



IT White Paper

**DATA CENTER SPACE OPTIMIZATION:  
*POWER AND COOLING STRATEGIES  
TO OPTIMIZE SPACE IN THE  
CONTROLLED ENVIRONMENT***



## Summary

With data center rack densities increasing exponentially and blade servers promising to push them much higher faster than many expected, IT managers are facing difficult decisions regarding critical system infrastructure.

System densities are now the main driver of data center space requirements. Newer high-density systems may have smaller footprints than the systems they replace, but they require more support in terms of power protection and cooling. As a result, many organizations find themselves facing a difficult decision – expand their data center facilities or limit their adoption of new technology.

Fortunately, a third alternative exists: re-evaluate support system design in light of the increasing pressure on data center space. There are several steps that can be taken to reduce the impact of power systems on the data center, while new approaches to cooling enable existing space to accommodate a greater number of high density racks.

*One of the ways data center managers are dealing with this situation is to increase rack spacing, essentially distributing the heat from the equipment over a larger space. This is, at best, an interim solution.*

## **Increasing Pressure on the Data Center**

Technology compaction is enabling equipment manufacturers to deliver more and more processing power in less and less space. The resulting “high density” systems consume increasing amounts of power and generate correspondingly high amounts of heat.

The impact of this trend on data center space utilization has taken some organizations by surprise. After all, shouldn’t smaller, more compact systems consume less space?

The answer, unfortunately, is no. Smaller footprints do not necessarily translate into reduced space requirements because performance increases have primarily been achieved by packing more transistors operating at faster speeds into smaller processors. This results in increased free-space requirements to remove the concentrated heat.

For example, the Intel Pentium 3 processor, introduced in 2000, had 28,100,000 transistors and a 1GHz clock speed, creating a maximum power consumption of 26 Watts. The Pentium 4, released just two years later, had almost twice as many transistors and three times the clock speed (55,000,000 transistors with a speed of 3GHz), creating a maximum power consumption of 83 Watts. So, while a Pentium 4-based system may have the same footprint as

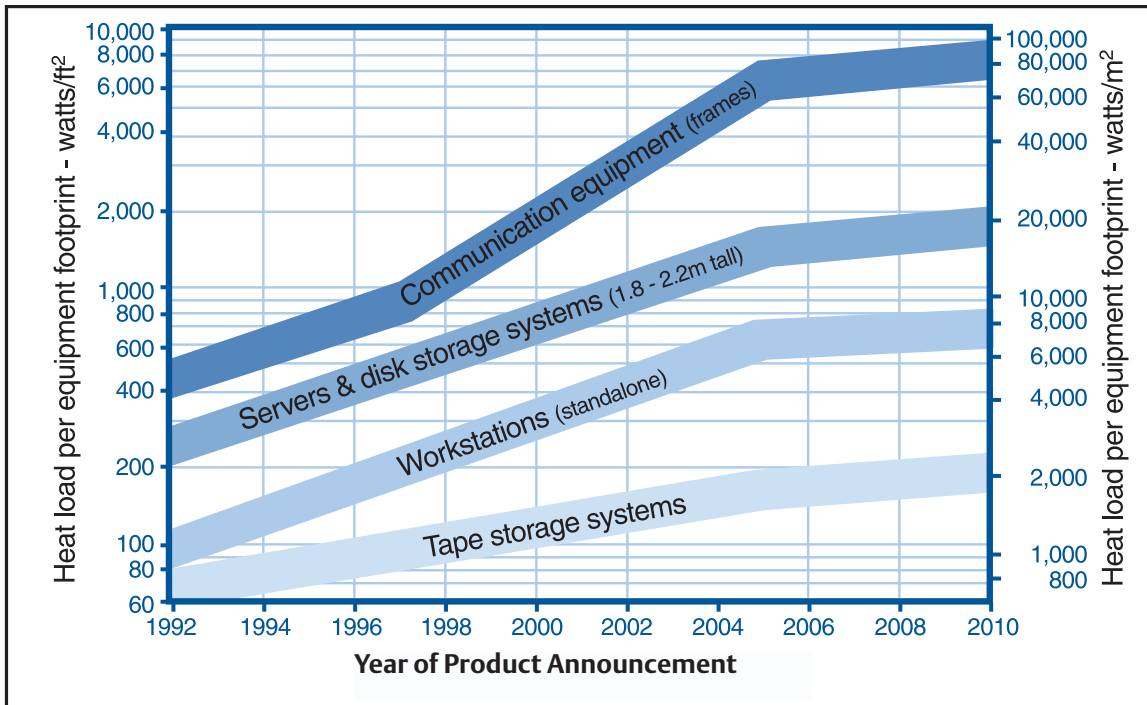
a Pentium 3-based system, it consumes significantly more power and generates more heat.

In parallel with processor advancements, server package sizes have been shrinking substantially. Much of this compression is enabled by smaller disk drives, power supplies and memory formats. New blade server packages are further condensing the computing package, creating even higher density racks.

This increases power and cooling system requirements. High-density systems may even generate so much more heat than the systems around them that they create “hot spots” in the data center, where temperatures directly above the equipment are hotter than the rest of the room. One of the ways data center managers are dealing with this situation is to increase rack spacing, essentially distributing the heat from the equipment over a larger space.

This is, at best, an interim solution. Considering the rate at which heat densities are rising (see Figure 1), relying on increased spacing alone will quickly consume available data center space and reduce the number of racks that can be supported in a particular facility.

The performance potential of high density systems can only be realized if the corresponding rise in heat density — and its implications on data center space — are successfully addressed.



**Figure 1. Heat densities of data center equipment are projected to continue to rise throughout this decade. (Chart reprinted with permission of The Uptime Institute from a White Paper titled *Heat Density Trends in Data Processing Computer Systems and Telecommunications Equipment, Version 1.0.*)**

### Data Center Economics

The data center is a unique environment within most organizations. Generally, it requires more precise environmental control, enhanced power protection and tighter security than other space within a facility. Consequently, its cost per square foot is much higher than general office space. This means the increased pressure on data center space, if not dealt with effectively, can have a significant economic impact.

Consider a typical 10,000-square-foot data center. Assuming average rack power densities of 1 kW, approximately 35 percent

of data center space is used for racks. The remaining 65 percent of space is required for aisles and support systems. Since a typical rack consumes about 7 square feet of floor space, this facility can support a maximum of 500 racks of 1 kW each.

If average power density increase to 10 kW per rack, with no other changes in data center infrastructure, increased rack spacing is required to spread the higher heat load across the room. Now only 3.5 percent of the space in the room is available to racks. The remainder is required for aisles and support systems. As a result, the facility could support only 50 racks.

***Unless an alternative cooling system that enables closer spacing of high density racks is deployed, it will be necessary to expand current facilities to support high density systems.***

Assuming a cost of \$175 per square foot for the data center shell, not including power and environmental systems, the shell cost for a 10,000-square-foot facility is \$1.75 million. This cost can then be divided by the number of racks in the room to calculate shell costs per rack. At rack densities of 1 kW, that cost is \$3,500 per rack (\$1,750,000 divided by 500 racks). When rack densities increase to 10 kW, the cost per rack jumps to \$35,000 (\$1,750,000/50).

This illustrates the potential impact of increasing densities on data center space and costs. In reality, this transformation is happening gradually and incrementally. However, it *is* happening. Unless an alternative cooling system that enables closer spacing of high density racks is deployed, it will be necessary to expand current facilities to support high density systems.

## **Cooling Systems and Data Center Space**

Traditional precision air conditioning units have provided effective cooling in thousands of data centers around the world; however, as system densities increase, they are being stretched to the limits of their practical capacity.

The key limitations involve the number of precision air conditioners that can be installed in a given room and the amount of air that can be pushed through perforated floor tiles in a raised floor.

Floor-mounted precision air systems take up data center floor space, limiting how

many systems can be installed in a data center facility. In addition, there is a physical limitation to how much air can be efficiently distributed through the raised floor. Trying to push too much air under the floor can create negative pressures that can actually draw cool air back down through the perforated tiles, rather than forcing it up into the room. In addition, the floor tiles themselves have physical limits as to how much air can actually pass through the perforations. Consequently, increasing cooling capacity will not necessarily result in a corresponding increase in cooling across the room. Computational fluid dynamics (CFD) performed for the Liebert Extreme Density Test Lab show wide variations in temperature across the room when raised floor cooling alone is used to cool high density racks.

There are several steps that can be taken to optimize the efficiency of the raised floor system. The first is an examination of the cabling running under the floor to ensure it is not obstructing air flow.

Floor height also plays a role. Doubling the floor height has been shown to increase capacity by as much as 50 percent. Data center managers planning new facilities should consider floors higher than the traditional 18-inch height. However, replacing the floor is usually not an option for existing data centers because of the disruption in operations it requires.

The hot aisle/cold aisle concept can be employed to increase cooling system efficiency. It involves arranging racks in a way that separates the cool air coming up from the floor from the hot air being discharged

from the equipment. Racks are placed face-to-face and floor tiles are positioned so cool air is distributed into this “cold” aisle, where it can be drawn into the rack. Heated air is then exhausted through the back of the equipment into the “hot” aisle. By supplying the cooling system with a smaller volume of hot air than a larger volume of mid-temperature air, more of the cooling system’s capacity is utilized.

Rack spacing can, of course, also be used to dissipate heat from high density racks, if data center space allows. In field tests, raised floor cooling systems have shown a practical cooling capacity of 2 to 3 kW of heat per rack. This means a 10 kW/rack system would require cold aisle widths of 10 feet to ensure adequate heat dissipation. Clearly, this will not prove to be a long-term solution as rack densities rise to 20 kW and beyond. A more effective long-term solution must be developed to support the continual deployment of new systems.

### Complementary Cooling

Complementary cooling has emerged as a way of supplementing room-level cooling, air filtration and humidity control provided by precision cooling systems that distribute air through a raised floor. Zone-based to efficiently focus cooling where it is needed most, complementary cooling units are placed close to the source of heat, often mounted at the top of the rack or on the ceiling.

These systems can be easily added to the existing infrastructure to reduce on a per rack basis.



**Figure 2. Complementary cooling systems, such as the Liebert XDV, bring targeted cooling close to the source of heat to supplement the capacity of traditional cooling systems.**

As mentioned previously, a 10,000-square-foot data center would have a shell cost of approximately \$1.75 million, excluding power and cooling. Cooling costs for a room can be calculated by multiplying the cost per kW of cooling by the cooling load.

Assuming cooling costs of \$925 per kW and a cooling system designed for densities of 50 Watts per square foot, shell and cooling costs for the 10,000 square foot data center used in the previous example would be approximately \$2.2 million (\$1.75 million in shell costs plus \$464,000 in cooling costs.)

This data center could effectively cool 500 racks with an average power per rack of 1 kW. Infrastructure costs amount to \$4,428 per rack.

When rack densities increase to 10 kW, shell and cooling costs stay the same, but now the same infrastructure can support only 50 racks. Consequently, the cost to support each rack grows by a factor of 10 to \$44,280 per rack.

*Complementary cooling systems can be easily added to the existing infrastructure to reduce costs on a per rack basis.*

*Adding complementary cooling will, in most cases, prove to be the most viable, cost-effective approach to managing high-density systems.*

Rack Density	No. of Racks Supported		Cost per Rack	
	Raised floor only	Raised floor with XD	Raised floor only	Raised floor with XD
1kW	500	-	\$4,428	-
10kW	50	386	44,280	11,644

**Figure 3. Complementary cooling can increase the number of racks that can be supported within a particular facility. This table shows the number of racks that can be supported, and the cost to support each rack, with and without complementary cooling in a 10,000 square foot facility originally designed for densities of 50 Watts per square foot.**

If complementary cooling is added, cooling system costs increase by about 65 percent. However, the additional cooling capacities allows for higher heat densities and closer equipment arrangement. Now, the 10,000 square foot data center can support 386 racks of 10 kW each – a 700 percent increase in capacity. This drives support cost per rack down to \$11,644 – 26 percent of the cost of a solution that relies on rack spacing. And, the reduction in the number of racks that can be supported is primarily a result of the increased UPS and power distribution equipment required to support the larger load, not a limitation in the cooling system.

Adding complementary cooling will, in many cases, prove to be the most viable, cost-effective approach to managing high density systems. Complementary cooling systems also have the advantage of addressing problems of efficiency created because higher heat loads may not be evenly distributed throughout the room, creating extreme hot spots. Further, this

approach has additional benefits not explored in this analysis in regard to the lower operating costs that will be realized from the efficiency of the supplemental cooling system.

### **Power Systems and the Data Center**

The location and footprint of power systems also have an impact on data center space, and these will vary based on whether a room- or rack-based protection strategy is being utilized.

The room-based strategy centers on a UPS system sized to provide backup power and conditioning for the entire room. Often, this approach benefits from the cost advantages that come with choosing a larger UPS system at the onset, rather than piecing together a system over time. It also provides added application flexibility by allowing the UPS system to be located outside the data center in a lower cost-per-square-foot area of the building.

A rack-based approach provides protection on a rack-by-rack basis, as individual UPSs are purchased and installed with each addition of network equipment. This approach is often adopted in smaller facilities that are not expected to grow and in cases where future capacities are impossible to project. While a rack-based approach may seem cost-effective, it is important to evaluate the implications of this approach in terms of both space and dollars.



**Figure 2. New UPS systems, such as the Liebert NX, are being designed with compact footprints to minimize their impact on data center space. A 30 kVA Liebert NX has a footprint of just 24 inches by 32.5 inches**

First, this approach typically does not provide the option of placing power protection equipment outside the data center. This means the UPS systems take up valuable floor space and add to the heat load in the data center.

Room-level UPSs may also be located inside the data center, but when they are, they consume less floor space and generate less heat than highly modular systems. That's because the larger the UPS, the higher its efficiency rating. This puts highly modular systems at a disadvantage because multiple, distinct UPSs are required to achieve a certain capacity. For example, it might take three modular systems to provide 120 kVA of protection, each operating at a lower efficiency than a single 120 kVA system would. This difference in efficiency translates directly into increased "heat dump" from the UPS system.

A similar scenario holds true in terms of footprint. A highly modular system will typically have a larger footprint than a fixed capacity UPS system – and the footprint differential increases with system capacity. A 40 kVA modular UPS requires 53 square feet of floor space, allowing for service clearances, while a 40 kVA fixed capacity system requires just 40 square feet of space – a 32 percent difference. At capacities of 120/130 kVA, the footprint of the modular system grows to 159 square feet, while the traditional system consumes just 75 square feet.

More importantly, the fixed capacity system will deliver greater reliability and availability than the highly modular system because:

*...the fixed capacity system will deliver greater reliability and availability than the highly modular system...*

- The modular systems utilize more parts than the fixed capacity system, increasing the UPS hardware failure rate and the risk of a critical AC output bus failure;
- The fixed capacity system includes module-level redundancy to enable concurrent maintenance while the highly modular system does not;
- The fixed capacity system provides longer runtimes using the in-cabinet batteries than the highly modular system, which typically require external battery cabinets to achieve desired runtimes.

Modular power protection systems may prove suitable for some applications, but facilities that expect to experience growth should consider the long-term impact on data center space and UPS system costs before embarking on a protection strategy based on this approach.

## Conclusions

Data center space utilization will be an increasingly important consideration as technological progress requires ever-increasing power levels and corresponding cooling requirements. The location and selection of support systems can have a significant impact on data center space.

Making decisions that take into consideration all contributing factors, from risk associated with heat spikes to placement of power systems to data center shell costs, can ensure that increasing rack densities do not drive the need for expanded or new data center facilities in the future. Careful decisions can also ensure that business needs — not support system limitations — drive a company's adoption of new technologies.



**LIEBERT CORPORATION**

1050 DEARBORN DRIVE

P.O. BOX 29186

COLUMBUS, OHIO 43229

800.877.9222 (U.S. & CANADA ONLY)

614.888.0246 (OUTSIDE U.S.)

FAX: 614.841.6022

**[www.liebert.com](http://www.liebert.com)**

While every precaution has been taken to ensure accuracy and completeness in this literature, Liebert Corporation assumes no responsibility, and disclaims all liability for damages resulting from use of this information or for any errors or omissions

© 2004 Liebert Corporation. All rights reserved throughout the world. Specifications subject to change without notice.

Trademarks or registered trademarks are property of their respective owners.

© Liebert and the Liebert logo are registered trademarks of the Liebert Corporation

© Keeping Business in Business is a registered trademark of the Liebert Corporation.

The Emerson logo is a trademark and service mark of the Emerson Electric Co.

Printed in U.S.A. 0704 WP401

