March 20, 2019

A mechanically routable data center

David E. Weekly

A mechanically routable data center

ABSTRACT

To reduce machine-to-machine latency, data centers are increasingly specified to have a high machine density. Per current practices, rack spacing is designed such that any machine in any rack can be serviced by humans. This limits the size and performance of a machine cluster. This disclosure describes data center configurations of high machine density, formed by arranging racks in tight two-dimensional or three-dimensional arrays. Provision is made to enable the lateral movement of racks. Machines are serviced by moving one or more racks up, down, or to the side, such that a temporary aisle is created leading to the rack under repair. A human or autonomous operator can reach the rack under repair via the temporary aisle. Once repair is completed, the aisle is closed.

KEYWORDS

data center; compute rack; movable rack; data center maintenance; data center servicing; catwalk; gantry; machine cluster; rack cluster; compute cluster; spherical rack configuration

BACKGROUND

To reduce machine-to-machine latency, data centers are increasingly specified to have high machine density. Per current practices, rack spacing is designed such that any machine in any rack can be serviced by humans. This limits the size and performance of a machine cluster.

Besides needing to move racks for service purposes service, it is often advantageous to physically rearrange a compute cluster to better co-locate tasks. For example, a rack of accelerators positioned physically proximate to a rack of data that is to be processed enables more efficient use of resources. A large cluster job can be run through an optimizer to see if such physical rearrangement results in efficiency. If there is an efficiency gain, the physical topology can be changed appropriately, e.g., by moving relevant racks physically close to each other (including physically cabling the racks to each other). Current practices do not lend themselves easily to making quick changes in data center topology to achieve efficiencies in computing.

DESCRIPTION

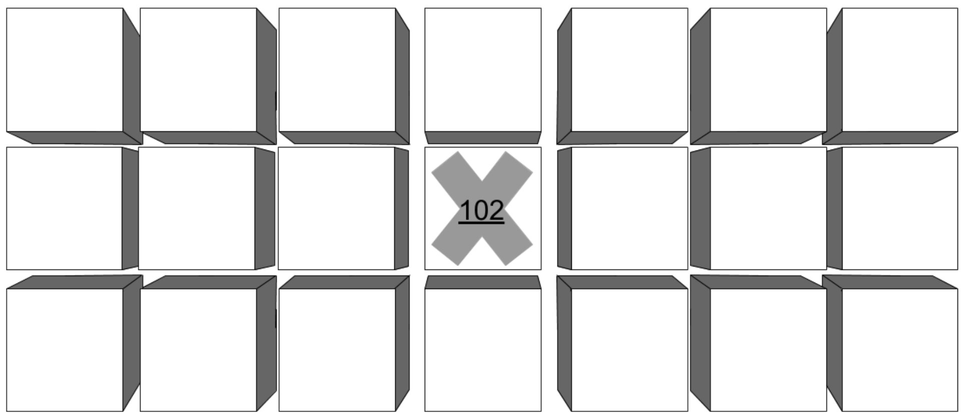

(A)

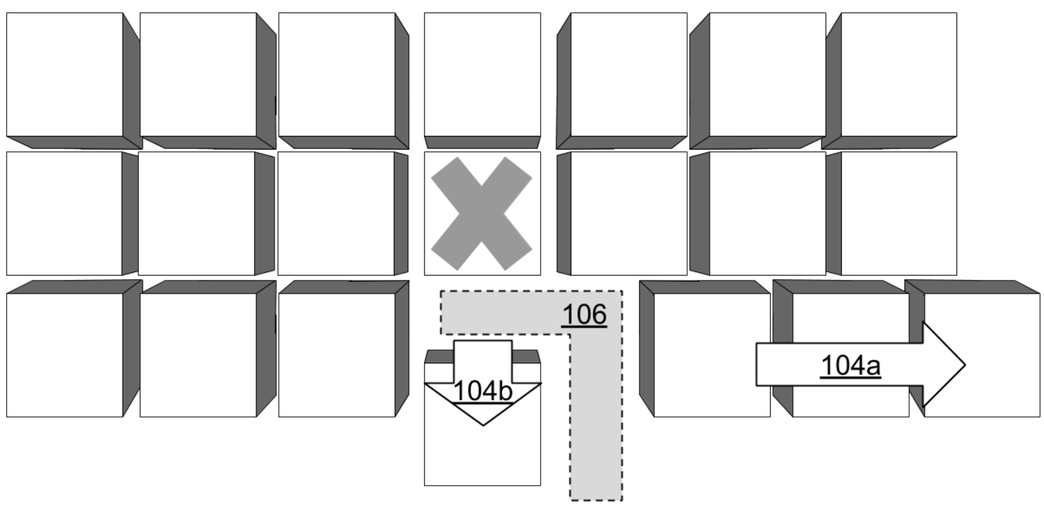

(B)

Fig. 1: (A) Top view of a dense compute cluster (B) Servicing of a rack (marked X) of the cluster

Fig. 1A illustrates the top view of an example rack configuration, per techniques of this disclosure. Racks are arranged in a dense NxM array. Individual racks are capable of being moved laterally, e.g., using sliders, passive casters, self-propelled casters, wheels, tracks, overhead rails, etc., to enable access to any rack in such a dense data center. When a rack (102) is to be serviced, other racks are moved to the side (104a-b) to create an aisle (106) for a human or autonomous operator to enter to service the rack. A repair agent can also arrive from the Z-axis (either emerging out of an access portal in the floor or descending from the ceiling), such that a full path to the opening is not required, only a space in front of the rack to be serviced. In such a case, only a single row or column of racks is moved. Thus, if the repair agent arrives from the Z-axis in Fig. 1B, just one move (104b) is sufficient to access the rack 102 under service.

While a rack is being serviced, spacing locks, e.g., small latches protruding from the floor, are raised to prevent the racks from accidentally re-compacting and damaging the repair agent. Descent/ascent into the repair space is contingent upon verification of correct deployment of these safety latches.

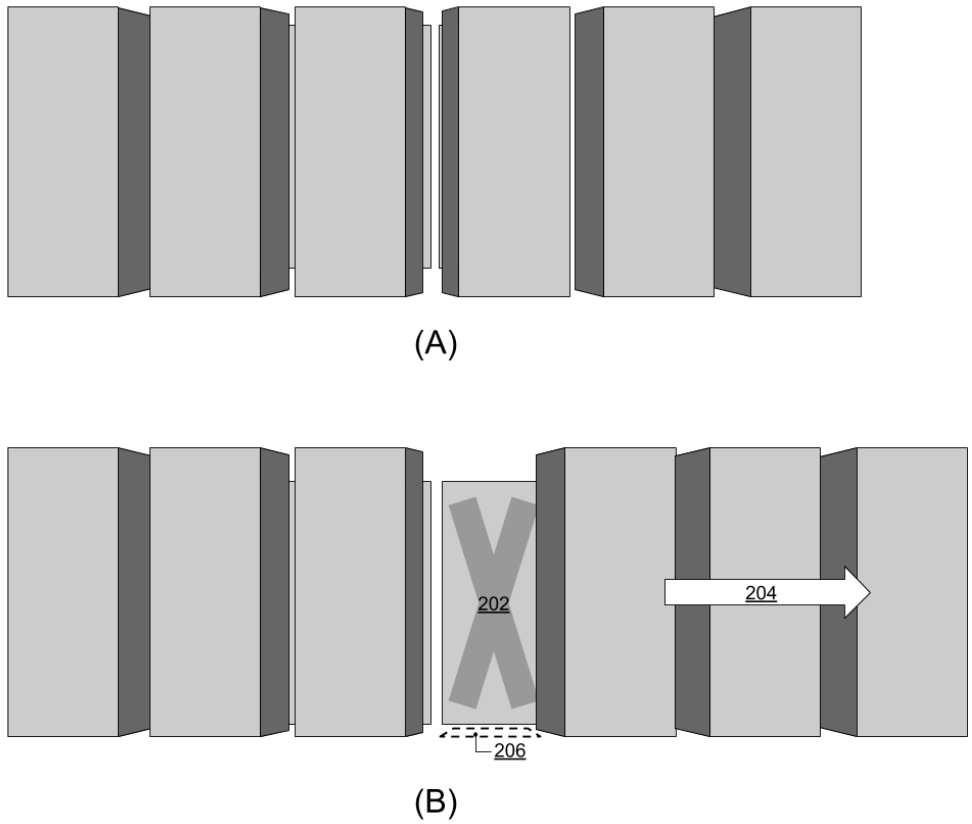

Fig. 2: (A) Front view of a dense compute cluster (B) Servicing of a rack (marked X) of the cluster

Fig. 2 illustrates the front view of another example dense rack configuration. When a rack (202) is to be serviced, other racks are rolled or slid to the side (204) to create an aisle (206) for a human or autonomous operator to enter to service the rack.

Racks in the dense configurations described herein advantageously have vertical connections for power, cooling, and connectivity. In one implementation, soft-mating floor vents direct cold air upward and a baffle routes the cold air to the front of the machines in a rack. Another baffle directs hot air upwards to a similar soft-mating hot air ejection port at the top. In a variant, airflow is forced continuously from bottom to top instead of side-to-side. If airflow occurs from bottom to top, the compute elements are positioned vertically, e.g., with the front of a server facing towards the floor and the rear of the server pointing towards the ceiling, to allow smooth airflow around the servers. In another variant, racks are cooled with a liquid, e.g., water, coolant, and have flexible cold-water input and hot-water output tubes attached to each rack. Valves enable the disconnection of the rack from the cooling system during servicing. Such valves, which serve as a safety mechanism to prevent disconnection while rack power is on, are activated after the rack is powered down. The valve-based safety mechanism avoids accidental disconnection of liquid cooling from a powered-on rack, which may cause overheating of the rack and its components.

In another topology, rather than individual racks being made movable, whole rows are made movable as a unit. This is similar to dense storage racks in libraries, where there’s enough lateral space for all the racks plus a single aisle, and the racks are laterally translated to open the aisle in the desired position. Such a topology reduces complexity of movement, and enables racks to be permanently affixed to each other in one dimension for purposes of power, networking, cooling, seismic bracing, etc. This also applies to multi-rack configurations, e.g., where racks are laterally connected to each other using connectors.

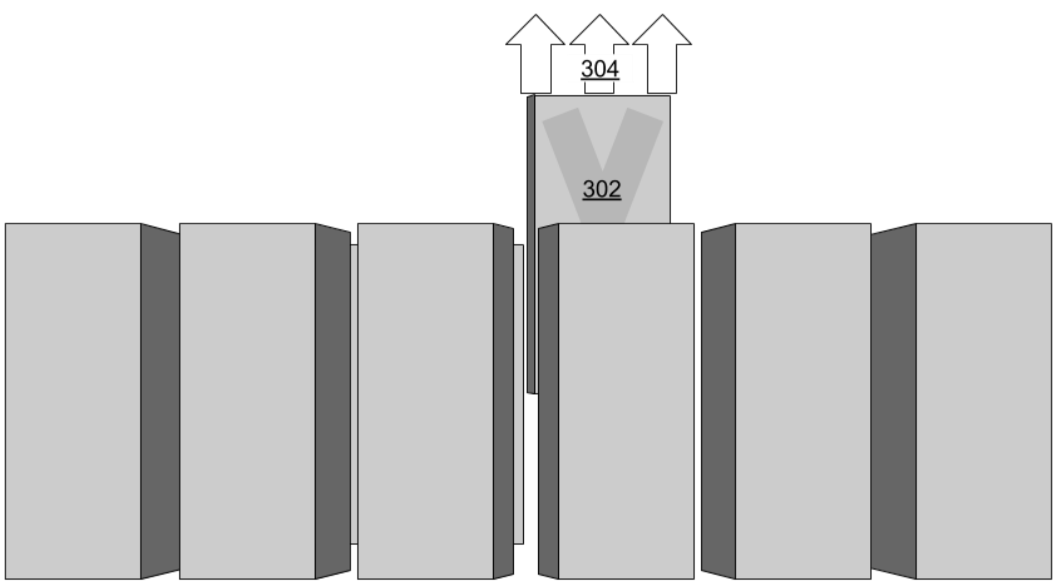

Fig. 3: Moving a rack vertically in order to service it

Alternatively, as illustrated in Fig. 3, instead of laterally moving several racks to make room for an operator in an aisle, a rack that is to be serviced (302) can be translated longitudinally or vertically (304) off-axis from other racks to allow for servicing from all sides. For example, this type of vertical rack ejection is done after disconnecting and powering down the rack. Alternatively, flexible cabling (for cooling, power, connectivity, etc.) can be utilized to allow a rack to be made available for service without taking it offline.

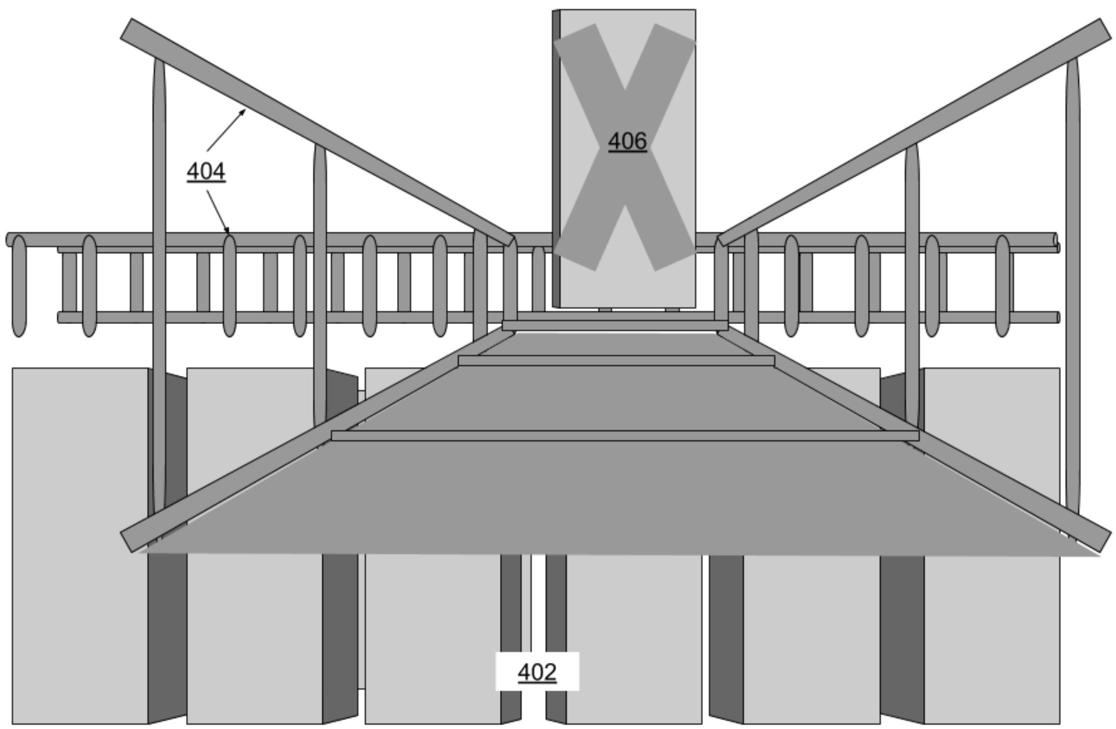

Fig. 4: Use of catwalks to service machines

Fig. 4 illustrates the use of catwalks to service a machine. A dense NxM grid of racks (402) is interconnected as a cluster, and above this grid sit catwalks (404). A rack to be serviced (406) is hydraulically lifted to its full height to allow it to be serviced from the second story. A device on the catwalks pulls up the rack which is braced by the data center floor below.

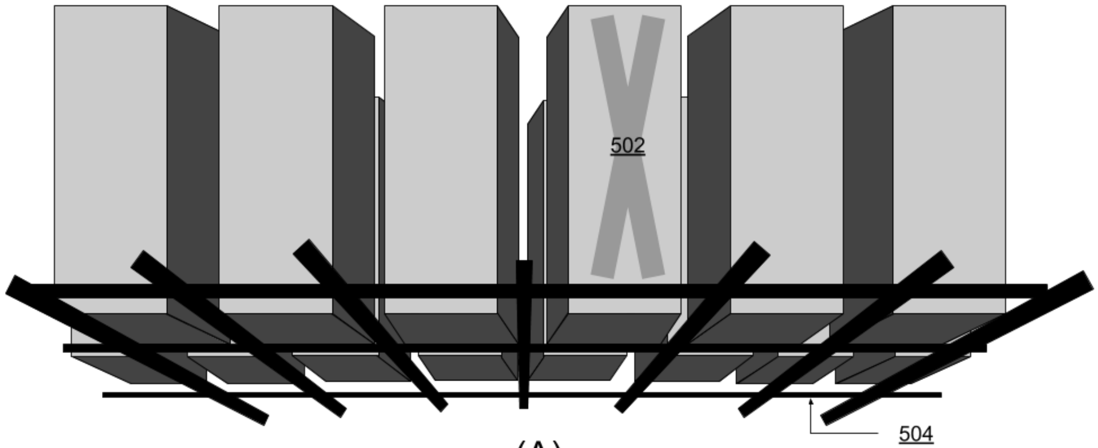

(A)

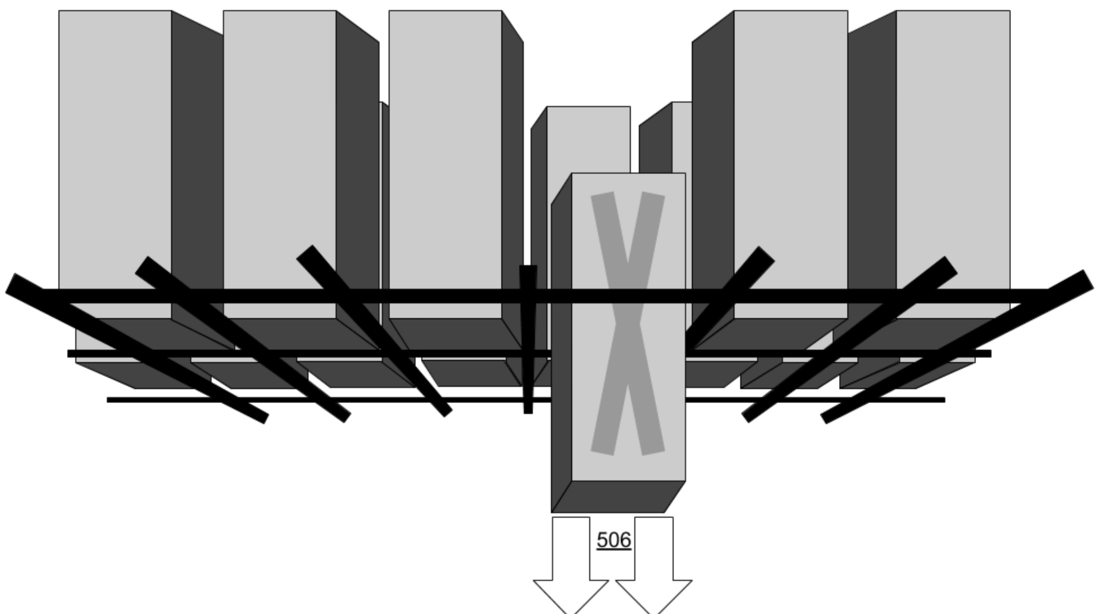

(B)

Fig. 5: (A) Front view of a dense compute cluster (B) Servicing of a rack (marked X) of the cluster by downward movement

Fig. 5 illustrates the use of downward movement to service a rack in a dense compute cluster. A rack that is to be serviced (502) is made to descend to a service level (506) below the data center floor in order for it to get serviced. Support for racks is advantageously provided from below. To ensure access to the rack without obstruction, mesh girders (504) slightly larger than the racks are provided in the floor of the data center and used as follows. A servicing tool positions itself on the floor below the data center under the desired rack and lifts the rack to a height enough to unlatch it from the cross-bracing to the data center floor. The rack is subsequently lowered to the service deck below the data center. The operation is reversed when adding a new rack.

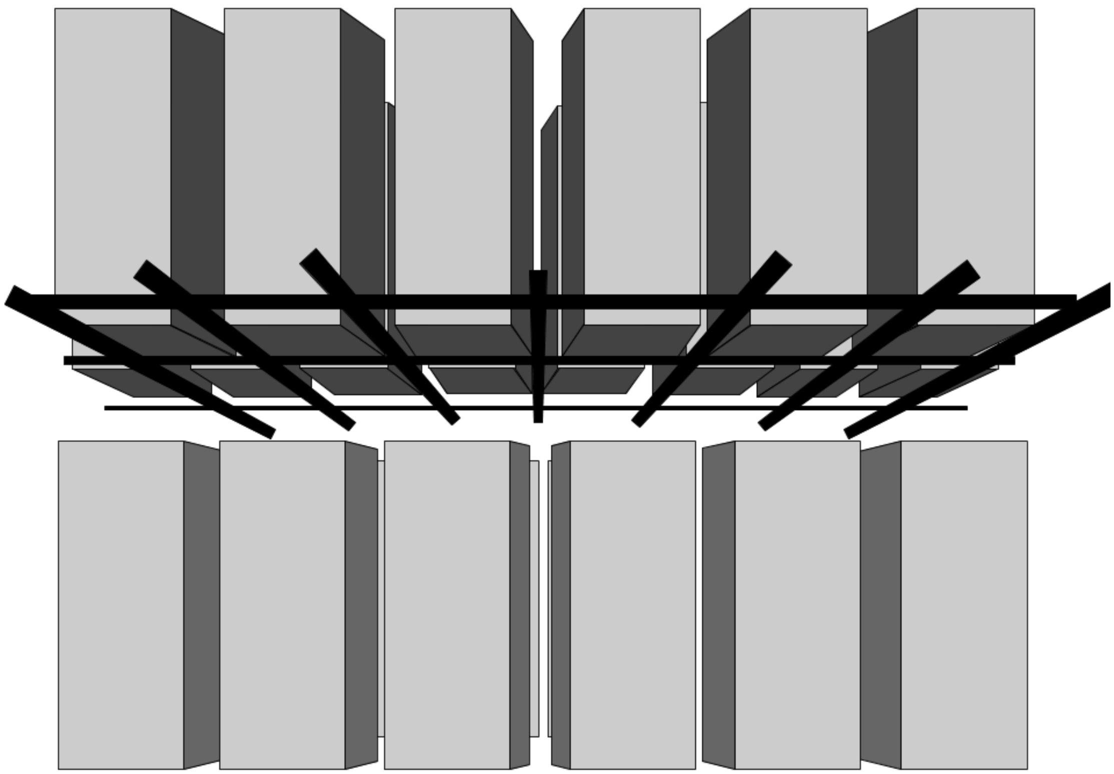

Fig. 6: A two-story rack configuration

Fig. 6 illustrates a two-story rack configuration in which the below-deck and the above-deck service platforms both host racks. Compared to a single-story rack configuration, which achieves a rack-cluster density of only NxM racks per unit area, the two-story configuration of Fig. 6 enables a higher rack-cluster density of NxMx2 racks per unit area.

A large data center can be subdivided into NxMx2 subunits. Racks, once disconnected, are mechanically transported to a common service location, such that a fixed, relatively small number of sites in a data center can be arranged to provide relatively sophisticated rack repair automation. In this case, the rack is brought to the rack fixer rather than the other way around. Such a technique enables a single automation unit to be used to populate, depopulate, and repair racks, resulting in higher utilization and efficiency.

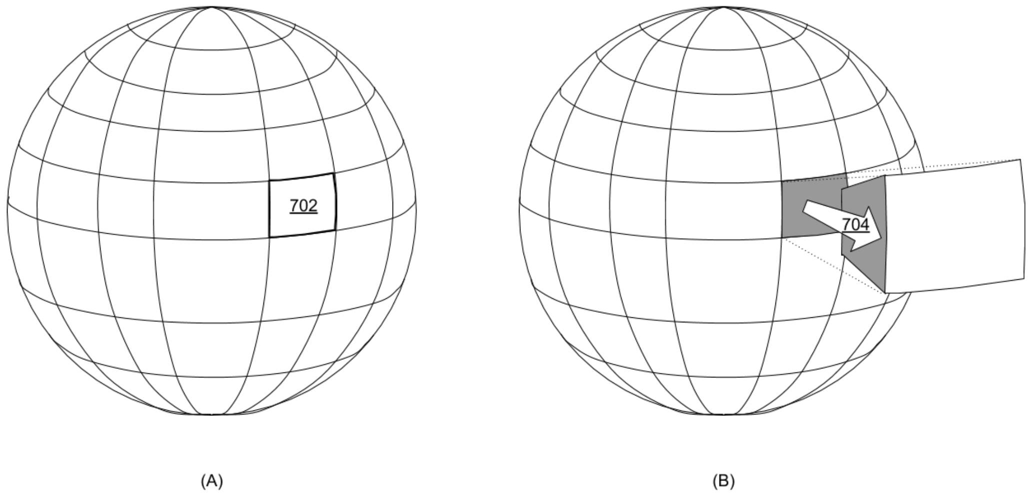

Fig. 7: (A) A spherical rack cluster (B) Servicing a rack of a spherical cluster

Alternative to racks in the form of rectangular prisms, data center components can be laid out such that routing, cooling, and power are supplied along a sphere, as shown in Fig. 7A. Serviceable components are in the form of spherical sectors, such as sector 702. A robotic gantry (not shown) capable of riding rails around the outside of the sphere can pick out a spherical sector, e.g., 704, disconnect it, and route it back to a point in the sphere where it can be serviced. A parallel gantry on the interior of a partially-hollow sphere can similarly pick out internal sectors and mechanically route them for service. With a large enough support structure, traditional rectangular prism racks can be placed in a spherical configuration. Since racks generally comprise solid-state electronics and are capable of operating in a variety of orientations, a spherical configuration of rectangular-prism racks is viable without significant modifications to the rack structure.

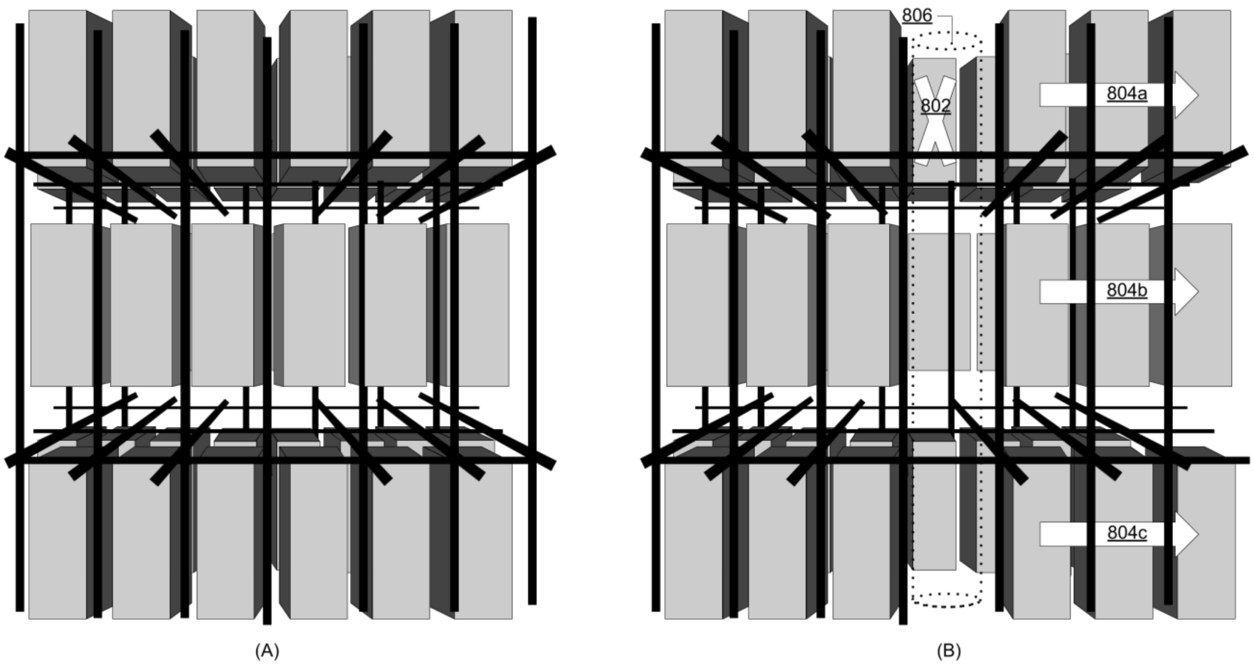

Fig. 8: (A) A multi-story rack configuration (B) Servicing a rack in a multi-story configuration

A multi-story rack configuration as shown in Fig. 8A can be achieved by orienting the aforementioned NxMx2 configuration perpendicular to the ground plane. Similar to the automated movement of passive pallets in industrial warehouses [1], a rack (802) is serviced by creating an aisle (806) towards it. The aisle is created by moving other racks (804a-c) in X, Y, or Z directions as appropriate. However, a multi-story rack as in Fig. 8A can also be serviced with static aisles. For example, a vertical free-aisle plane can allow access to a robotic gantry that pulls a rack from either side of the aisle between two NxM walls of racks. The robot picks an individual rack, disconnects it from power, network and cooling connections, and either ejects it from the cluster or places it in a different spot in the cluster. Similarly, a newly populated rack can be placed into the datacenter in an empty slot and connected to power, network, and cooling, to rapidly turn up capacity.

The multi-story rack can be adapted to use automatic plugs for power, networking, and cooling in each bay of a high-bay configuration thereby enabling ultradense compute clusters. In a variant of the multi-story rack configuration, computing can be made container-based, e.g., an automated and self-contained high-bay configuration with cooling, power, and networking attachments sufficient for a standard container. This advantageously simplifies logistics and reduces time to deployment. The container comes off the back of a truck or rail and gets dropped into the data center almost immediately.

CONCLUSION

This disclosure describes data center configurations of high machine density formed by arranging racks in tight two-dimensional or three-dimensional arrays. Provision is made to enable the lateral movement of racks. Machines are serviced by moving one or more racks up, down, or to the side, such that a temporary aisle is created leading to the rack under repair. A human or autonomous operator reaches the rack under repair via the temporary aisle. Once repair is completed, the aisle is closed.

REFERENCES

[1] SSI Schäfer, “The success story at JYSK goes on by taking a further step into the future with SSI Schäfer” https://youtu.be/KDzxEuxPlWk accessed on Jan. 8, 2019.