Data Center of the Future: Agile, Lean, Green and Unbounded

RFG Perspective: The cloud computing model is forcing IT executives to reevalaute their approaches to data centers, resulting in new designs and systematic ways to reduce personnel and processing costs. While IT executives have begun to address the key issues holistically, especially power consumption, most are executing initiatives targeted at narrowly defined objectives and are tactical. But there will be massive changes over the next decade that will make future data centers agile, lean, green, and unbounded. Therefore, enterprises must have a vision of what data centers will look like 10 years from now if they are to minimize capital outlays when transformating their data centers. Executives should aggressively pursue a data center transformation that will withstand the business, economic, and technology changes that will occur over the next decade.

As RFG wrote in a recent research report, "ever since IT executives started betting their business on **Intel Corp.**-based servers, IT executives have counted on the price/performance ratio improving year-over-year. It was a simple business model – more processing power for less money every year." In fact, IT executives count on the price/performance ratios improving year-over-year on all of the IT hardware they acquire. But what is not getting equal attention is the impact the business, economic, and technology changes that will occur over the next decade will have on their data center requirements. Unless IT and facilities executives focus on these trends companies will overspend in capital expenditures (capex) and operational expenditures (opex) when architecting, constructing or upgrading data centers.

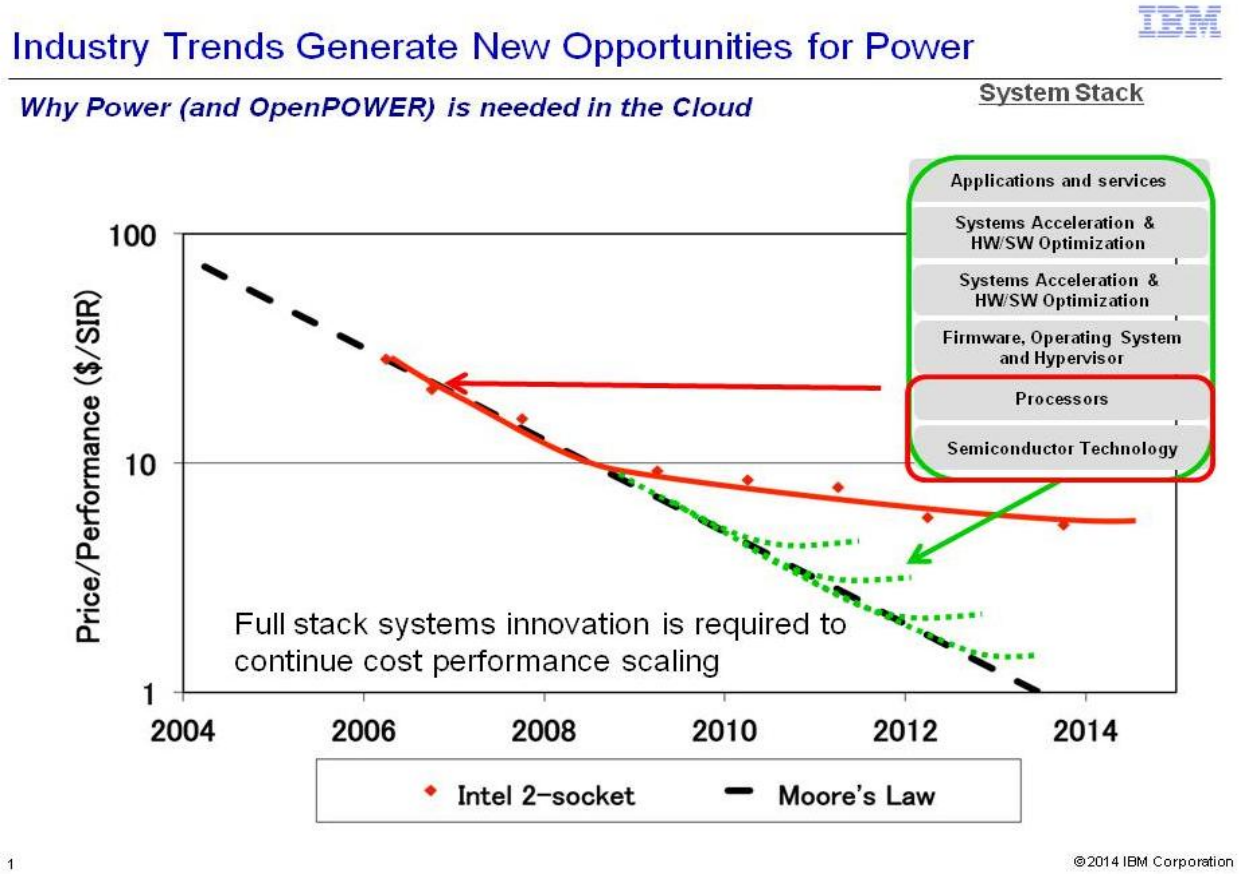

Source: IBM

For example, the above logarithmically-scaled chart shows the price/performance of 2-socket servers. But an almost identical chart could be shown that only looks at the performance gains for Intel and other processors. The footprint required to support today's data center processing power could be reduced by up to 70 percent in four years. In another four years it will shrink again by approximately the same ratio and after 10 years could occupy about five percent of the original space. It could be even less if users switch from 1- and 2-socket servers to 8-, 16- 32-socket servers or scale-up solutions like mainframes and Unix servers. Now, of course, the data center space needs will not shrink that much since the demand for processing power is growing but, for most firms, growth will be less than the rate of annual performance improvement. Therefore, from a server side, it is highly probable that the floor space requirements will shrink from today's needs.

A similar story can be told on the storage side. As the below chart shows, over the course of 10 years what currently occupies about 14,000 square feet of floor space and 1,250 racks will nicely fit in a single rack covering 11 square feet of space. While the storage capacity growth has been known to reach 40 percent per year at many companies, the fact is that physical floor space allocated to storage has been shrinking over the past decade and will continue to do so.

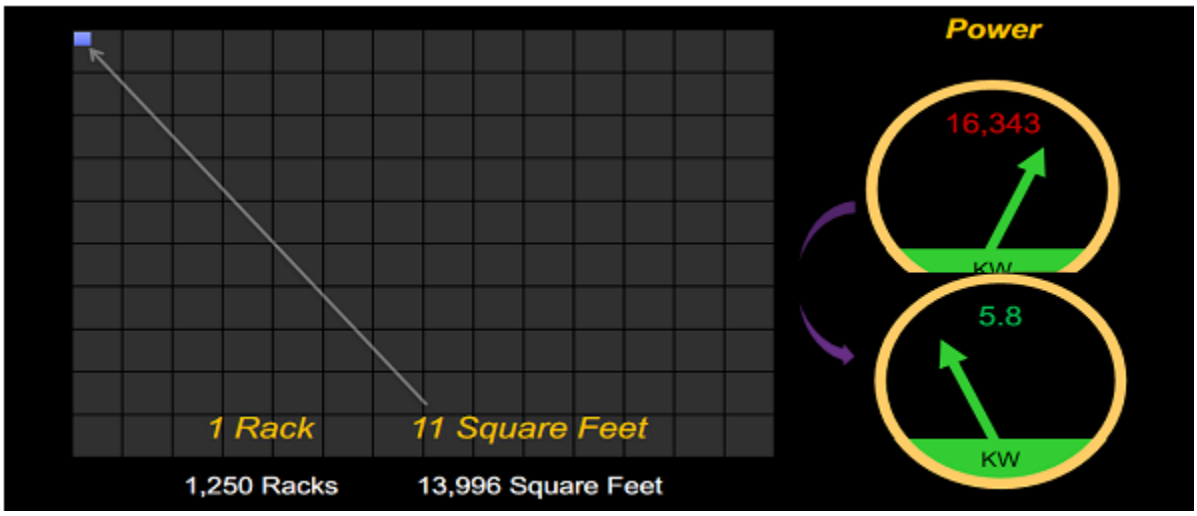

Source: IBM

The space consumed by routers will be as dramatic as storage. Most companies today have not virtualized their routers and, with the new technologies including software defined networking, the number of physical router ports will shrink by 90 percent over time, assuming network traffic remains constant. This also assumes no change in bandwidth. But companies are moving towards 10 and 40 GB Ethernet switches from 1 GB switches; additionally storage networks are being combined with external networks – driving even more shrinkage.

The bottom line is that unless one's enterprise is undergoing rapid business, application, or storage growth, the total floor space needed over the next decade will shrink, not grow. However, floor space is not the biggest issue when it comes to future data center design.

Lean

The bigger issue in the data center will be power consumption and capacity. The density (kW/sq. ft.) is rapidly increasing as all hardware providers improve capacity while shrinking the footprint. Servers have gone from 8 kW/sq. ft. for a mainframe to more than 24 kW/sq. ft. for certain bladed Intel server environments. On an overall basis it has increased from approximately 0.08 kW/sq. ft. to 1.36 kW/sq. ft. in the past decade. This is a 17x increase in density in 10 years. RFG does not expect this exponential curve to lessen, but it could accelerate. This could be trouble for planners, not so much for the building site but acquiring power sources that will be able to meet the demand over the life of the data center. Thus, it could force data centers to become power grid independent, or partially so, by developing their own power sources.

Additionally, IT executives are now trying to optimize their data centers through automation, consolidation, thin provisioning, standardization, virtualization and use of the latest technologies. Companies are actively pursuing strategies that are creating least-cost operating environments. The use of various cloud computing platforms is just one method of satisfying the need for lean. ***RFG has found through various studies that IT executives can cut their operational expenditures by 50 percent while simultaneously doubling capacity and potentially cutting power consumption by more than 40 percent.***

Green

Power usage effectiveness (PUE) has been the predominant measure of power consumption, even though it only evaluates the non-IT infrastructure and is thrown out of whack when new, more power-efficient IT hardware is installed. Nonetheless, various studies put the average PUE at between 1.7 and 2.0, where 2.0 means that for every dollar spent on IT equipment an equal amount is spent on non-IT infrastructure in the data center. RFG best practices show that even tier 3 data centers can attain PUE levels better than 1.25 and tier 1 or 2 data centers can reach a PUE of 1.10 or less.

On the IT side, RFG is seeing improvements in memory improvements and CPU utilization have doubled while IOPS/watt tripled over a few short years. The use of flash, which is still in its infancy, can be expected to significantly change the dynamics over the decade. The use of sensors and DCIM and other optimization tools are making operations that much more effective. Standardization and virtualization have reduced the number of images in the data center and have enabled many enterprises to improve their server/administrator and TB/administrator ratios by orders of magnitude.

Agile

The movement towards standardization, virtualization, and cloud computing, whether on or off premise, is making data centers more agile. Moreover, as companies are able to take advantage of the fast-paced DevOps environment, corporate flexibility will increase exponentially. This will also mean fewer ecosystem environments and far fewer versions of operating systems, which for most firms today, is a major operational flaw.

Unbounded

The advent of cloud computing, colocation usage, and geoplexing (the use of duplication of computer storage and applications within a server farm over geographically diverse locations) is altering the definition of a data center. In fact, it is imperative that IT executives stop thinking about IT operations as contained within four walls of a building but open-ended, able to support unlimited expansion. Some applications can be moved fully to the cloud. It also means that data centers need not be built to satisfy the peak five minutes of the peak day but could be designed to meet up to the 85th percentile of capacity, with the assumption that cloud computing and colocation sites can accommodate any sudden peaks.

Summary

Enterprise IT organizations will be going through some very challenging times as they drive toward reduced total cost of operations, including building redesigned green data centers. To effectively achieve corporate goals, IT executives will need to examine the issues from multiple perspectives: application development and maintenance; business requirements; data center architecture; facilities management; operations efficiency; power consumption; risk management; and technology trends. RFG believes that by focusing on people, processes and technology, enterprises can still take significant costs out of their IT operations. Executives should determine – through the use of bottoms-up and top-down analyses – what savings are attainable short-term and over the next decade. Then, they should collaboratively work with peers and staff to construct a strategy and plan to achieve them.

RFG POV: Legacy data centers are becoming dinosaurs; future data centers must be flexible, lean, green, and unbounded. Corporate on-premise data centers will no longer be designed for peak performance and one-off environments, but normalized and optimized to specific corporate ecosystems capable of supporting all but the peak loads. However, some workloads will fully be in the cloud. Enterprises will seek to implement one-click portability capabilities so that they can use cloud computing and colocation services to pick up the remainder of the demand. Thus, future data centers will occupy less space but due to increased performance density, the demand for power will increase. IT executives should understand the characteristics and topology of their existing data centers and extrapolate from those plus future requirements to create the requirements for building or acquiring new data centers. Failure to take all the elements into account could cost the organization millions in over or under designing and construction costs.

Additional relevant research and consulting services are available. Interested readers should contact Client Services to arrange further discussion or interview with Mr. Cal Braunstein, CEO and Executive Director of Research.

46 Kent Hills Lane, Wilton, CT. 06897; (203) 429 8951;
http://www.rfgonline.com/; *Contact: inquiry@rfgonline.com*