WP

Server Technology White Paper

Managing variable data center rack densities

White Paper STI-100-015
March 2015

HEADQUARTERS - NORTH AMERICA
Server Technology, Inc.
1040 Sandhill Drive
Reno, NV 89521
United States
1.775.284.2000 Tel
1.775.284.2065 Fax
sales@servertech.com
www.servertech.com
www.servertechblog.com

Western Europe, Middle East and Africa
Server Technology
Fountain Court
2 Victoria Square
Victoria Street
St. Albans
AL1 3TF
United Kingdom
+44 (0) 1727 884676 Tel
+44 (0) 1727 220815 Fax
salesint@servertech.com

Central Europe, Eastern Europe and Russia
Niederlassung Deutschland
Server Technology LLC
42119 Wuppertal
Germany
Tel: + 49 202 693917 x 0
Fax: + 49 202 693917-10
salesint@servertech.com

APAC
Server Technology
Room 2301, 23/F, Future Plaza
111-113 How Ming Street,
Kwun Tong, Hong Kong
Direct line: +852 3916 2048
Fax Line: +852 3916 2002
salesint@servertech.com

Server Technology
Quality Rack Power Solutions

Managing variable data center rack densities

Overview

There is no shortage of predictions regarding the future of data center power densities, but their track-record over the last decade has been poor. At the same time, there is also a general inconsistency in the way different groups define and discuss the topic of the high-density data center. This paper will touch on these predictions and definitions. Following that, it will tie the density discussion into the other main data center goals: efficiency, capacity planning, and uptime. Finally, it will provide a bit of guidance for cutting through the fog where it comes to how power and equipment densities vary within a data center and for planning how to handle the varying densities at the data center rack.

Current State of Data Center Density

In the first years of the 21st century, the data center was most often a mishmash of old mainframes, networking racks, and various server/storage cabinets with no pre-designed structure. Power density at the rack was not a topic of discussion as most were in the 1-3kW range and the power and cooling infrastructure were overbuilt for maximum uptime (see Figure 1). By 2005, the typical data center manager was exposed to increasing rack power due to new technology, especially blade servers. Yet, the new 5kW racks still proved to be only a minor challenge to the power distribution design with some moving to 3-phase circuits, and a more considerable challenge to the cooling design, bringing about the need to pre-design the data center for hot-aisle / cold-aisle configuration, watch for hot-spots, and provide for backup cooling.

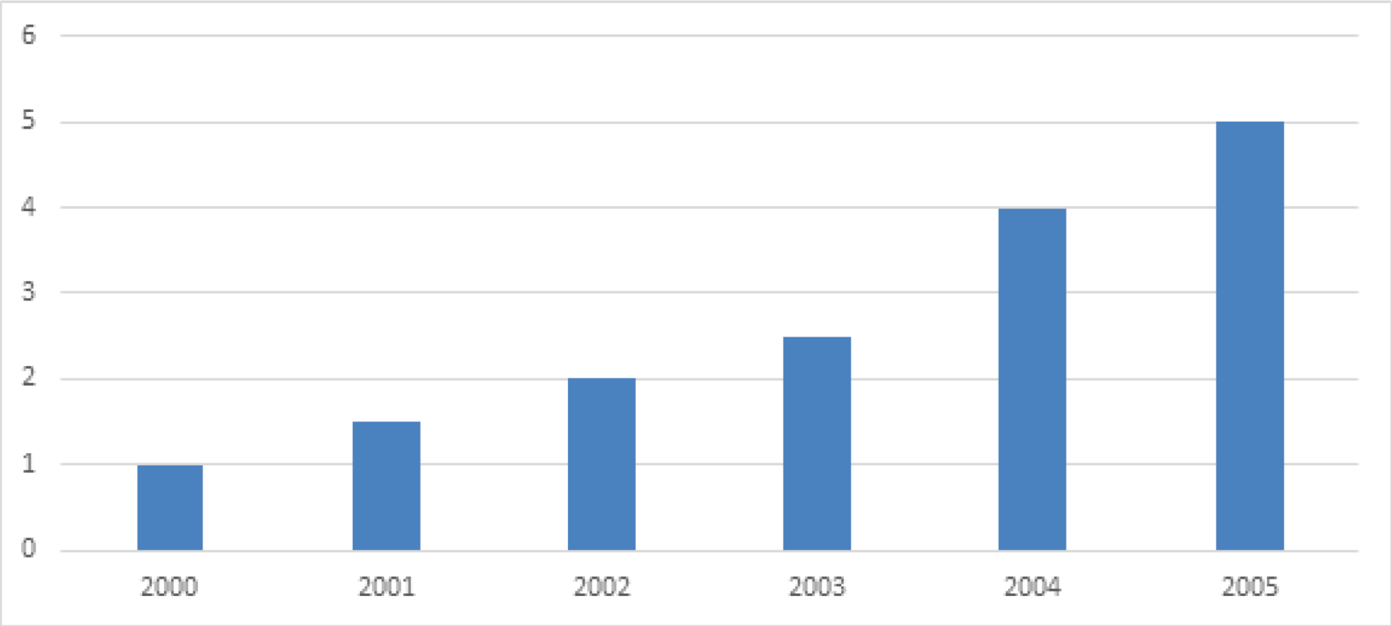

Figure 1: Average power density (in kW) per rack (IDC, 2007 et al.)

Managing variable data center rack densities

Now, a decade later, there remains the discussion of the potential for sudden rapid increases in data center power density. Of course, it is very difficult to increase power density in existing data centers, but the continual call for consolidation is leaving this possibility open for new data center construction. Additionally, the cooling strategies are becoming much more efficient and tuned for minimal needs rather than being significantly overbuilt as in the past.

What Exactly is “High Density”?

The term “High-Density” could refer to any number of things including more equipment in a rack, more racks in a data center, or even more applications on a server. For the purposes of this paper, we will concentrate mostly on the increases in total power utilization per rack and increases in equipment count per rack, though we will also touch on other needs in the data center.

The threshold for high power density set forth by Strategic Directions Group (2014) and endorsed by AFCOM was 8kW per rack, with an extreme density defined at >15kW per rack. Intel (2014) pushed the limits with 43kW per rack extending out to 1100 W/ft^2 using custom-built narrow racks and free-air cooling at elevated temperatures. This leads us to the conclusion that the term “high-density” will always be relative, no matter what any single group tries to peg it to.

Emerson Network Power reported in fall 2009 that the average power density of their surveyed respondents was 7.4 kW per rack. These respondents expected this to grow to 12 kW per rack by 2011 and 16.5 kW per rack by 2019. Strangely, by the fall 2014 survey, the average had dropped to 5.83 kW per rack and projection for two years hence had become 8.9 kW. Indeed, this semi-annual survey of the Data Center Users Group (DCUG) shows that the peak was 7.7 kW in 2012 (see Figure 2). Assuming the variation through the years from 2006 to 2014 was simply due to varied respondents, it seems there has been no increase in rack density on average. So why do these same data centers keep expecting the density to rise? And what might be keeping it stable?

High power density is often linked to high equipment density. That is in general, the more equipment you install in a rack, the more power must be delivered to the rack. We have seen this to be true with the immediate popularity of the Server Technology® HDOT™ cabinet PDU which has become the industry standard for rack equipment density. It is important to understand that the link between high equipment density and high power density is only valuable in some instances. For example, comparing one rack to another with similar equipment or analyzing one rack evolving over time will demonstrate the link between the two forms of density. On the other hand, any single install could be of very low power equipment which fully loads the RU space in the rack, or a small number of blade chassis drawing much higher power.

Another driver in terms of both rack equipment density and rack power density is increases in the heights of racks. Based on recent trends, the 42RU rack is no longer the standard for large purpose-built data centers. 45RU to 52RU have been adopted by many organizations to relieve the cost of floor space. By allowing more equipment to be installed in a rack, the typical average

Managing variable data center rack densities

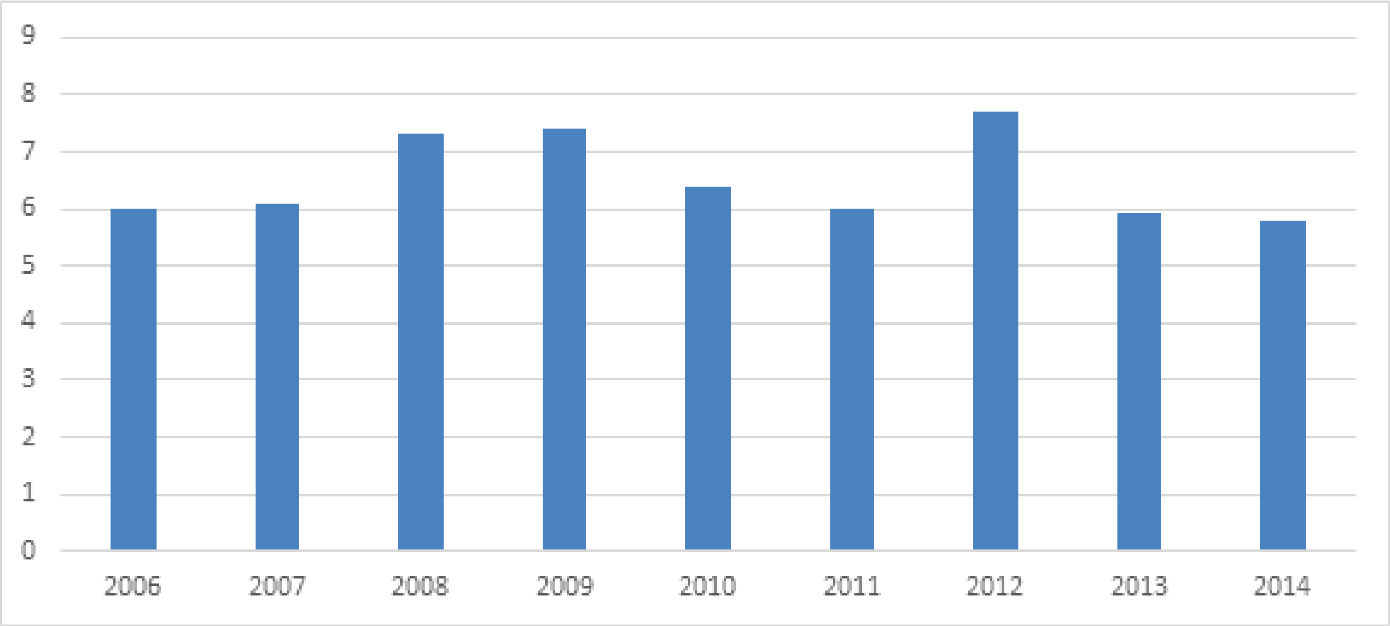

Figure 2: Average power density (in kW) per rack - Emerson DCUG data from 2006 to 2014

power density will increase from previous builds. We may be coming upon the time where we start seeing the DCUG respondent's projections come true – higher average power densities may finally be at hand.

Average Power Density vs Peak Power Density

It is important to understand and differentiate between the average rack power density and maximum rack power density across a data center floor. It is equally important to understand and differentiate between power density that is averaged over time and power density peak within a period of time. We might call the first topic "spatial power density variation" where the average spatial power density is dependent upon size of the data center and is often tied to infrastructure capacities, and the peak spatial power density is dependent on individual components within the system which is tied to the specific design aspects. For example, a 345 W/ft^2 data center (based on white space) might be specified by the fact that 1 MW of power and cooling is available for use over 2900 ft^2. Using the AFCOM standard rack area of 25 ft^2, 116 racks are deployed at an average power load of about 8.6kW each.

The second topic might be called "temporal power density variation" where the average temporal power density is dependent upon regular application loads, and the peak temporal power density is dependent upon sporadic application loads. For example, any given rack, POD, circuit, or data center will have a peak allowable power load. Within this, there will be an average power load over time at each level based on application. Continuing the example of the 8.6kW racks: if the data center simply provisions 8.6kW for each rack, the total peak load will never exceed the data center allowable, but also any single rack will not be allowed to exceed 8.6kW. This is not optimal as experience shows that the average rack load over time will be much less than the peak. This not only results in the

Managing variable data center rack densities

average load of the data center over time being much less than the allowable, but also the peak load of the data center over time being much less than the allowable for the data center. This is because only a few racks will peak at any given time, not the whole of the data center. The final result of this design is that it is a great over-provision of power.

It would be possible then to say that the average of 8.6kW per rack is useless unless you allow individual racks to exceed the 8.6kW. Thus overprovisioning of some or all racks allows the overall data center to reach toward the allowable peak. The strategies used for power overprovisioning and for maximizing usage to the limits set by electrical codes require continual monitoring during growth stages and will be further considered later in this paper.

Cost Considerations for High Density Deployments

In Schneider Electric's white paper (2014), Choosing the Optimal Data Center Power Density, the conclusion is made that exceeding about 11kW per rack in design capacity has significantly diminishing returns on reductions in cost with additional design and operational complication outweighing those small cost benefits. This is, of course, a generalization which assumes that space costs are minimum. On the flip side, average rack heights have been increasing and 17kW rack PDUs have become more and more popular. Presumably, these deployments have been determined to be best for these organizations.

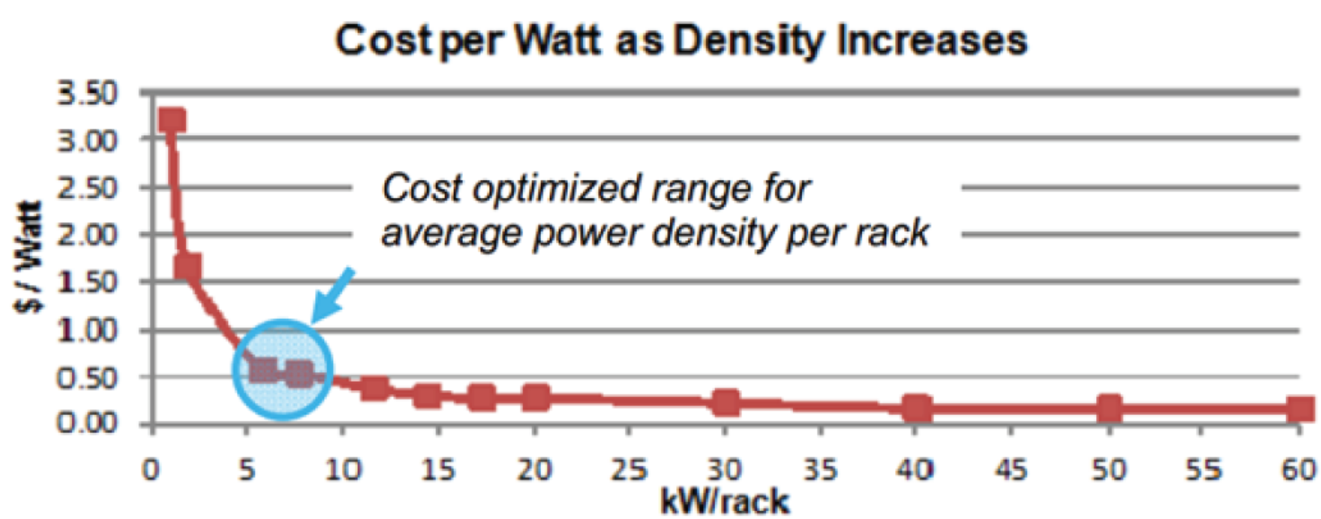

Figure 3: Cost per watt vs power density – Schneider Electric (2014)

The problem with many of these types of analyses is that they leave out the most important factor. That is, they omit the devices such as servers, storage, and network gear that are the power load in the data center. The process of arbitrarily dividing power load between a certain number of racks and then analyzing the cost of infrastructure and cooling does not take into account the IT transactions taking place within the data center. These transactions include the IT refresh cycle in which newer equipment is installed to provide much better performance but with a nontrivial increase in power consumption. They also include variable application loading over time that may be driven by changes in projects, personnel, or market growth. They even include “green” goals such as driving PUE downward by pushing more IT load. These density-based IT transactions will be revisited later in this paper.

High Power Density vs Other Data Center Business Drivers

The discussion of increased power density at the rack does not lie in a vacuum. In fact, it may be a cause or an effect of various other changes or improvements in the data center. Virtualization, consolidation, automation, cloud deployment, modularity, free cooling, big data, hot- or cold-aisle containment and any number of other meaningful concepts interrelate to power density through the topics of efficiency, capacity planning, and uptime.

Efficiency and PUE

Efficiency of the power system, including the cooling components has been the hottest topic in the data center for the last several years. Calculation of PUE is being performed by the majority of medium to large scale data centers today. To truly maintain an understanding of the efficiency of the data center, the data center manager needs a continuously operating monitoring system such as Server Technology's SPM (Sentry Power Manager). Through data trending and reporting, power management at the rack PDU, and rack environmental monitoring, data center personnel can keep a close eye on the relationship between efficiency and power density. Figure 4 shows that the average PUE in the data center has leveled off over the last few years. Adoption rate of PUE measurement continues to hover around the 70% mark. Increasing power densities may push for increased measurement and reduced PUE.

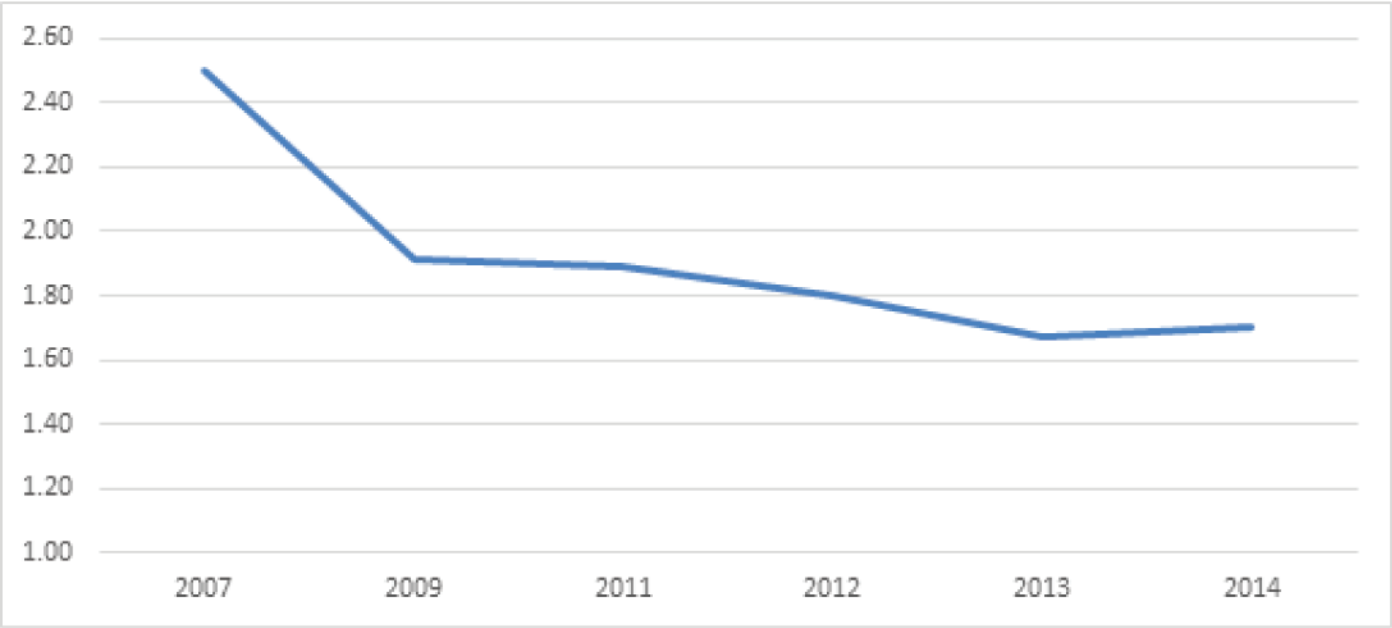

Figure 4: Average PUE – Uptime Institute Surveys (Symposium 2014)

Managing variable data center rack densities

Quite often, the mantra that consolidation and virtualization leads to better efficiency is driving the data center to higher power densities. If a rack has been loaded with equipment such that it is already pushing the allowable power load, virtualization that increases utilization will also reduce the number of devices that will fit within the power limits of the rack. This may be more efficient in the sense that less power will be used to perform the same job, but the sparser distribution of equipment may reduce the efficiency of the cooling system. Determining the overall level of the benefit of this change can only be achieved through measurement.

One of the hot topics in efficiency is the idea that cooling costs are reduced by allowing temperature to be increased. This may be independent or coincide with increasing rack power density. Higher rack power densities directly increase the total heat load per unit volume of the data center. Allowing the inlet temperatures to increase at the same time has the potential to dramatically reduce the fault response time. This pushes the data center manager to opt for more complete monitoring of temperature all the way down into the rack, and for more levels of control including outlet management using Switched cabinet PDUs.

Capacity Planning

DCIM and other capacity planning tools have taken hold in many data centers today. Effective use of the available power infrastructure and building in a modular fashion have led to the data center operating closer to peak power capacity at any given time. One possible way to increase capacity in an existing data center is to increase the rack power density. Of course, the power and cooling infrastructure throughout the data center would typically need to be upgraded to allow for this. Keeping track of the percentage of capacity used as it increases over time is important to properly time these types of upgrades, or of modular expansion of the data center. Tools such as SPM can provide a means to monitor direct measurements over time as well as providing a means to predict when capacity will hit its limits in the future via predictive trending. This can be accomplished at every power distribution point within the data center from UPS down to the outlets of the cabinet PDUs, not only for capacity predictions, but also as guidance for deployment of new equipment within the established power limits. Similarly, keeping track of the free physical RU space in the cabinets within the data center can be tracked using the same monitoring tools.

Understanding the curve of power density growth is critical to the further planning of data center deployments and application installations. By watching trends of power usage cycles over days, weeks, months and beyond, expansions of the infrastructure to allow for power density increases can be executed in a controlled manner. Finally, this predictive method can be extended to the financial side of the business for long-term projections regarding needs for future data centers, including guidance for deciding the design-to power densities.

Uptime

The top priority of most enterprise data centers remains uptime. Increased power density can put pressure on the power redundancy components and reduce time to react on cooling component faults. This has usually been the number one reason to monitor power conditions in the data center with tools such as SPM. Though this may still be strictly true for some data centers, it has been quite different for many others. With the increased power density due to consolidations and virtualization, additional data and network redundancy has been implemented. This additional redundancy, in some cases, means that the uptime requirement for any particular rack is greatly reduced. In some very large organizations, whole data centers can go down without affecting the business, and are therefore designed with low levels of redundancy (e.g. Uptime Institute Tier 2 instead of 3 or 4). These are very interesting cases, but for the rest of us, the cap-ex of creating such redundancy is hard to come by.

The typical enterprise data center is built with power redundancy delivered from the source power plants all the way down to the rack devices through the use of a pair of cabinet PDUs. If any one of the distribution points fails to handle the loss of power to the redundant partner, significant downtime is likely to occur. This makes it even more important to have a monitoring system to keep track of each “side” of the power chain and the combined loads per phase along the way. Features such as the Circuit modeling and Cabinet Redundancy integrated into SPM help data center personnel maintain awareness of this as power and equipment densities increase over time.

Variable Power Density at the Rack

Taking into account the goals in the enterprise data center we have discussed such as efficiency, capacity planning, uptime, and power density along with the current state of the industry and the expectation of future technology, we come to the question of what can and should be done in a new data center design. What is the optimal density to provision for and what should be the expectation of variations in actual power usage across the data center and over time?

Rack to Rack Variation

According to Schneider’s data, oversizing racks to handle peak to average power ratio of 2:1 comes at < 1% cost premium to the full data center when keeping all upstream distribution built to the spatial average power density, oversizing a group of racks or POD tied to a large PDU comes at a 4% cost premium to the full data center, and oversizing the full data center comes at a 20%+ cost premium. Their recommendation is to oversize the racks as long as there are procedures to prevent overloading of upstream systems. As they define this strategy, “Rack PDUs, breaker(s), cables sized to handle expected peak density in any rack within the pod, but the pod will not exceed the average overall.”

Some may object to the strategy with the concern that any particular individual rack could be loaded such that it will be safe while pushing a pod over the safe limits. This speaks to the logistics of the deployment process. In fact, this strategy is standard practice at the rack PDU level already. Case in point: a 30A single-phase power strip in North America is rated at 24A continuous usage, but it will be designed to have two (2) branches protected by 20A circuit breakers or fuses. An installer could easily trip the panel breaker without tripping the internal branch circuit protection. This all comes down to risk tolerance vs logistical control, which is a topic I will leave to others.

Managing variable data center rack densities

This is where we need to be very clear on the definition of "average". The consequences of high usage at the rack level oversubscribing the pod may be very damaging to the business. The pod must still be designed to handle the peak loading expected at any given time. Let's assume a pod of 20 racks is designed to handle a peak loading of 100kW which gives a spatial average loading per rack at temporal loading peak of 5kW. This corresponds to the North American safety rating of a 208V/24A circuit. It is useful here to take a tangent to recognize that the 100kW pod must be built to handle 125kW per the NEC in the USA, so we should only analyze this either from the full maximum rating standpoint or from the safety rating standpoint. For consistency, we will use the safety rating. If the racks in this pod are deployed with the 208V/24A circuit, each rack must always be below this value. It is a good time to note that rack power distribution units do not lie on a continuum of power capacities. There are discrete values based on standard connections for various parts of the world.

N. Amer.		1-phase			3-phase	
Voltage	16 A	24 A	48 A	16 A	24 A	48 A
120V	1.9 kW	2.9 kW		5.8 kW	8.6 kW	17.3 kW
208V	3.3 kW	5.0 kW	10.0 kW	5.8 kW	8.6 kW	17.3 kW
240V	3.8 kW	5.8 kW	11.5 kW	11.5 kW	17.3 kW	34.6 kW
Europe Voltage	16 A	32 A	63 A	16 A	32 A	63 A
230V	3.7 kW	7.4 kW	14.5 kW	11.0 kW	22.1 kW	43.5 kW

Table 1: Standard circuit capacities for data center racks – 3-phase circuits must be balanced

Short-term Variation

Variation in power load across the data center is only the beginning of the design plan. An understanding of power usage over time based on application load variability is needed to minimize the oversizing of power distribution components. Load balancing and batch process scheduling will help keep the temporal variation in power density from exceeding the peak allowable load. We will find in practice that the temporal average power usage of the rack, and thus the pod will be closer to 50% of the design max, or 2.5kW per rack and 50kW per pod. Of course, this is the safe route, but it is not the most efficient use of resources. Instead, if the racks are all deployed with 208V/24A 3-phase which provide 8.6kW and those racks loaded with 4.3kW on a temporal average, the pod will see 86kW average over time. This allows any three of the twenty racks to double their load or all twenty racks to increase load by 16%, while the pod remains under the 100kW rating.

Figure 5 shows how a particular server might vary its power draw based on its CPU utilization. This utilization is directly tied to application processes and shows that any single server could see an increase in power draw between a typical load state of 30% and "full" utilization of 80%. Also shown is the maximum number of servers of this type that can be racked on a 30A 3-phase 208V circuit based on these various levels of utilization. This data is important when trying to optimize the design densities.

Managing variable data center rack densities

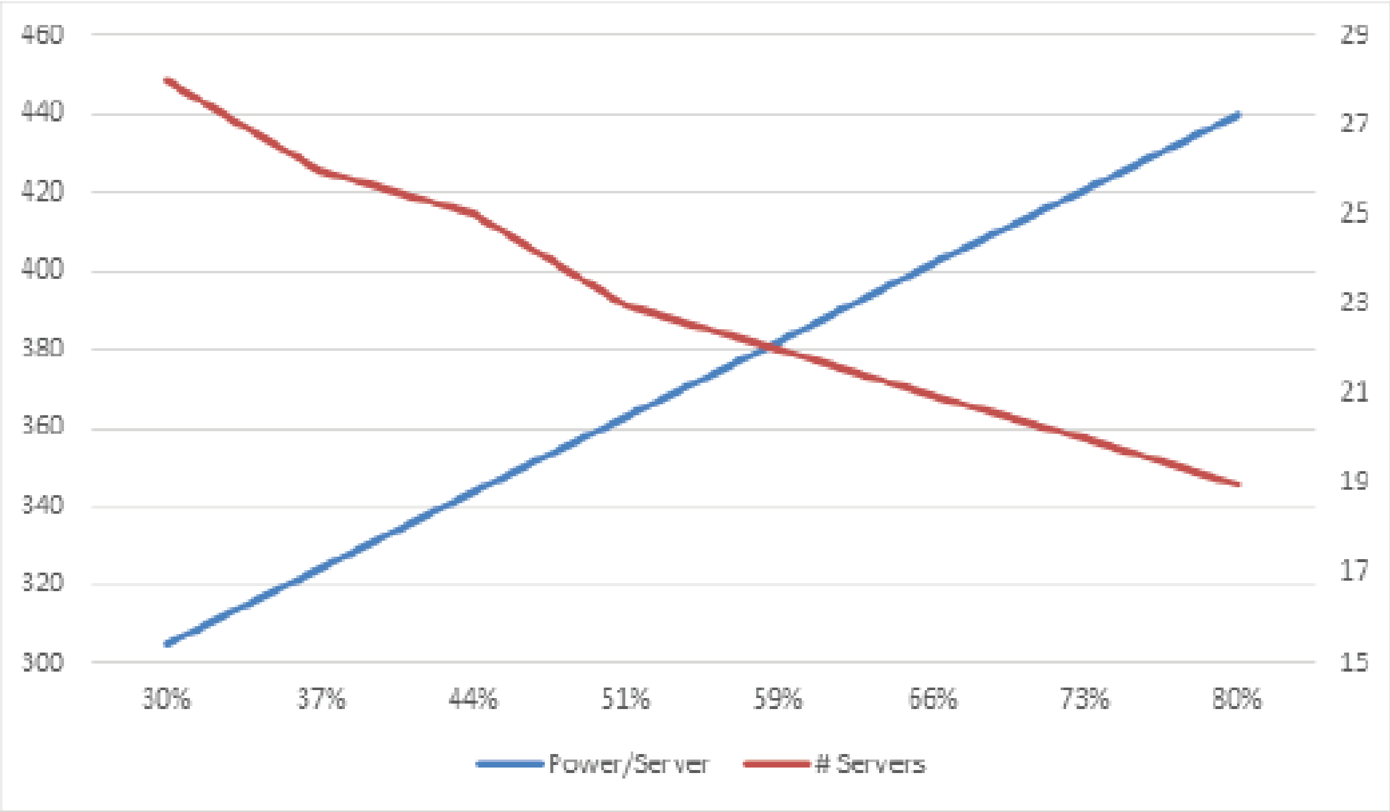

Figure 5: Power per server and Servers per rack vs. CPU usage – 30A 3-phase 208V circuit.

Figure 6 shows the effect of Virtualization using an example that peaks at an 8:1 ratio using the same server model as above. Here the virtualization causes the utilization to rise upwards toward 80%. The trace shows the effect on the percentage of the number of racks to obtain the same performance as the reference servers running at 30% utilization. As the utilization increases to 80% through the addition of more virtual servers, the number of racks full of servers reduces to 18% of the reference design.

Long-term Variation

Planning for the future of the data center is difficult but must be a part of the density discussion. With modular deployment being all the rage, it is important to plan for the process of growth. The strategy of oversizing the racks allows one to continue the modularity at the pod level. That is, a proper modular data center construction will allow additions of pods with their corresponding electrical and mechanical infrastructure as long as physical constraints such as floor space and weight limits allow. The trickier part of this prognostication is the equipment refresh cycle. A typical refresh of computing servers on a three-year cycle might see compute performance quadruple while power load increases at 25% (about 50% is typical for the processors themselves). From

Managing variable data center rack densities

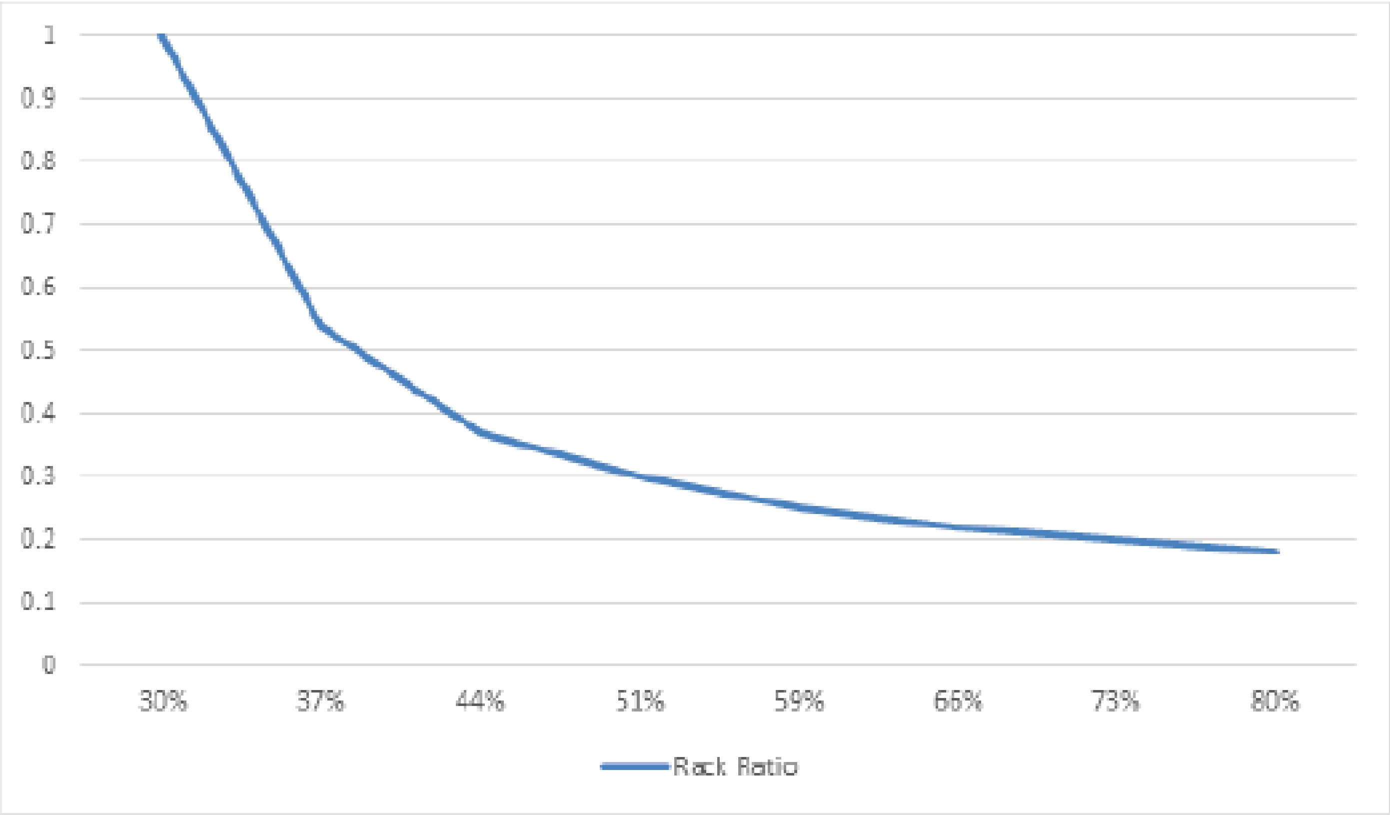

Figure 6: Effect of virtualization in regards to consolidation.

the power efficiency standpoint, this is a no-brainer. The consequence though is that the racks in the example will see these increases and we already know that we can only absorb a 16% increase in overall load in the example given. This results in lower server/rack densities and thus reduces the overall performance boost of the refresh at the data center level. If this data center is expected to operate for 5 refresh cycles (15 years), it will see power loading increase to 3x its original load for its servers (much less for its network and storage gear) for 1024x performance boost. Even if only half of the equipment are servers, there will probably be double the original total data center power load if space is not sacrificed.

Figure 7 shows the effect of server refresh cycles over time when the goal is to keep the racks full of equipment. This provides full access to the performance improvements over time, but forces an infrastructure design that allows for double the power loading, as previously mentioned. This method is difficult to achieve in practice, though it does provide for consistency in the cooling system, thus keeping PUE minimized.

Managing variable data center rack densities

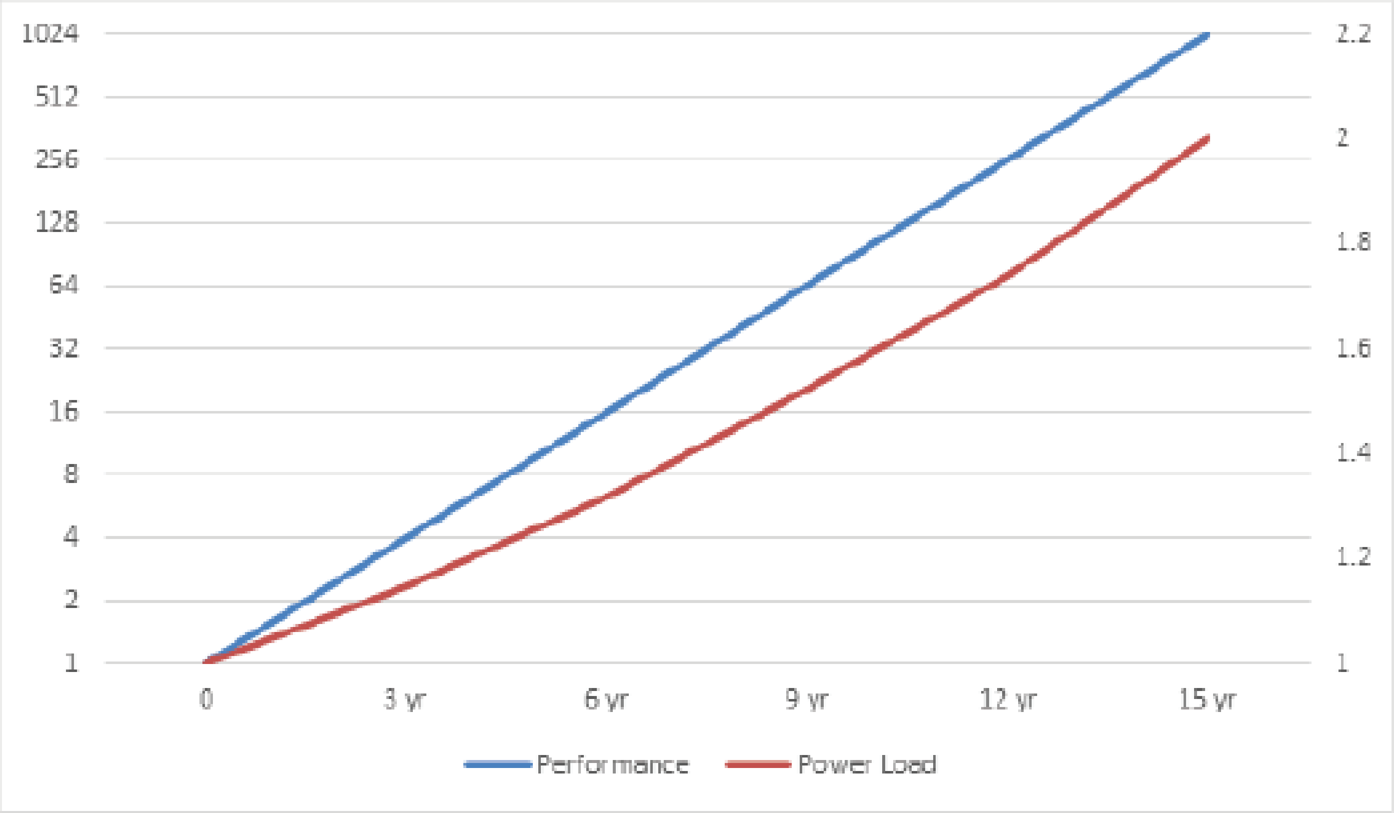

Figure 7: Refresh cycle effects over data center life – constant space utilization

Figure 8 shows the effect of server refresh cycles over time when the goal is to keep power usage constant. In other words, the infrastructure is designed and loaded to the maximum from the outset. The result of this strategy is that space utilization falls over the life of the data center and performance is only about half of the Figure 8 case after 15 years.

This leads to the crux of the density question and the fact that there is not one right answer. The strategy of simply oversizing rack PDUs to allow for high density racks or moments of high power usage does not provide sufficient long-term results. At the same time, designing for a very long-term potential power density has been shown to be at the expense of near-term cap-ex and op-ex. The best solution probably lies somewhere in the middle. Oversizing the power distribution at the POD level and building all infrastructure in a highly modular fashion will allow for the most flexibility and agility in the long-term.

Managing variable data center rack densities

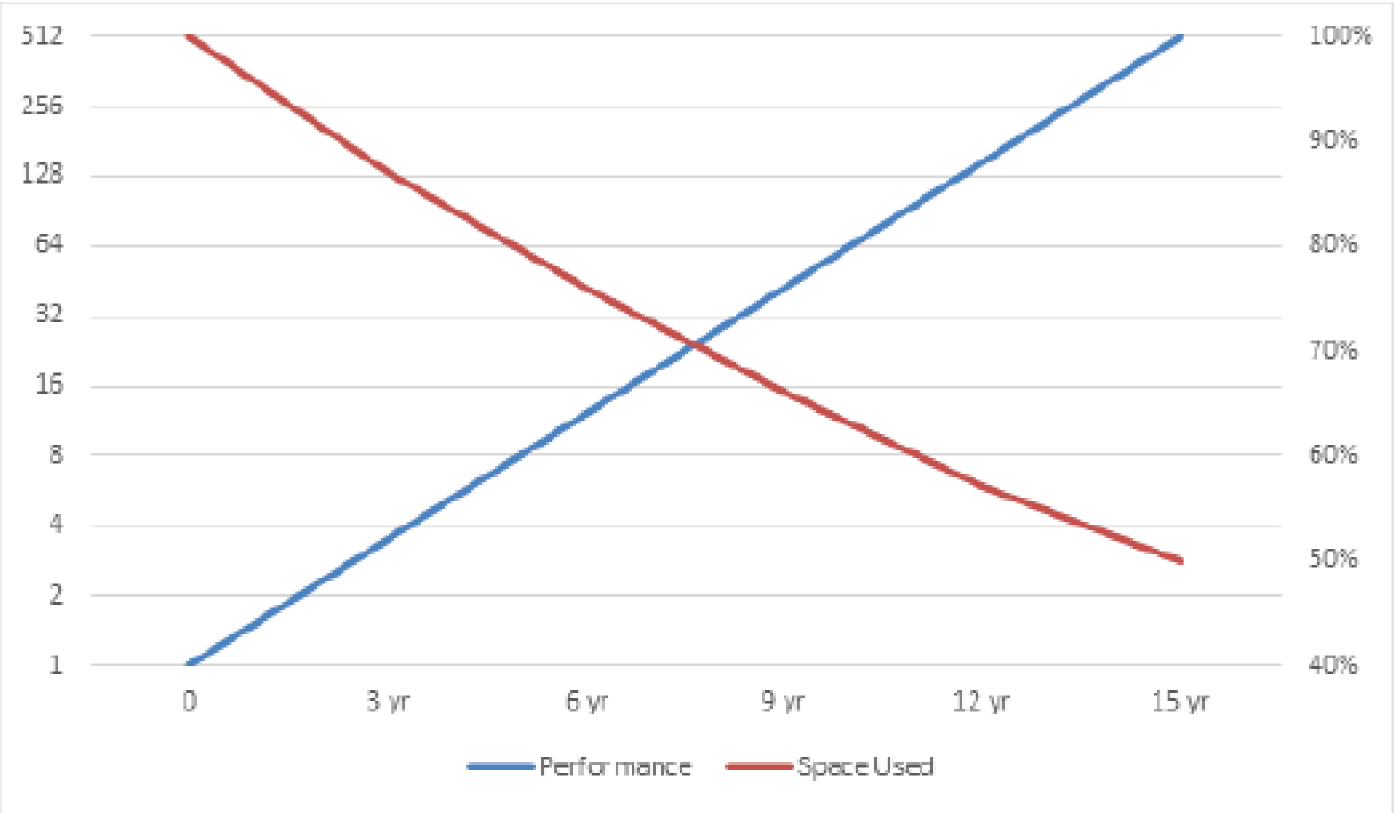

Figure 8: Refresh cycle effects over data center life – constant power utilization

Summary

The data shows that, over the last decade, actual data center operators have greatly overestimated the expected increases in power density. The expectation seems reasonable from the standpoint that server performance per watt and watts per RU continue to increase exponentially. The data center; however, seems to have been stuck in the 5-8kW per rack range. The reason may simply be that the old data centers continue to operate with the same infrastructure without upgrades. In this case, racks actually become less dense in terms of their equipment count.

This bottleneck may be about to open up with newer data centers built on modular platforms and consolidation initiatives driving enterprises out of their old infrastructure. Indeed, the 52RU 17kW per rack data center is becoming more and more common. Because of this increase in power and equipment density, continual monitoring of power and environmental conditions through DCIM products such as the SPM by Server Technology has become ever more important to align with the goals of efficiency, capacity planning, and uptime.

Managing variable data center rack densities

With the understanding of the variation that can occur from rack to rack, from changes in application loading within the racks, and from IT refresh cycles, it is clear that the question of the optimal design-to power density does not have merely one right answer. It is important to model the expectations in physical equipment density in the rack as refresh cycles change the power density per device. This is likely to lead to an initial deployment that maximizes the device count in the rack which will result in the need for high outlet density cabinet PDUs such as the HDOT by Server Technology. As refresh cycles reduce the device count within the racks, overload risks of higher power equipment can be mediated and 3-phase load balancing can be easily maintained with color-coded Alt-Phase Switched outlets available on these cabinet PDUs.

References

1. IDC 2007. John Humphreys, Jed Scaramella, Enterprise Platforms Group. The Impact of Power and Cooling on Data Center Infrastructure.
2. Strategic Directions Group Pty Ltd. 2014. Data Center Standards – Data Center Size and Density Version 3.0. Level 9, 145 Sinnathamby Blvd, Springfield, Qld 4300, Australia. ABN 58 103 746 872.
3. Intel 2014. John Musilli, Paul Vaccaro. Intel IT Redefines the High-Density Data Center: 1,100 Watts/Sq Ft.
4. Emerson Network Power 2009. The Four Trends Driving the Future of Data Center Infrastructure Design and Management. 1050 Dearborn Drive, PO Box 29186, Columbus, Ohio 43229. WP167-39, SL-24637
5. Emerson Network Power 2013. Integrating People, Process and Technology to Transform Data Center Operations and Performance. 1050 Dearborn Drive, PO Box 29186, Columbus, Ohio 43229. SL-25674
6. Emerson Network Power 2014. Fall 2014 Data Center Users' Group Survey Results – October 2014. Data Center Users' Group.
7. Schneider Electric 2014. Choosing the Optimal Data Center Power Density. White Paper 156.
8. Uptime Institute 2014. Symposium Data Center Survey Results.