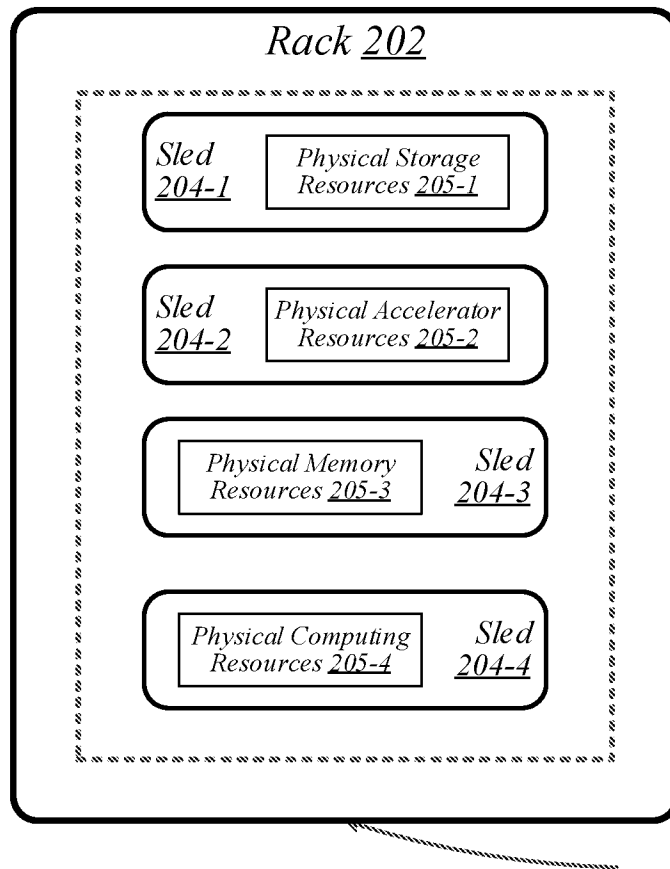


FIG. 2



Physical Resources
206

FIG. 3

Data Center 300

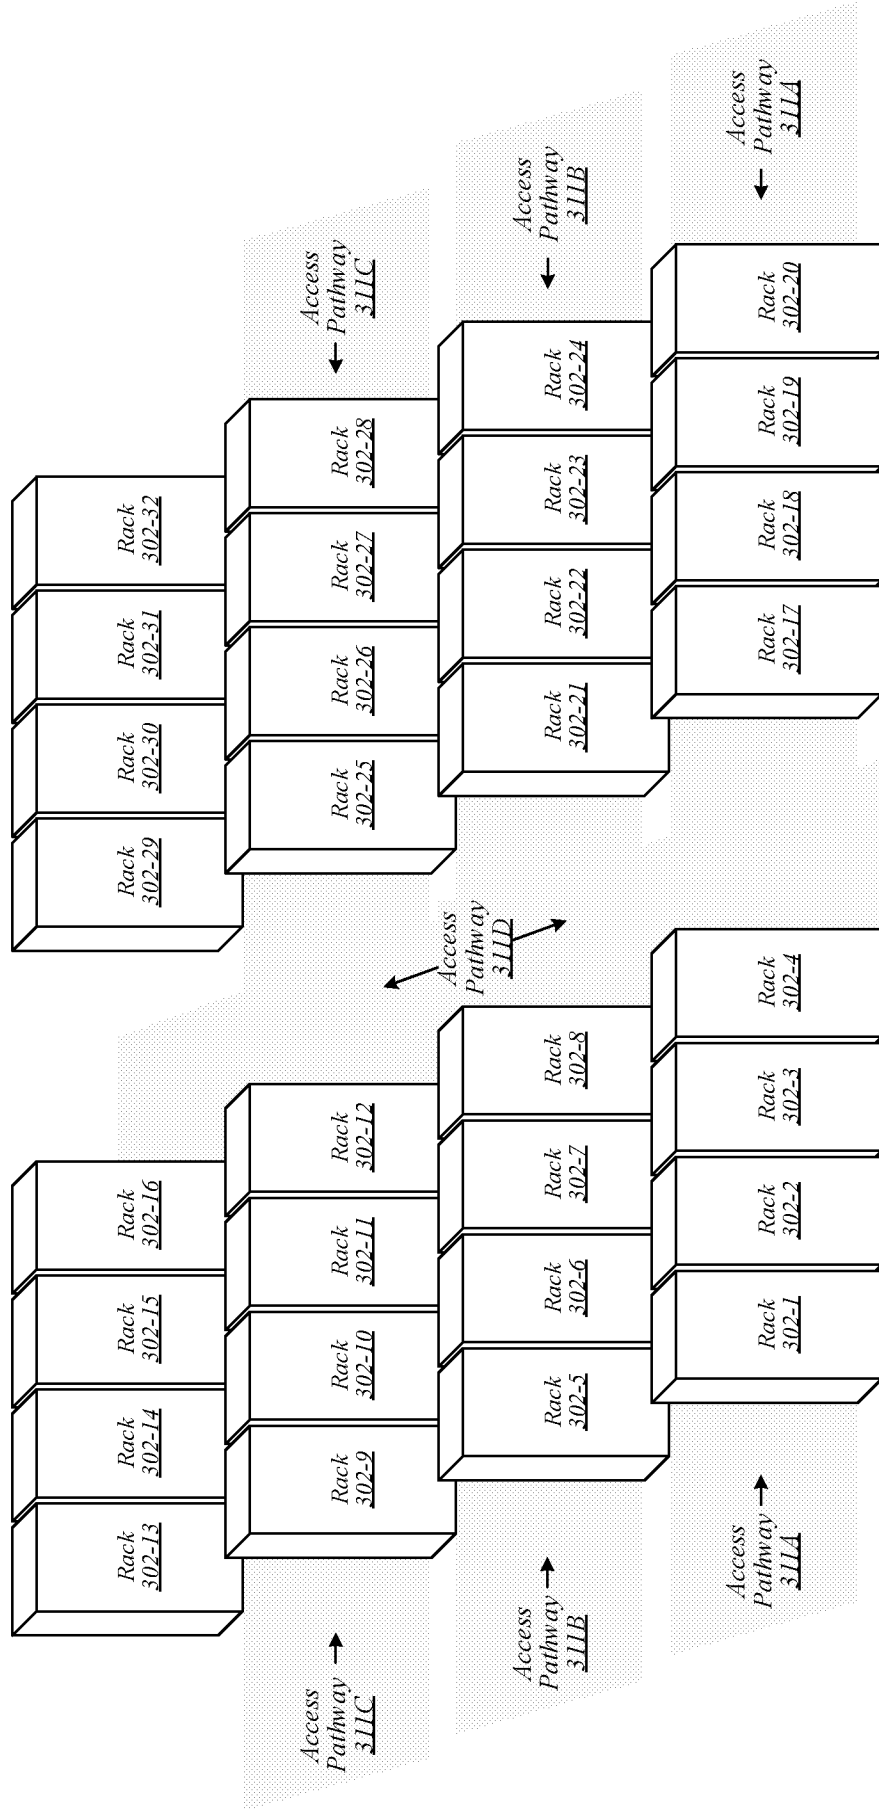


FIG. 4

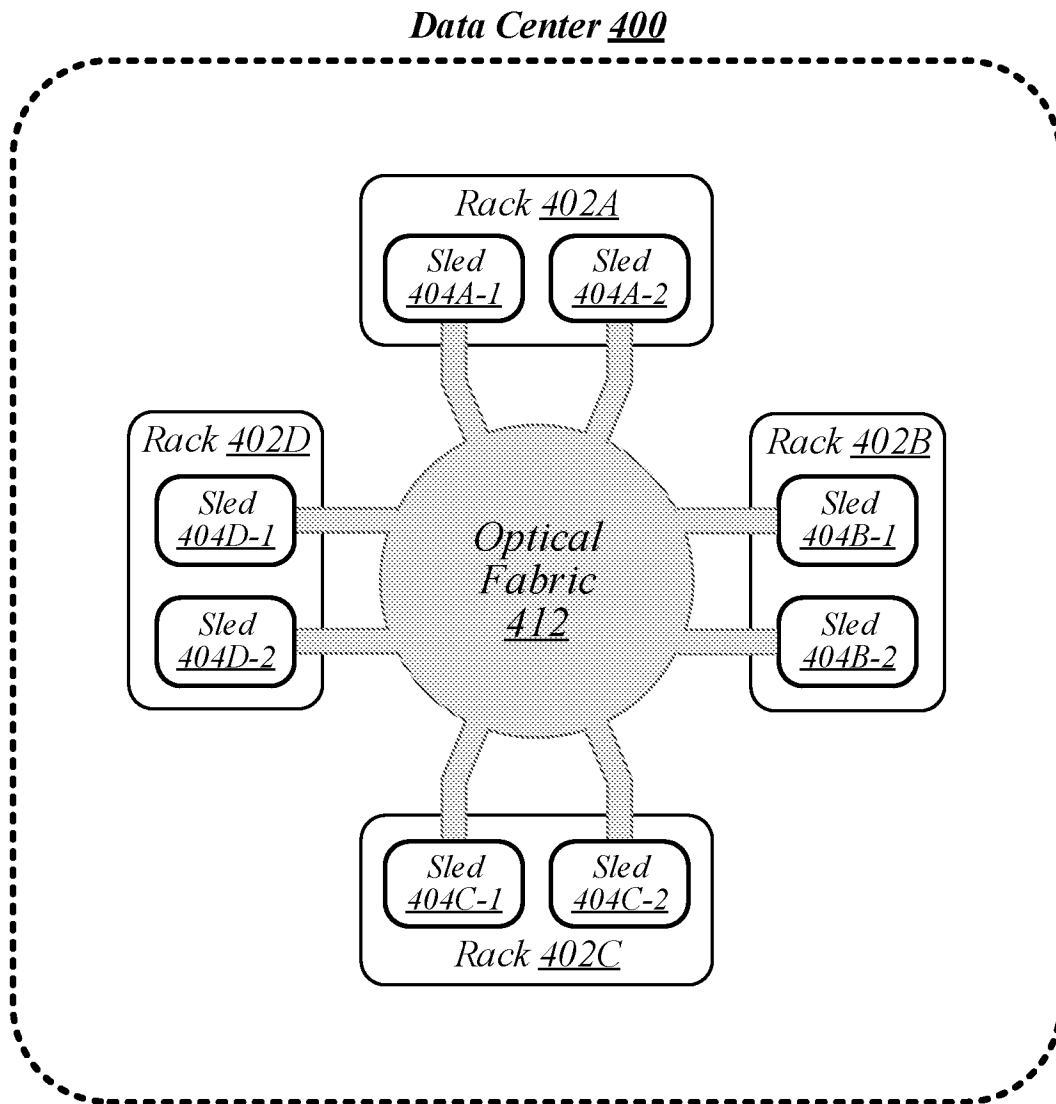


FIG. 5

500

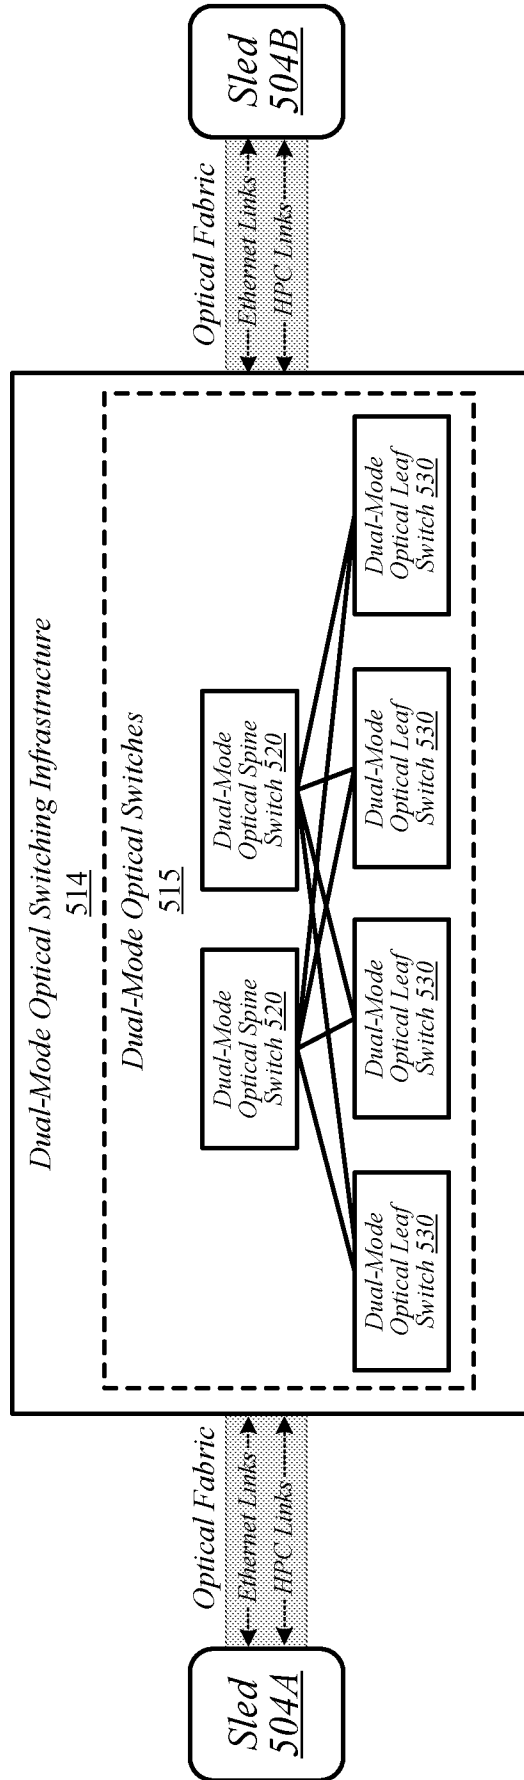


FIG. 6

600

Rack Access Region 601

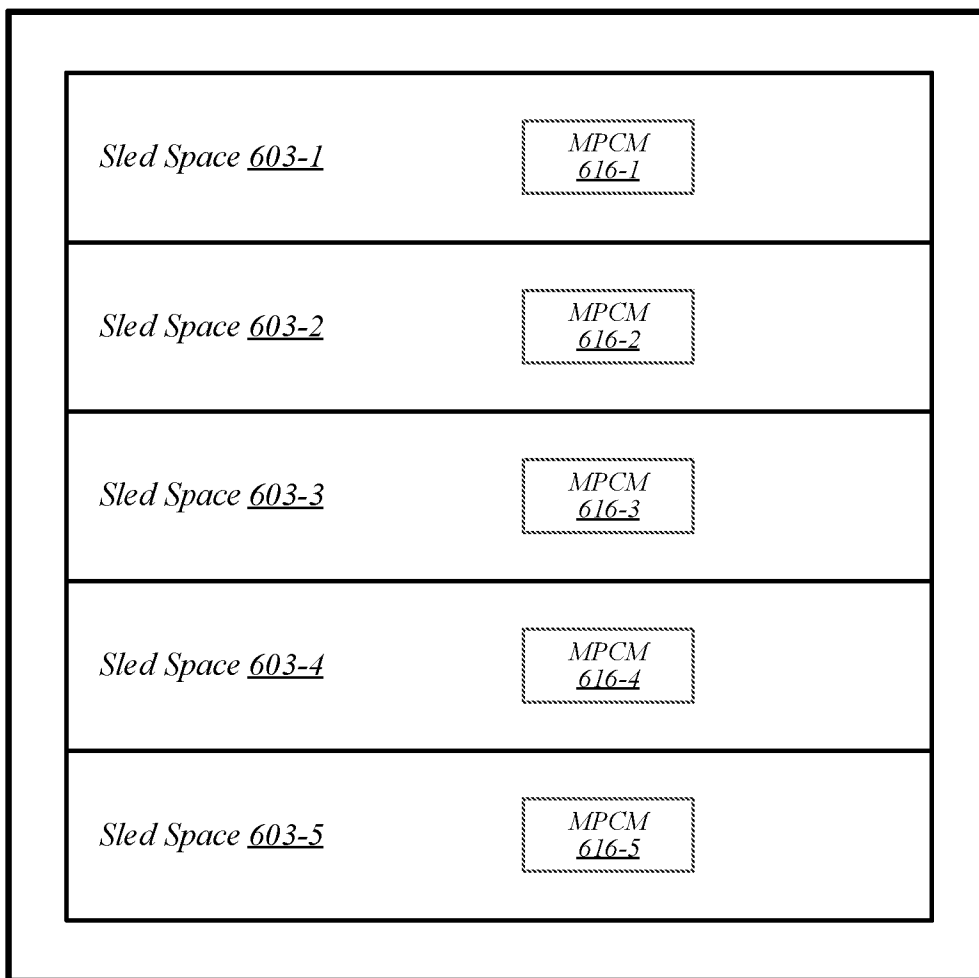


FIG. 7

Sled 704

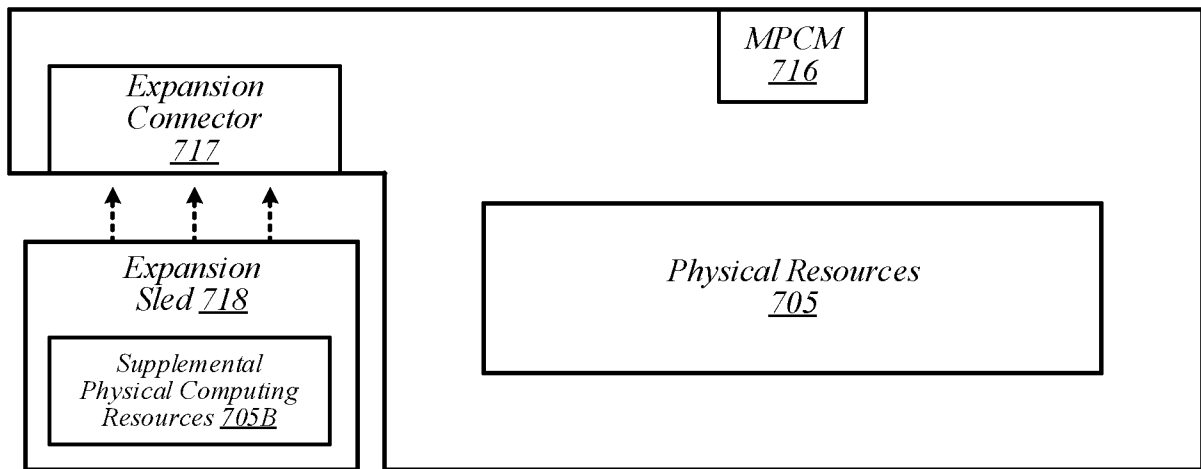


FIG. 8

800

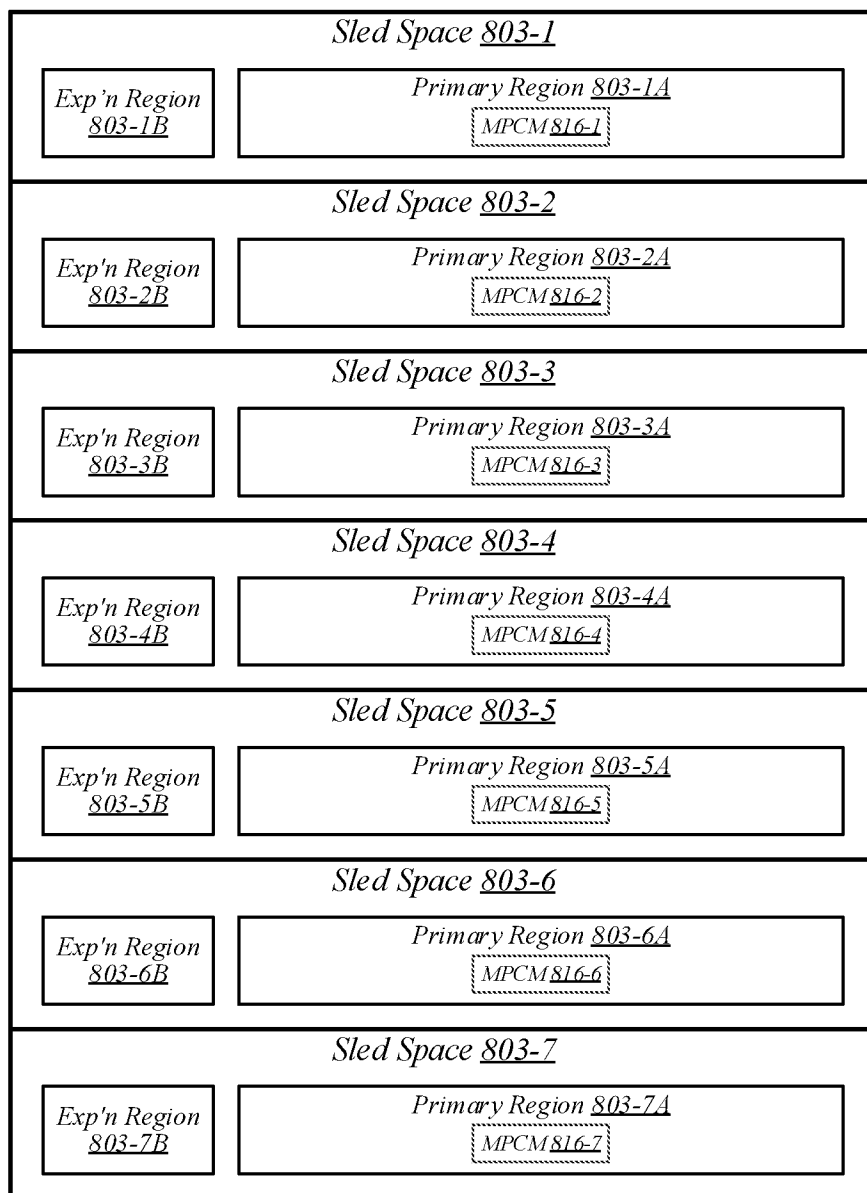


FIG. 9

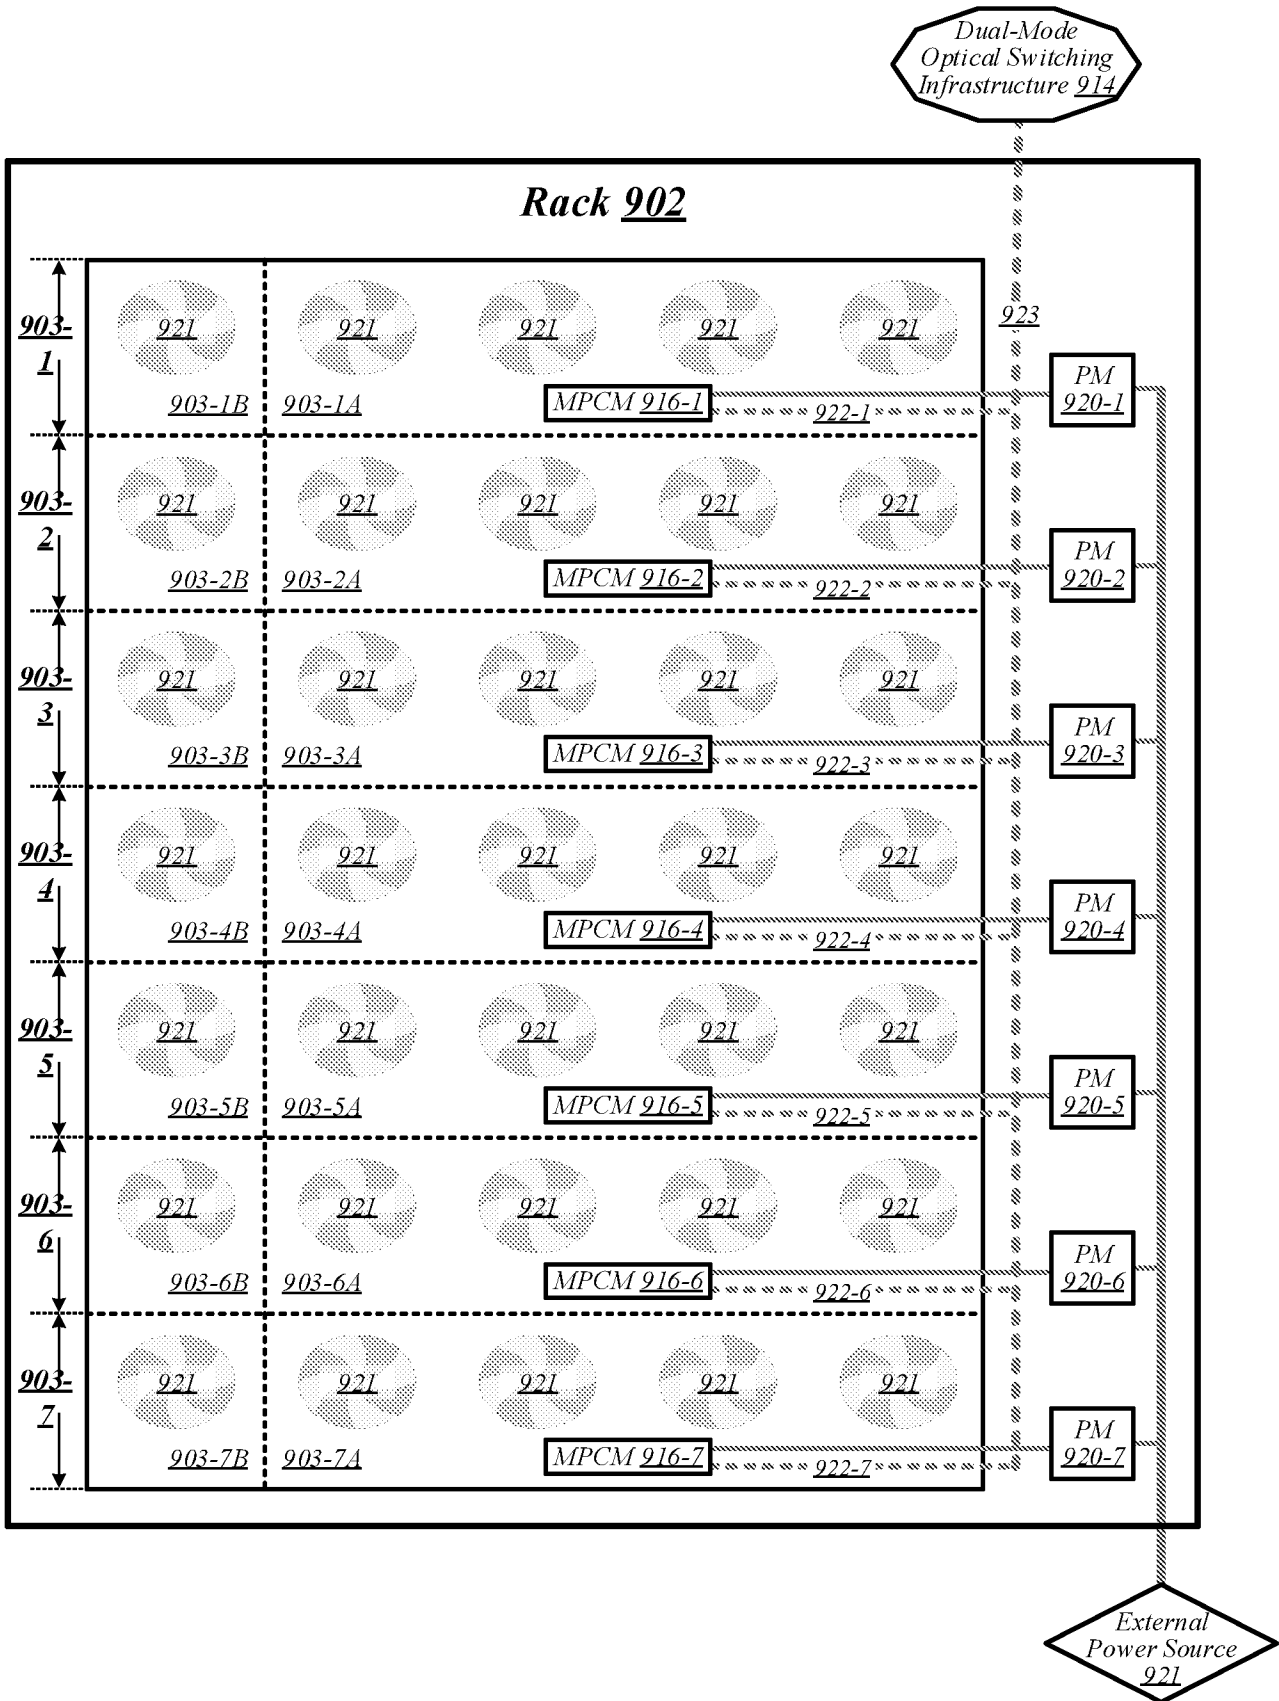


FIG. 10

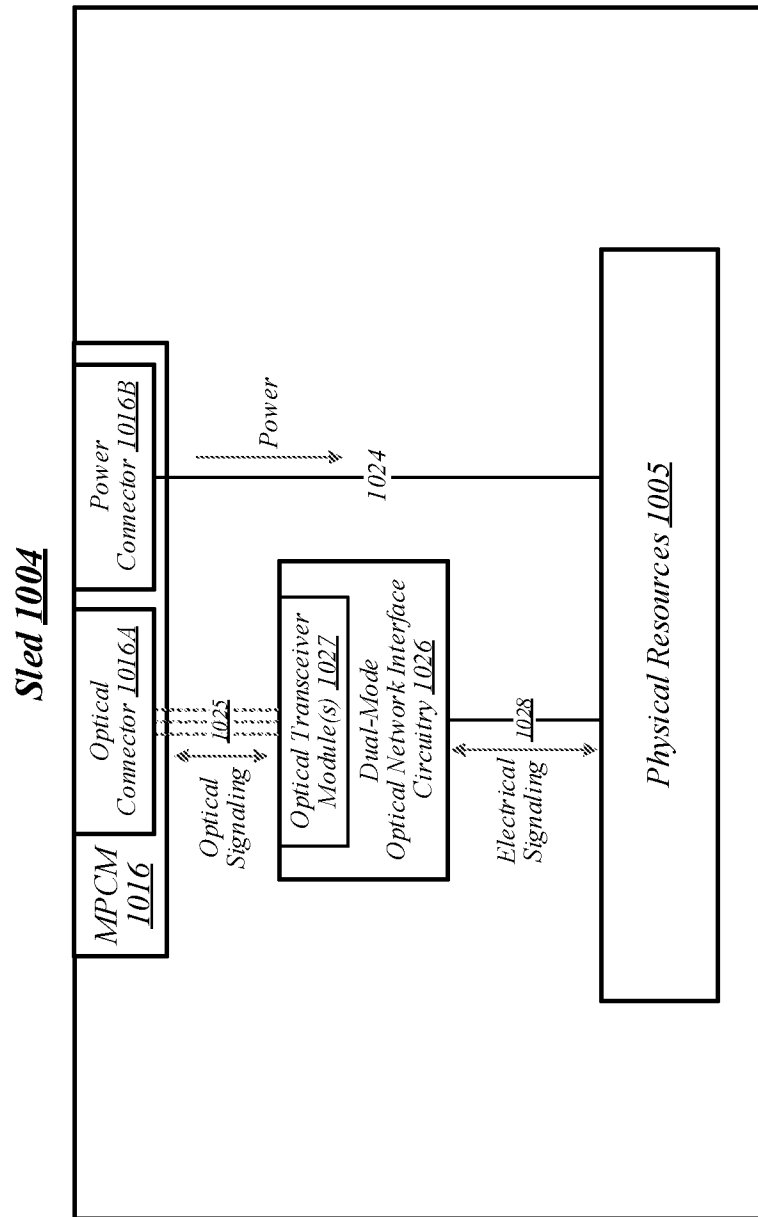


FIG. 11

Data Center 1100

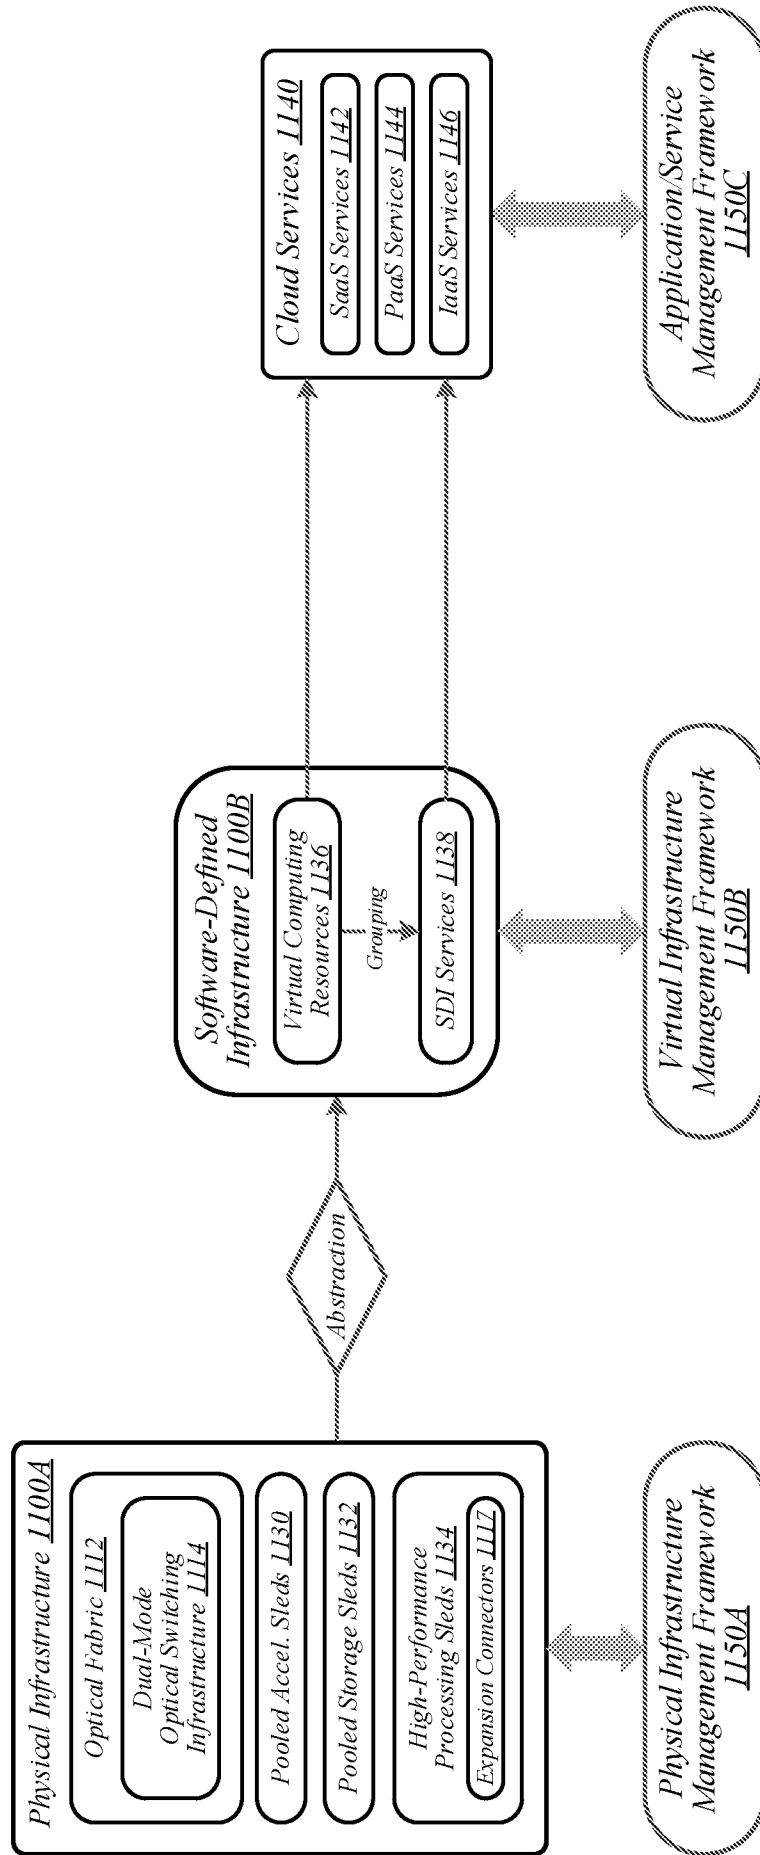


FIG. 12

1200

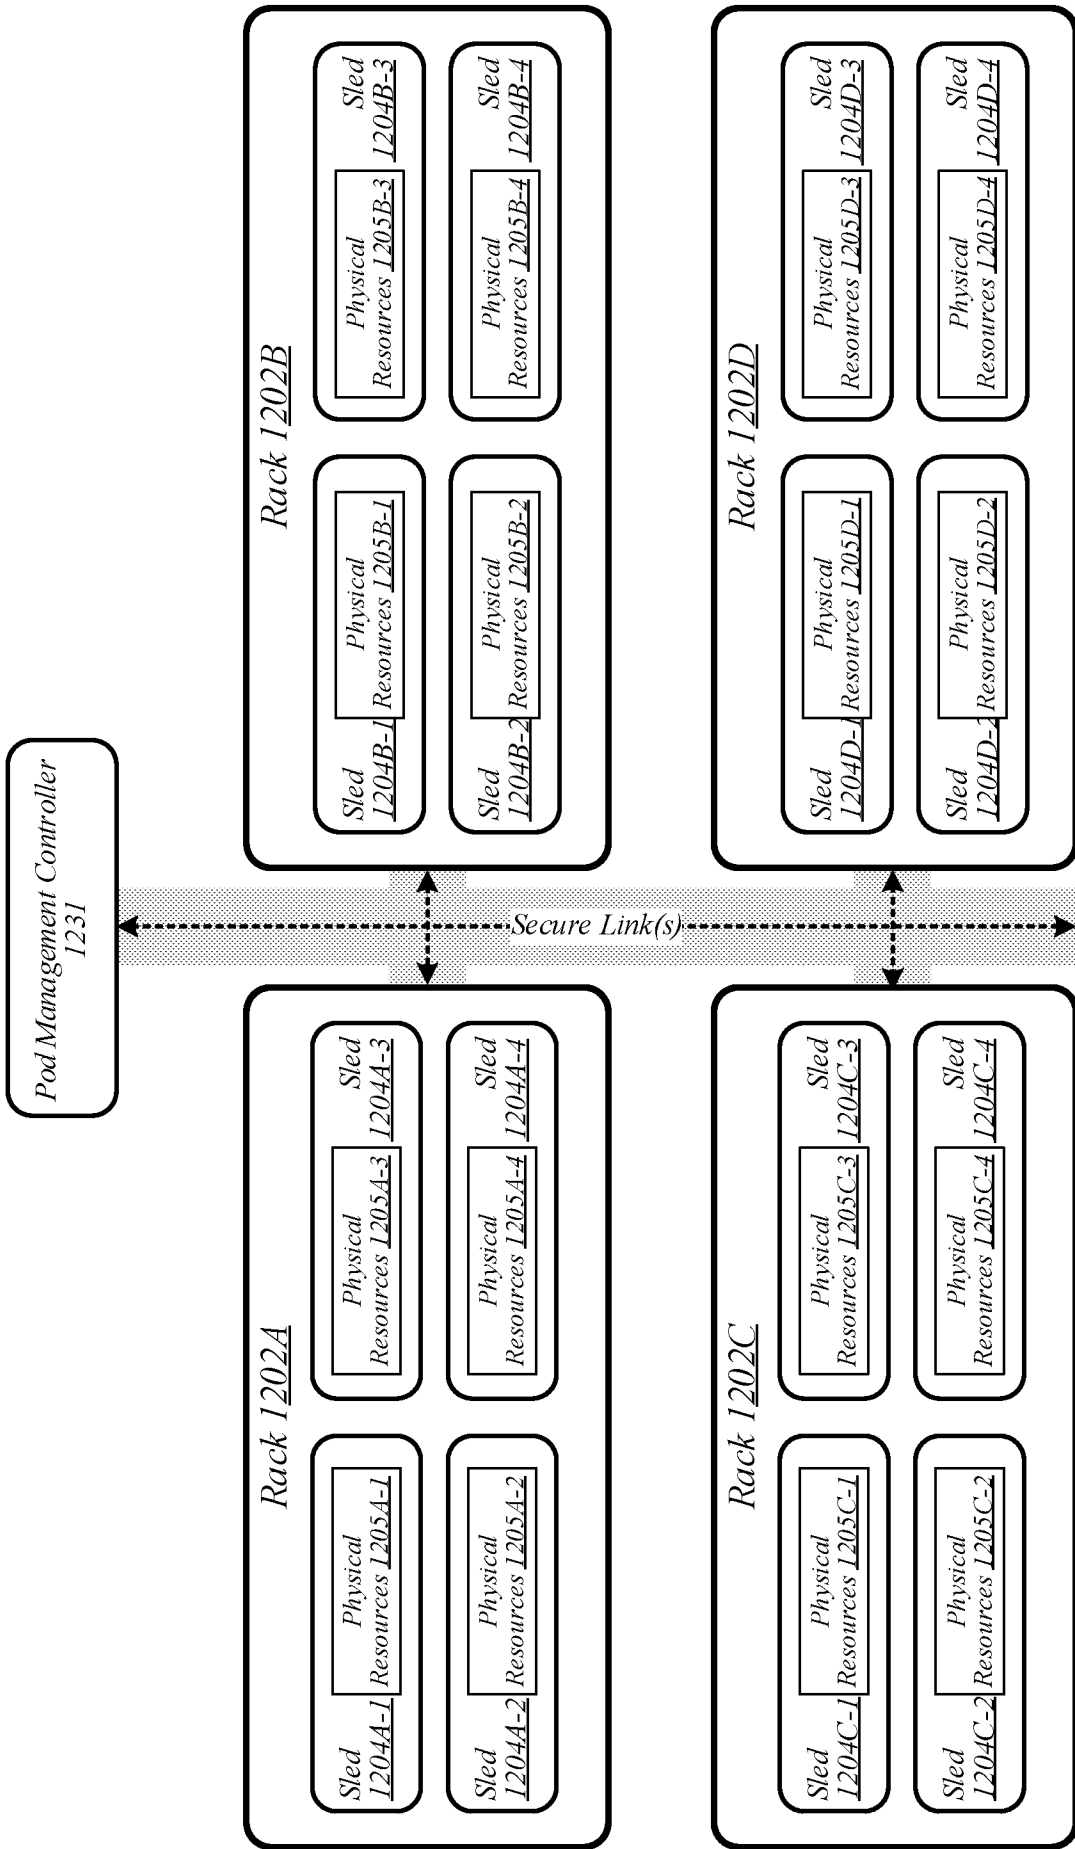


FIG. 13A
Sled 1304

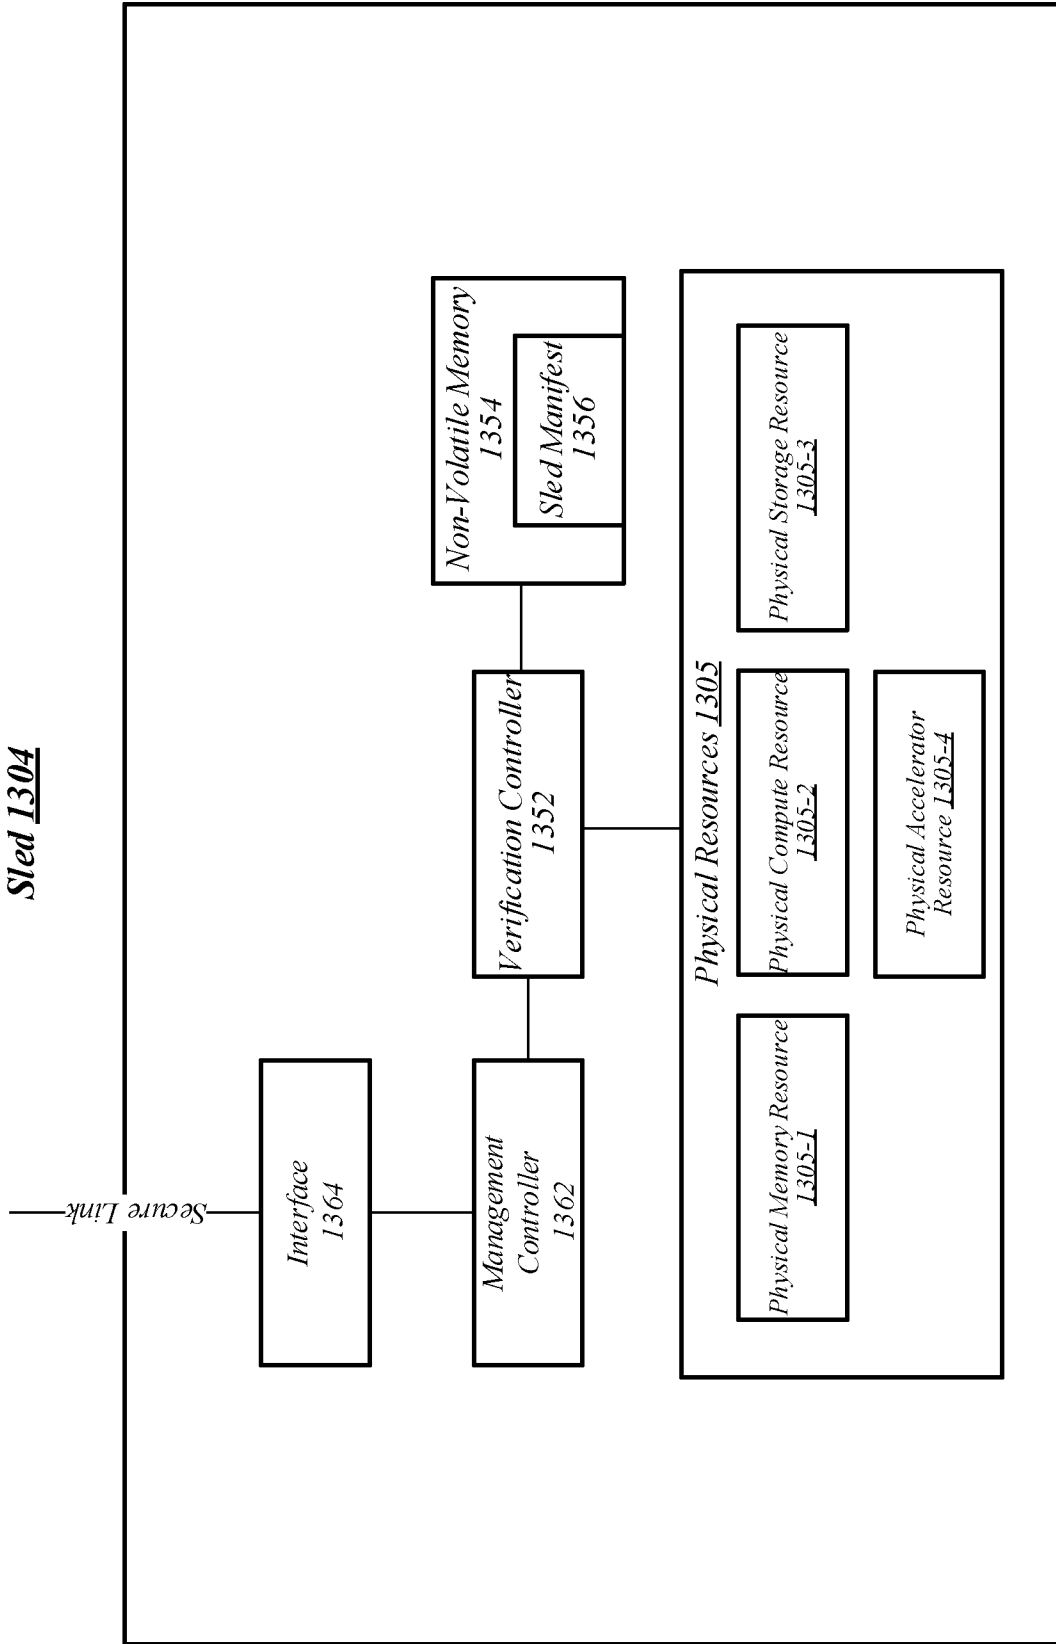


FIG. 13B
1320

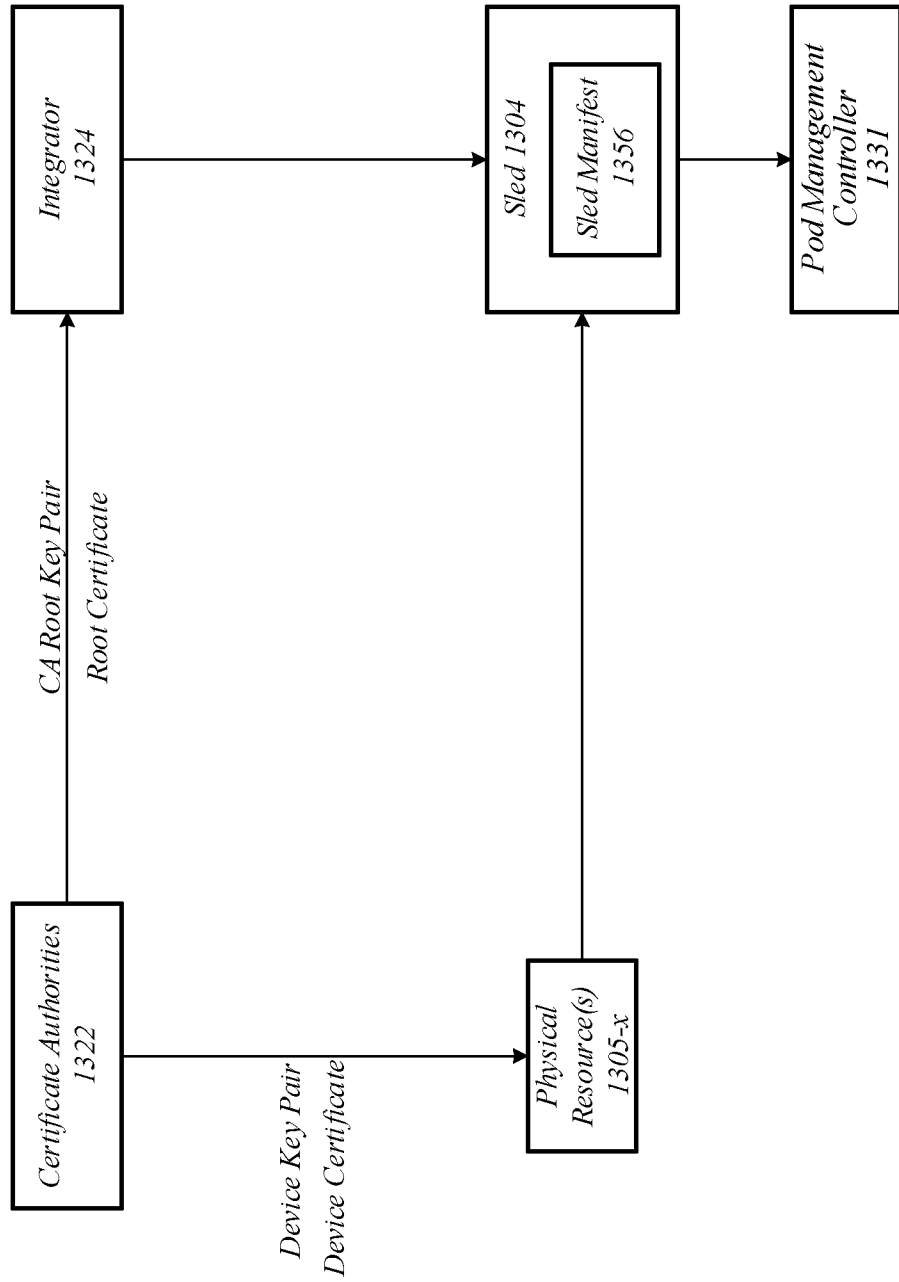
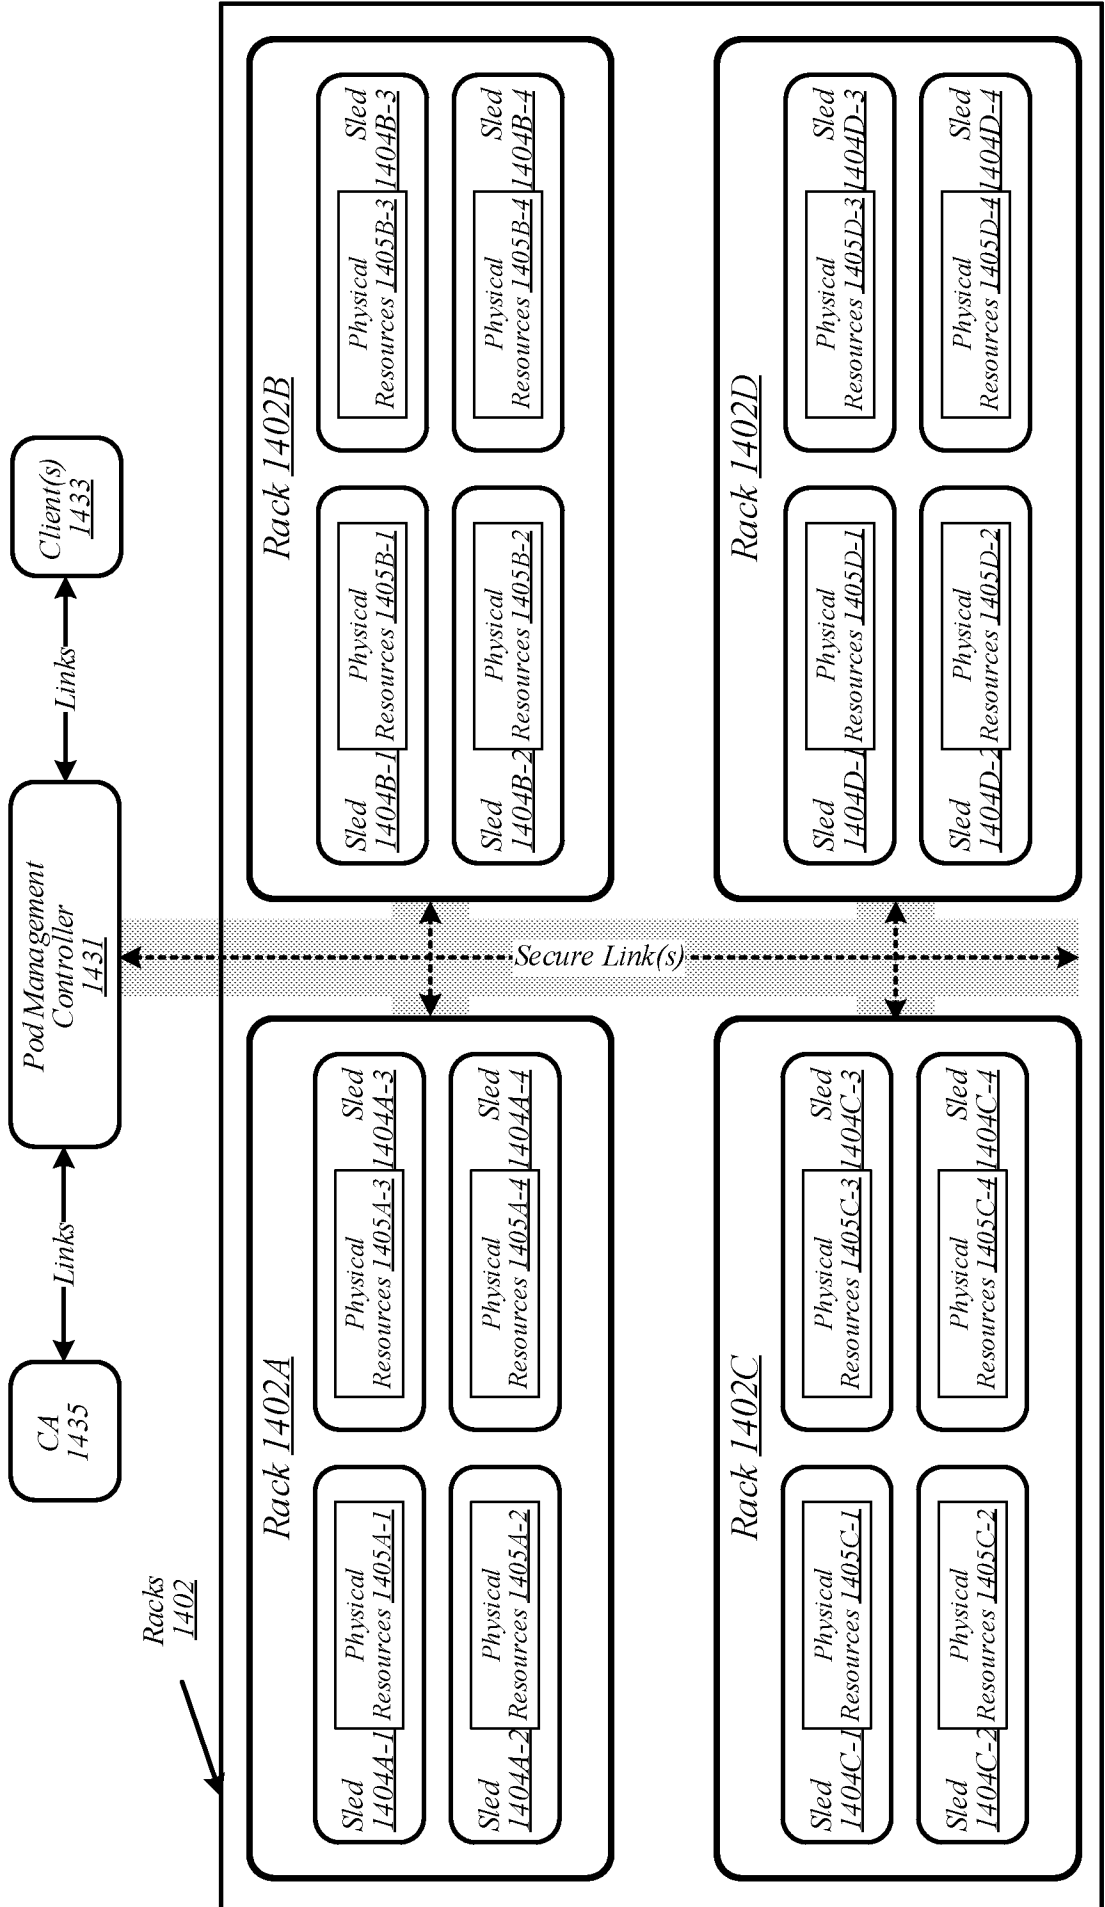


FIG. 14

1400



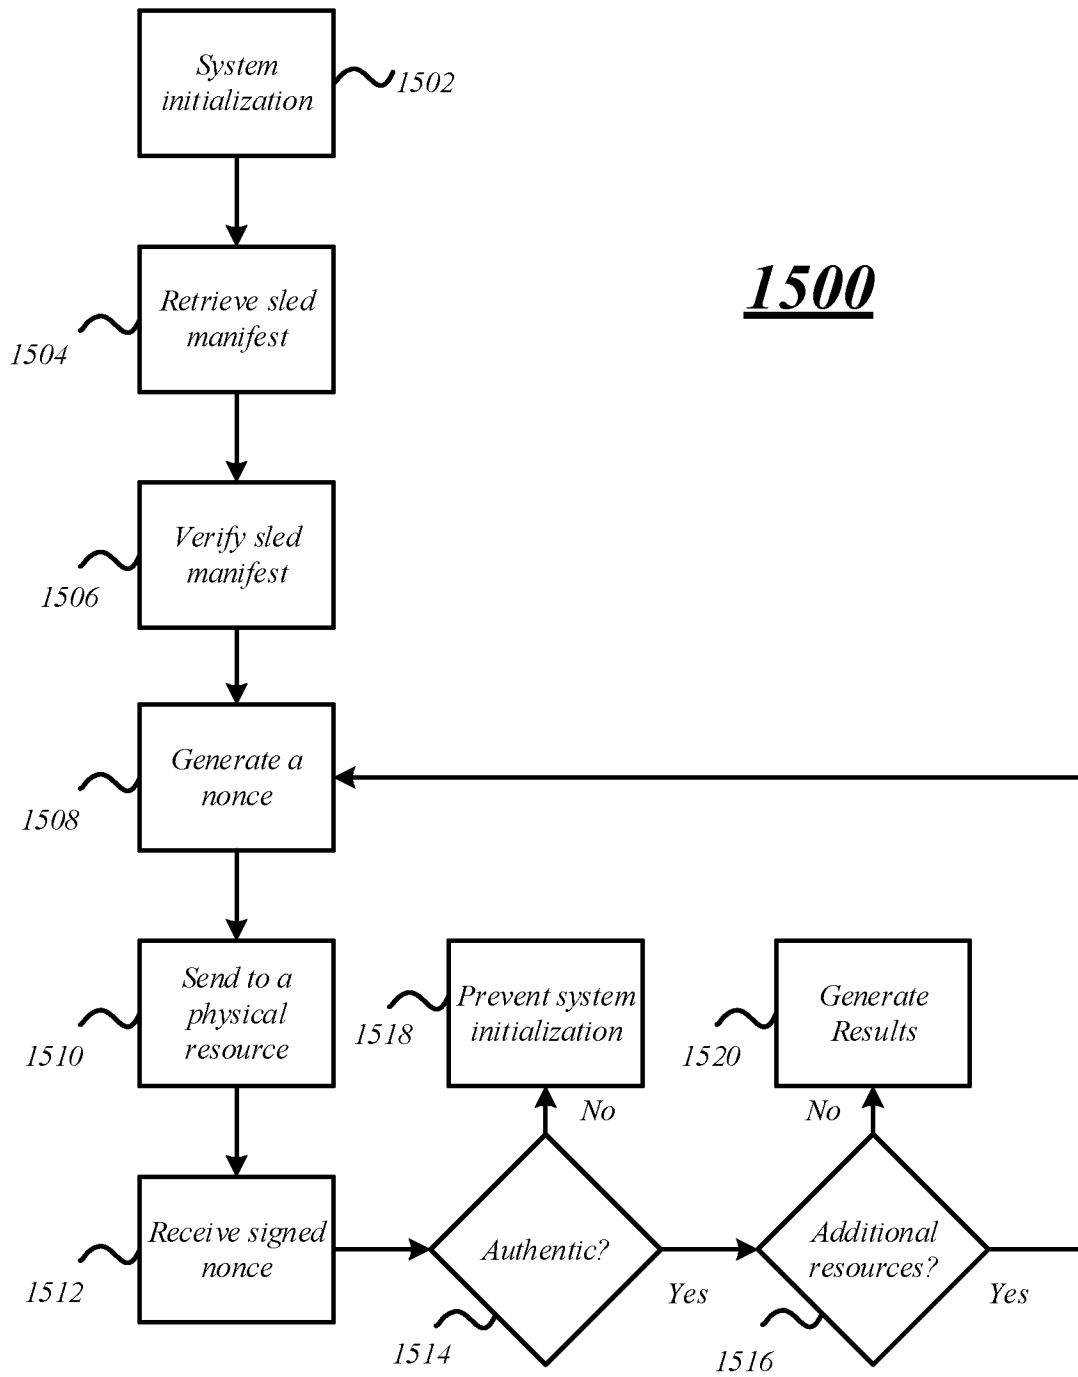


FIG. 15

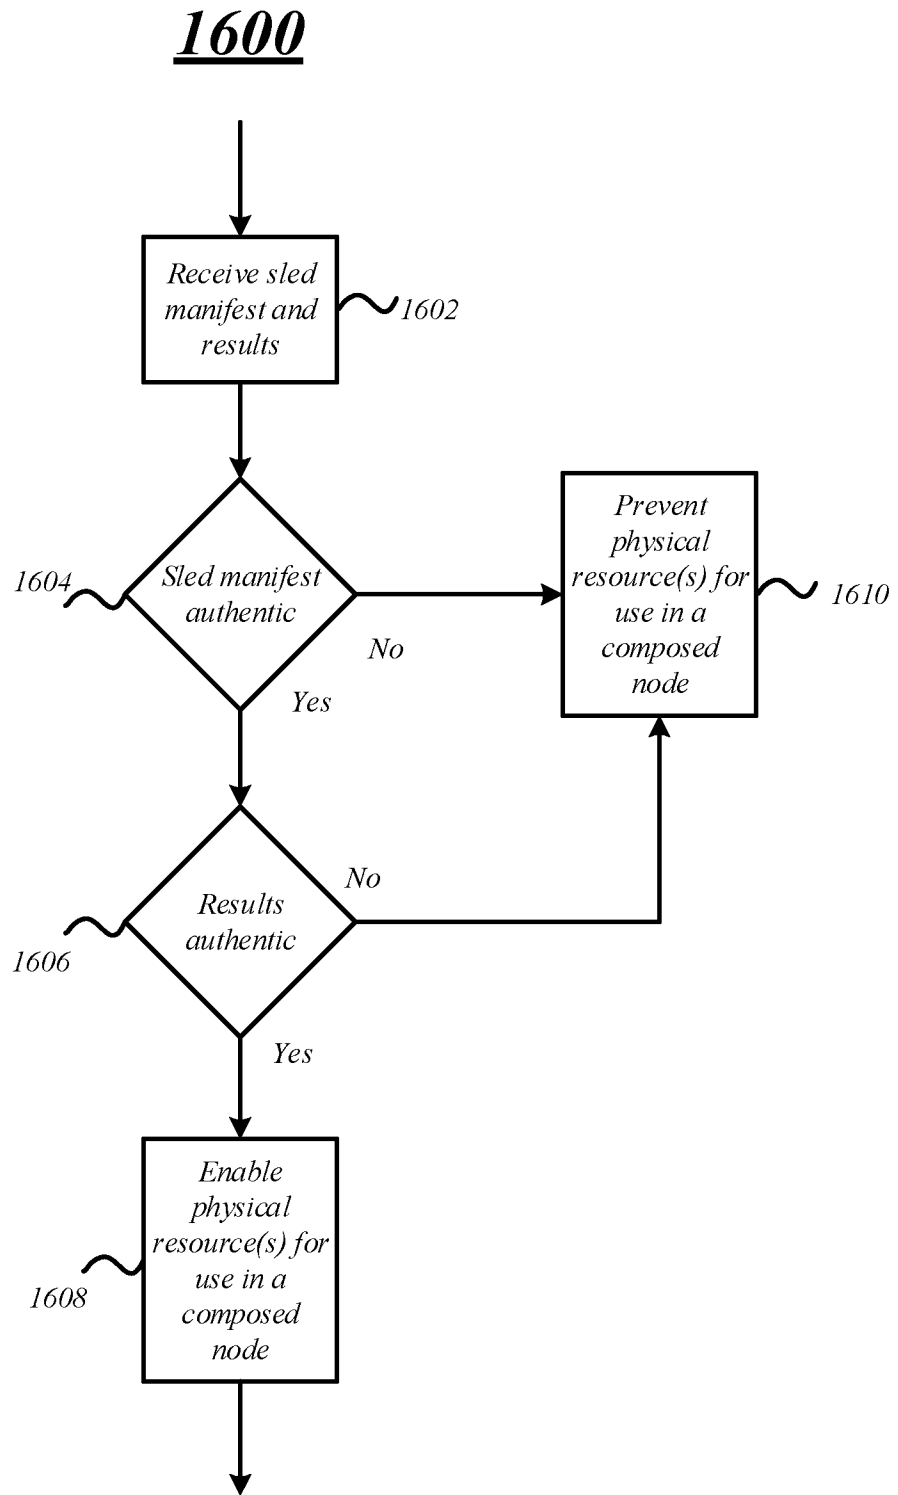
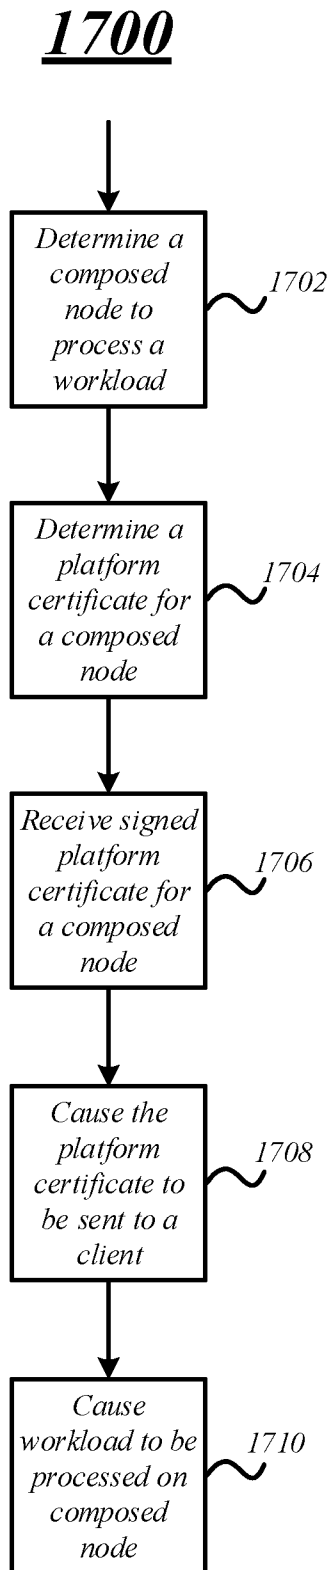


FIG. 16



A. CLASSIFICATION OF SUBJECT MATTER**G06F 21/50(2013.01)i, G06F 21/60(2013.01)i, H04L 9/08(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 21/50; G06F 11/30; G06F 17/60; H04L 29/06; G06F 12/00; H04L 9/32; H04L 9/00; G06F 21/60; H04L 9/08

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: integrity, check, manifest, authentication, validation, determine, hash, key

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|---|-----------------------|
| Y | US 2007-0005992 A1 (TRAVIS SCHLUESSLER et al.) 04 January 2007 See paragraphs [0029], [0038], [0042], [0055]; claims 14, 18; and figure 2. | 1-25 |
| Y | US 2010-0293383 A1 (CHESLEY B. COUGHLIN et al.) 18 November 2010 See paragraphs [0013]-[0014], [0016]-[0017], [0035]; and figure 1. | 1-25 |
| A | US 2003-0009365 A1 (DERMOT TYNAN et al.) 09 January 2003 See paragraphs [0006], [0044]; and figure 2. | 1-25 |
| A | US 2015-0236856 A1 (XILINX, INC.) 20 August 2015 See paragraphs [0017], [0028]; and figure 1. | 1-25 |
| A | EP 2129077 A1 (HITACHI LTD.) 02 December 2009 See paragraphs [0006], [0021]; and figure 2. | 1-25 |

 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

08 November 2017 (08.11.2017)

Date of mailing of the international search report

08 November 2017 (08.11.2017)

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. +82-42-481-8578

Authorized officer

KIM, Seong Woo

Telephone No. +82-42-481-3348



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/043343

| Patent document cited in search report | Publication date | Patent family member(s) | Publication date | | |
|--|------------------|-------------------------|------------------|--------------------|------------|
| US 2007-0005992 A1 | 04/01/2007 | CA 2761684 A1 | 18/11/2010 | | |
| | | CA 2761684 C | 14/07/2015 | | |
| | | CN 102428448 A | 25/04/2012 | | |
| | | CN 102428448 B | 02/09/2015 | | |
| | | EP 2430548 A1 | 21/03/2012 | | |
| | | EP 2430548 A4 | 21/05/2014 | | |
| | | JP 2012-527188 A | 01/11/2012 | | |
| | | JP 5443596 B2 | 19/03/2014 | | |
| | | SG 175987 A1 | 29/12/2011 | | |
| | | US 2016-0170908 A1 | 16/06/2016 | | |
| | | US 9270683 B2 | 23/02/2016 | | |
| | | WO 2010-132647 A1 | 18/11/2010 | | |
| | | US 2010-0293383 A1 | 18/11/2010 | US 2007-0006175 A1 | 04/01/2007 |
| | | | | US 2008-0022129 A1 | 24/01/2008 |
| US 2011-0231668 A1 | 22/09/2011 | | | | |
| US 2012-0226903 A1 | 06/09/2012 | | | | |
| US 7953980 B2 | 31/05/2011 | | | | |
| US 8132003 B2 | 06/03/2012 | | | | |
| US 8499151 B2 | 30/07/2013 | | | | |
| US 8601273 B2 | 03/12/2013 | | | | |
| US 2003-0009365 A1 | 09/01/2003 | IE S20010015 A2 | 17/04/2002 | | |
| US 2015-0236856 A1 | 20/08/2015 | CN 106031082 A | 12/10/2016 | | |
| | | EP 3108609 A1 | 28/12/2016 | | |
| | | JP 2017-506850 A | 09/03/2017 | | |
| | | KR 10-2016-0123336 A | 25/10/2016 | | |
| | | US 9270469 B2 | 23/02/2016 | | |
| | | WO 2015-126967 A1 | 27/08/2015 | | |
| EP 2129077 A1 | 02/12/2009 | EP 2129077 B1 | 01/06/2011 | | |
| | | JP 2009-290648 A | 10/12/2009 | | |
| | | JP 5053179 B2 | 17/10/2012 | | |
| | | US 2009-0300349 A1 | 03/12/2009 | | |
| | | US 2012-0159158 A1 | 21/06/2012 | | |
| | | US 8176316 B2 | 08/05/2012 | | |
| | | US 8819417 B2 | 26/08/2014 | | |