



White Paper

Data center TCO; a comparison of high-density and low-density spaces

M.K. Patterson, D.G. Costello, & P. F. Grimm
[Intel Corporation, Hillsboro, Oregon, USA](#)

M. Loeffler
[Intel Corporation, Santa Clara, California, USA](#)

Paper submitted to THERMES 2007 for publication in Santa Fe, NM Jan. 2007

Contents

Abstract 2
Motivation 3
Background and Definitions 4
Example Data Center 6
Benchmark Data 9
Conclusions 11

Keywords:

- Data Center
- Thermal Management
- TCO
- High-Performance Computing

Abstract

The cost to build and operate a modern Data Center continues to increase. This Total Cost of Ownership (TCO) includes capital and operational expenses. The good news in all of this is the performance or compute capabilities in the same Data Center (DC) is increasing at a much higher rate than the TCO. This means the actual cost per unit of compute performance is coming down in the Data Center.

While that is a positive trend the increasing densities still present a challenge. This challenge though is primarily one of design and operation. One of the most common misconceptions in this period of growth is that the TCO of a new data center is lower with a low density design. We look at the construction and design of both types and present results demonstrating that high-density DCs are a better choice for reducing the owners cost. These results apply to new construction and mostly-unconstrained retrofits. Densities of 1000 watts per square foot of work cell are being achieved with good efficiencies. Modern designs of 200 to 400 watts per square foot of work cell are much more common, but cost more. Costs of the architectural space, power systems and cooling systems are reviewed as are the operational costs for these systems. High-density DCs do cost less. The challenges for the high-density DC are also called out and suggestions for successful operation are made.

Motivation

ASHRAE (2005) provides projections for datacom density trends, as shown in Figure 1. Of particular interest in this paper is the trend for compute servers, both 1U & blades, and 2U; these are the primary building blocks of scale-out data centers. The 1U trend for 2006 indicates a heat load of roughly 4000 watts / sq ft of equipment floor space. A typical rack has a foot print of 39 by 24 inches, this represents a 26 kW rack. Very few racks of this power are in place. Is Figure 1 incorrect or are there other factors? The ASHRAE guide represents the peak value, or what could be expected in a fully populated rack. But DCs are not being built to this density. Instead DCs are still being built to the 1999 ASHRAE values of compute density. Is DC development lagging behind and not technically capable of supporting 26 kW racks? Or would a data center at that density be too costly, and more expensive than the ones currently being built? These issues are analyzed and it is shown that the technology for power and cooling for racks per the ASHRAE trend does exist, and that a data center built to this standard would have a lower TCO. The authors believe that the problem has to do with the life of datacom equipment (3~5 years) as compared with the lifetime of these facilities (~15 years) and the inertia that lifetime builds into data center strategies and design.

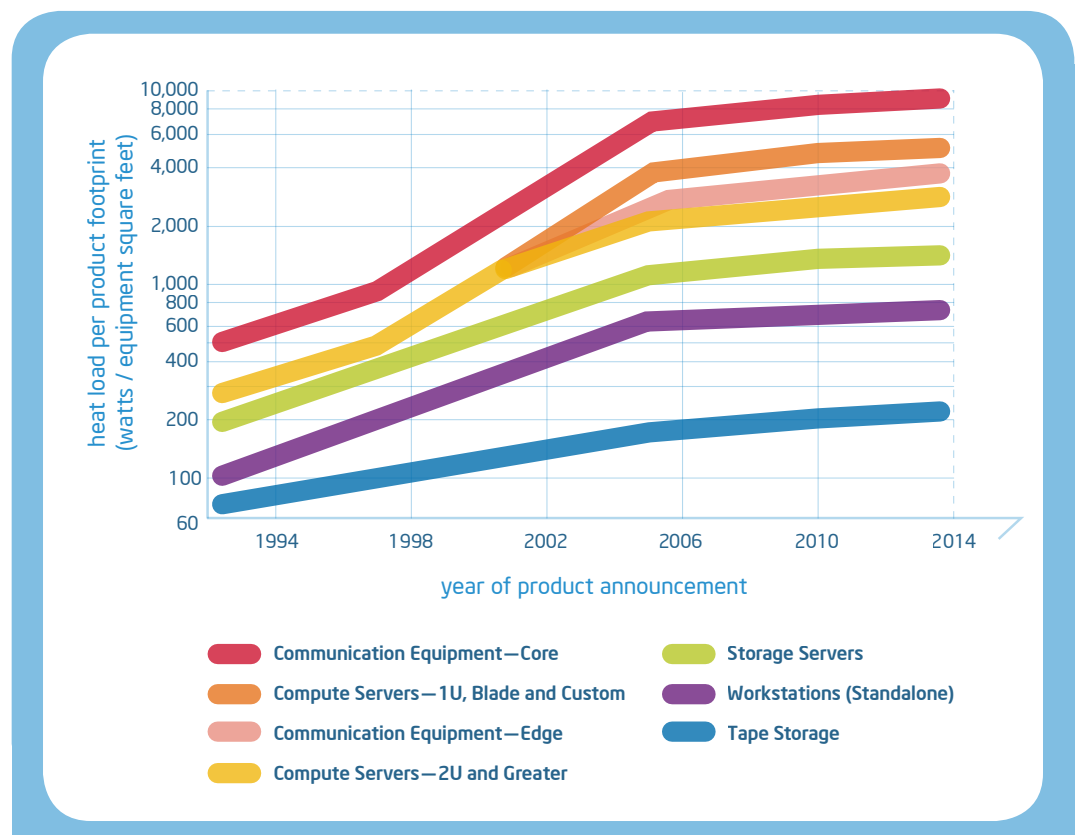


Figure 1. ASHRAE Datacom Trend Chart showing increasing density over time

Background and Definitions

Best practices for today's data center layout is repeating rows of racks side-by-side with alternating cold aisles and hot aisles. The cold aisle supplies cool air to the servers, with each rack discharging into a hot aisle shared with the next row of servers. Raised floors provide cool supply air to the cold aisles with overhead returns to the air conditioning system for the warm return air. In this hot aisle / cold aisle configuration, varying numbers of servers can be fit into each rack based on many factors; cooling capability, power availability, network availability, and floor loading capability (the rack's loaded weight). Other configurations can also be successful (ASHRAE 2006).

Definitions

Prior to evaluating the benefit or drawbacks of various metrics there are definitions that need to be presented. The first is that of the **work cell** (see Figure 2).

The work cell is the repeating unit of cold aisle, rack, and hot aisle. This represents the square footage directly attributable to a specific rack of servers.

High-density data center is taken to mean a data center with racks at 14 kW and above, with a work cell of nominally 16 to 20 square feet. Further, high density does not imply or require liquid cooling. It has been reported in some trade journals and elsewhere that anything above 14 kW will need supplemental cooling or liquid cooling to be able to handle these types of loads. This is not the case. High density can be cooled successfully with standard hot-aisle / cold aisle design as shown in Figure 3 (16 sq ft work cells and 14 kW to 22 kW racks with 40+ servers in each).

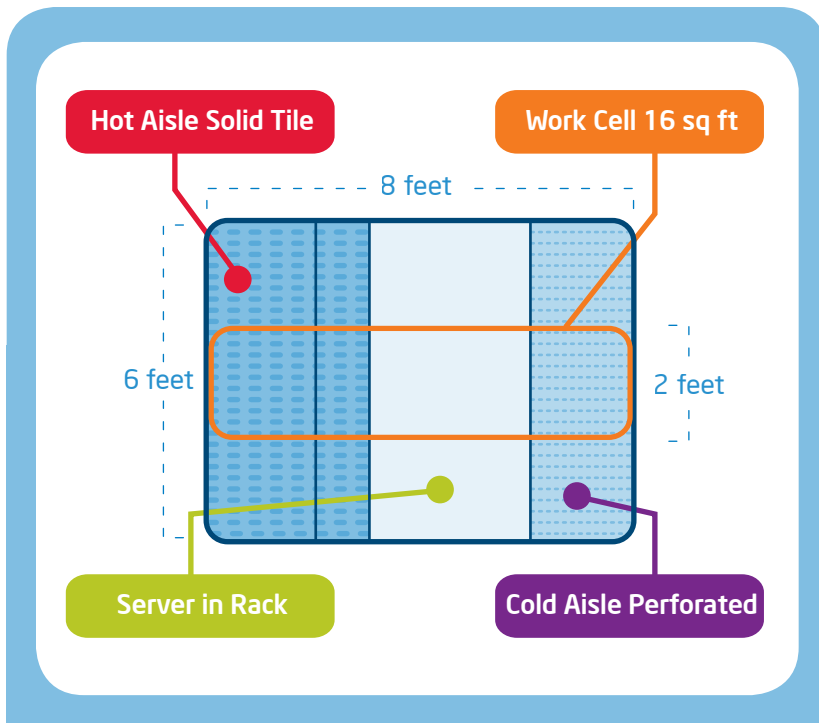


Figure 2. Single work cell in a row of servers in a Data Center, looking down from above




Figure 3. Photograph of a cold aisle in a high-density Data Center




Figure 4. Infrared photo of 14 kW racks at the end of a cold aisle in a high-density data center

The limit of air-cooling in high-density data centers is a matter of much debate. Figure 4 shows an infrared image of the data center in Figure 3. This HPC data center is easily carrying the high-density racks with no recirculation problems. On-going internal analysis by the authors show that supporting 30 kW racks is feasible with air cooling and the results of this paper applicable there.

Layout Efficiency measures data center square footage utilization. This is defined as racks per thousand square feet. This value is a measure of the efficiency of the DC layout. The typical range is 20 to 30, higher numbers being better. This is similar to the Rack Penetration Factor of Malone and Belady (2006); however, this metric is independent of the size of the rack itself.

The metric most often used when discussing DC density is **watts/sq ft**. However this metric is often misused because it lacks specificity in the denominator. It is typically assumed that the value refers to raised floor area, but it could refer to the DC and support area, or even the entire campus.

Watts/sq ft of work cell is the preferred metric for data center to data center benchmarking, or in infrastructure discussions. The metric is particularly well suited in the evaluation of cooling load as the ability to move the requisite cooling air through the raised floor and the exhaust air through the hot aisle is all carried in the size of the work cell. A high-powered rack surrounded with a large amount of raised floor in the work cell is not high density.

Watts/rack is useful in sizing power distribution to a rack and for determining how full a rack can be. It should not be used by itself to define a data center unless the square footage of the work cell is known, then the two metrics become functionally equivalent.

Watts/sq ft of raised floor is also of limited value. The layout efficiency of the different data centers can vary and this can greatly affect the value. One use for this metric is for infrastructure sizing (e.g., watts/sq ft times the total raised floor can give the total needed cooling). But any method of calculating utility system sizing must consider average watts/sq ft or watts/rack. Sizing infrastructure using max power from each server will result in oversized facilities equipment.

Total Cost of Ownership (TCO) represents the cost to the owner to build, as well as the cost over time to operate and maintain the data center. These costs are all brought back to a present value using appropriate engineering economics. The right TCO metric is **cost/server** when the specifics and number of servers has been determined. Alternately, **cost/kW** can be a useful metric, particularly when the servers to be installed are not known. In this metric, kW is the power available to the servers, rather than the power into the site (which includes UPS losses and power for the cooling system). The use of cost/sq ft would not be valid as the high-density DC will have a greater cost/sq ft. Unfortunately, low density DCs are often chosen based on this faulty comparison. Cost/sq ft is not valid, as compute capability for each square foot is not the same for both DC types.

Two analyses were done in determining the TCO of low and high-density data centers. First an example data center was considered with implications of each density option. The requirements of each were compared and TCO impacts calculated. Second, a benchmarking effort's results are compiled and plotted, and compared with the results of the example data center analysis. Finally, specific design and operational considerations for high-density DCs are reviewed.

Example Data Center

Consider a new DC where the owner has determined 10,000 1U dual processor servers are needed. This DC and the server choices will serve to demonstrate the difference in TCO of different density options. (1U = 1.75 inch height, equipment racks are measured in U, with a typical rack being able to hold 42 1U servers, or 21 2U servers, etc.) The supplier of the equipment has adhered to the ASHRAE Thermal Guideline (ASHRAE, 2004) and published the cooling load and required airflow of the server. Table 1 shows a summary of the server and the two data centers.

	Low-Density Data Center	High-Density Data Center
# of servers	10,000	10,000
Watts / server	400	400
CFM / server	39	39
kW/rack	6.6	17
Servers / rack	16	42
Total racks	625	238
Sq ft / work cell	16	16
Layout Efficiency (rack/Ksf)	~22	~22
Sq ft of raised floor needed	28,571	10,880

Table 1. Data Center and Server Definition

Low density will be taken as the median of a survey done of 28 new or retrofit data center projects being considered by a leading DC design-build firm. The data (Aaron 2006) is shown in Figure 5. The values range from 75 watts/sq ft to 250 watts/sq ft. The median is between 125 and 150 watts/sq ft, which represents a 6.6 kW

rack. For the high-density DC the main goal is full racks and minimizing square feet. 42 of the selected servers in a rack would require just under 17 kW.

The data center can now be further detailed. The total airflow required is the 39 CFM for each server plus a 20% safety factor for leakage and bypass. Note that the total airflow in each DC is the same; it is driven by the number of servers and not the space density. The value of 20% is low and would only be successful in a data center with good airflow management. A value of 35% may be required in a DC where less care is taken in design, CFD analysis, and operational acumen.

The raised floor in the low-density room is 18 inches, while the high-density room will need a 30-inch raised floor to handle the higher per-rack flow rate for the high-density cabinets. Both will require 4 MW of UPS power to drive the servers. Total server power is independent of density, as long as the assumption of equal work outputs for each room is adhered to.

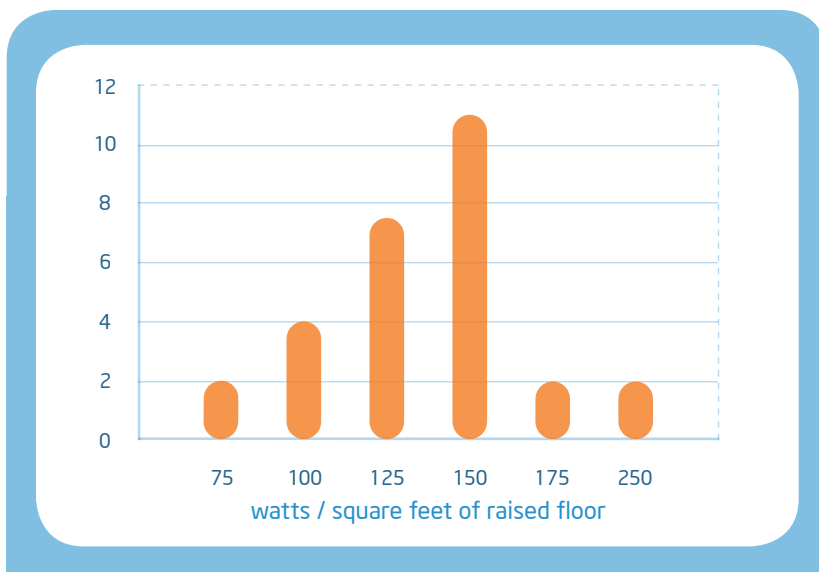


Figure 5 New data center projects reported by national data center design-build firm

	Low-Density Data Center	High-Density Data Center
Total Airflow CFM	468,000	468,000
Raised floor height	18 inches	30 inches
CFM / rack	749	1966
Total UPS power needed	4 MW	4 MW
Cost of power	10¢/kW-hr	10¢/kW-hr

Table 2. Data Center Design Results

The data center has essentially been scoped out. Costs associated with the various segments of the data center can be determined. There are five major areas for consideration; civil, structural, and architectural (CSA), power, mechanical (primarily cooling), safety and security, and finally the IT equipment itself. The density will impact each of these in a different way.

CSA is highly impacted by the density. Turner and Seader (2005) provide a useful two component cost model. First it includes a cost/sq ft value which is independent of data center density or “Tier” (or critical-ness). The second component is a value based on cost/kw of useable UPS power. This cost does vary by Tier. The cost/sq ft metric is \$220/sq ft and is primarily associated with the CSA portion of the facility. This is shown in Table 3. A ~\$4M savings on the CSA portion of the facility can be had by building a high-density data center with fewer square feet.

The high-density data center requires nearly 1/2 acre less land. Also the permitting and fees for any project are often based on square footage. The specifics of the individual site location would dictate the magnitude of additional savings associated with higher density.

There are additional costs incurred by the high-density design. First is the cost of a CFD analysis. It could be argued that any data center with 10,000 servers warrants a CFD analysis. However, low density designs often go without. The cost of this analysis is difficult to estimate as it depends more on complexity than on the square footage, but for a new, homogenous DC \$5/sq ft is fair.

The other high-density penalty is that of the higher building height required. The raised floor is higher

and this will also need a similar height increase in the return air plenum. The overall building height will be roughly 30" greater. Building cost is far less sensitive to height than to area. The marginal cost to increase the height of the building as required is assumed to be 10% of the building cost. Experience in high-density data centers is that the cost delta for the 30" raised floor is small compared to the 18" raised floor. Even for the most challenging seismic area the design is the same with the exception of the longer pedestal. This cost delta is on the order of \$1 / sq ft.

The required power total to the room and the airflow total in the room are the same for both concepts. This can be extended to the major facilities equipment (Cook 2006). Both the high and low density options will use the same chillers, cooling towers, chilled water pumps, utility connection, UPS, and major switchgear. It follows that the utility spaces for both are identical. In the DC fewer, larger, higher-capacity fan-coil units will have a lower cost per CFM than many small distributed units but the cost determination is beyond the scope of this work. Similarly, power distribution units capital cost favor the high-density space, as does the allowed shorter cable runs.

The electrical operational costs are approximately equal. There will be a difference in power used, with less power drawn in the high-density room due to shorter cable runs and larger (and typically more efficient) equipment, but the value will not be significant.

Cooling costs will have an impact on the TCO. The central cooling plant will be the same for both. The major difference in cooling costs comes from the power needed to move the air

in the data center. Both have the same total flow rate. The configurations were modeled using a commercially available CFD program. In the high-density room, with the 30" raised floor and >50% open grates, a static pressure of 0.10" wg. is required to drive the proper airflow through each grate. In the low density design, with the 18" raised floor and the 25% open perforated tiles, 0.13" wg. is needed. The additional power needed for the low-density DC can be determined from the fan laws. Increasing the pressure required by 30% for the same airflow will have an impact of

$$\Delta power_{airflow} = (\sqrt{1.3})^3$$

The power for the high-density case is from kW/CFM values for an actual high-density DC.

The solution to this cost penalty would seem to be using the more open grates instead of the restrictive perforated tile. Unfortunately that solution would not work. Grates, with their non-restrictive flow/pressure characteristic, need a deep plenum under the raised floor with minimal flow restrictions to ensure an even flow distribution. Grates used with a lower, more restrictive plenum would not provide uniform airflow and are not a good airflow management tool for shallow raised floors. Perforated tiles work best in that application because their higher pressure drop is the controlling factor in the flow distribution and results in uniform flow. (VanGilder and Schmidt, 2005)

Safety and security cost is somewhat sensitive to area (e.g., smoke sensors / sq ft) but it is not the primary driver of these systems so no specific credit is taken for high density. Also, a desired higher level of monitoring due to the higher density could offset this.

Lighting will cost more in the low density space. There are ~18,000 more square feet that need to be lit at roughly 1.6 watts / sq ft. The increased maintenance would also add to the cost.

The IT equipment is a large portion of the total budget. Both options hold the same number of servers. The difference is the number of racks. The partially full racks carry an economic penalty. A per-rack cost of \$1500 is assumed, with another \$1500 to move in and install the rack. This value is conservative; with higher values for some of the more advanced, interconnected DCs.

The summary in Table 3 captures the major differences in costs for the two subject DCs. It does not represent the total cost, but covers the areas where there are differences. The cost for a high-density DC is measurably less. In the example above savings were \$5.2 million dollars. This savings is equal to \$520/server, or 1700 more servers (at ~\$3K/ea), or a year's electricity cost.

	Low-Density	High-Density	Notes
Capital Cost - Building	\$6,285,620	\$2,393,600	\$220/sq ft for CSA
Design Cost for CFD	\$0	\$54,440	Assumes \$5/sq ft;
Capital cost taller DC	\$0	\$239,360	Assumes +10%
Capita cost for 30" RF.	\$0	\$10,880	\$1 sq ft
Lighting	\$126,000	\$0	NPV(5 yr, i=5%)
IT Equipment (Racks)	\$1,875,000	\$714,000	\$1.5K/ea + \$1.5K/install
Oper Cost - Cooling	\$1,091,000	\$736,000	NPV of 5 yr with i=5%
Total Cost Delta	\$9,377,620	\$4,148,240	\$5.2 M savings

Table 3. Data center TCO comparisons and major cost deltas

Benchmark Data

The example data center analysis showed that a high-density data center does have a lower TCO than a low density DC. This result is also supported by a review of Figure 6.

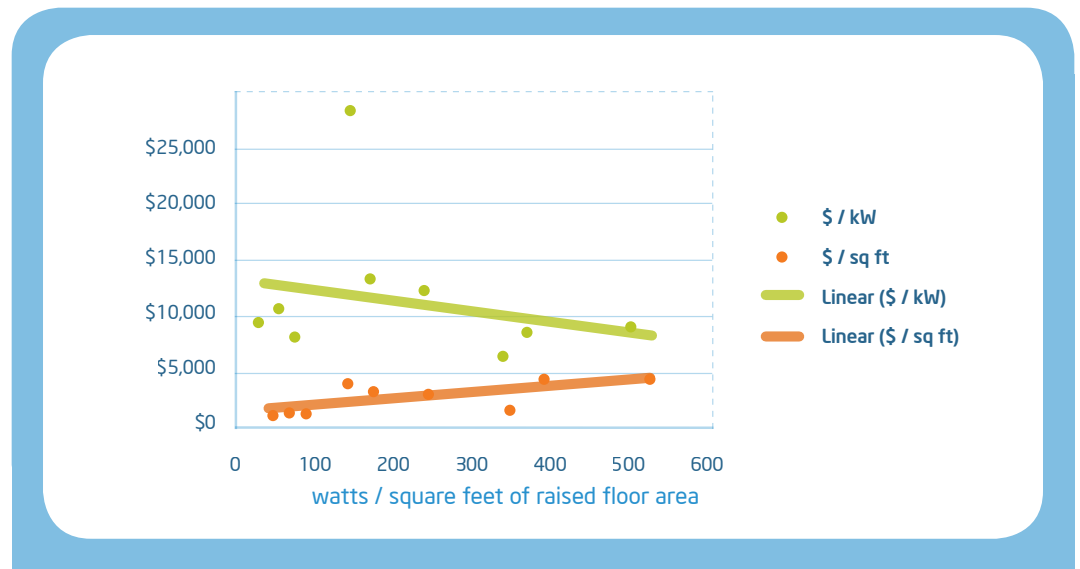


Figure 6. Data center cost metrics as a function of data center density

This shows the result of a benchmarking study of eight different internal and external DCs completed by Costello (2005). As expected, cost/sq ft is higher at greater densities. But for the same computing capability less square footage is needed, making the cost per square foot metric a poor predictor for TCO. The better metric is cost/kW. The kW basis accounts for the DC's ability to support a given compute workload. The datacom equipment will need the same total power regardless of density. Cost/kW shows a negative slope indicating high density does have a lower first cost. Note that the cost here is based on the total data center, including design, land, and all infrastructure and is different from the metric used by Turner and Seader (2005).

High Density Considerations

High-density DCs require different methods in both design and operations on the part of the owner, but when these are weighed against a lower TCO they are usually a good investment.

Design and Construction: The design of the high-density data center requires a greater focus on airflow in the room. The challenge is not the increased volume per rack but a much greater control of airflow distribution. The air must be delivered to where it is needed. Low-density rooms often exist with poor airflow management. CFD analysis of the high-density room is a must. Patel (2002) reviews the importance and opportunities provided by this level of analysis.

High-density rooms often have higher cooling loads than the typical CRAC units prevalent in low-density data centers can provide. The CRAC units, specifically designed for data centers, provide benefits such as monitoring, reliability & redundancy, and capacity tailored to high sensible heat ratio loads found in data centers. Denser applications often require industrial grade units based on size alone. The goal is incorporation of the CRAC unit's benefits into the larger system.

Another challenge in the design of high-density spaces is uniformity of the servers. Uniform server loading makes the task of airflow management simpler. Non-uniform loading can still be handled but the details have to be understood. Zoning the DC into homogenous server types can facilitate airflow management.

Operations: A frequent (but incorrect) concern voiced over a high-density data center is that the hot aisle will be too hot for personnel. The hot aisle temperature is independent of density in the ranges being discussed here. Assume that in each DC the inlet air to the servers is within specification and proper designs have precluded recirculation. The servers, whether there are 42 or 16 in the rack, will pull the needed amount of air per server from the cold aisle and discharge it to the hot aisle. The temperature rise across any server is based on the individual workload and thermal control algorithm in place. But that delta T, assuming the servers are the same, will be the same. The hot-aisle temperature, which is a direct result of server temperature rise, is independent of density.

What is more likely the cause of the cooler hot-aisle phenomena in low density DCs is a greater chance of air-flow mismanagement with leakage or bypassed cool air being delivered to the hot

aisle. In a properly designed, built, and maintained DC, regardless of density, the hot aisle will be hot. If the hot aisle is not hot, cooling capacity, energy, and money are being wasted.

Another issue is data center airflow velocities. Consider the 42-server rack discussed earlier. With a 20% safety factor, the rack itself will need approximately 2000 CFM. That flow, when delivered through a 2x2 floor grate will have a nominal velocity of 500 fpm. Tate Access Floors (2002) provides flow rate-versus-static pressure curves of a typical floor grate. 500 fpm is in the middle of the operating range of a typical grate (~56% open area), so the pressure and flow are not extreme.


500 fpm (5.7 mph) is above what would normally be considered a comfort space, particularly at data center supply-air temperatures, however it is not unworkable. The Beaufort Scale (National Weather Service, 2006) defines this velocity as a light breeze, and not until the velocity reaches a moderate breeze (1144 – 1584 fpm) is the velocity noted as “raises dust and loose paper.”

Conclusions

High-density data centers will provide the DC owner with a reduced cost of ownership when compared with that of a low density DC.

High-density data centers require specific design considerations, most notably a path for the higher volume of air. Grates can replace perforated tiles. Raised floors of 30 inches are needed in the cold-aisle/hot-aisle strategy. This could preclude some legacy data centers from moving to high density without local enhanced cooling, but new data centers and those with sufficient height for a retrofit can benefit from increasing densities.

There are risks and ergonomic negatives to a high-density configuration, but these can be overcome by proper design and recognition that modern data centers do not require continuous staffing. If the DC can be designed or retrofit to support the infrastructure for high-density computing, the owner will be able to have a smaller DC with the same computing performance at a lower TCO.



For further information, please visit:

www.intel.com/technology/eep

Acknowledgements

We would like to express our gratitude to Kelly Aaron of Nova for her support of this study.

References

- Aaron K., e-mail message to author, Mar. 6, 2006.
- American Society of Heating, Refrigerating and Air Conditioning Engineers (ASHRAE) 2004. Thermal Guidelines for Data Processing Environments. Atlanta: ASHRAE
- ASHRAE. 2005. Datacom Equipment Power Trends and Cooling Applications. Atlanta: ASHRAE
- ASHRAE. 2006. Design Consideration for Datacom Equipment Centers. Atlanta: ASHRAE
- Cook, D. 2006. DC TCO Report, Hillsboro, OR: Intel Internal Report
- Costello, D. et. al. 2005. Data Center Benchmarking, Hillsboro, OR: Intel Internal Report
- Malone C. and Belady C. 2006. Data Center Power Projections to 2014. iTHERM 2006, San Diego
- National Weather Service. 2006. <http://www.srh.noaa.gov/mfl/hazards/info/beaufort.php>
- Patel, C.D., Sharma, R, Bash, C.E., Beitelmal, A. 2002. Thermal Considerations in Cooling Large Scale High Compute Density Data Centers, 2002 Inter Society Conference on Thermal Phenomena, pg 767-776
- Tate Access Floors, GrateAire 24 specification sheet, http://www.tateaccessfloors.com/pdf/grateaire_panel.pdf, Tate Access Floors Inc, Jessup, MD (accessed July 15, 2006)
- Turner, W.P and Seader, J.H. 2006. Dollars per kW plus Dollars per Square Foot are a Better Data Center Cost Model than Dollars per Square Foot Alone, Uptime Institute White Paper, Santa Fe
- VanGilder, J.W. and Schmidt, R.R., 2005, Airflow Uniformity through perforated tiles in a raised-floor Data Center. ASME Interpack 05, San Francisco, 2005

Copyright © 2005 Intel Corporation. All rights reserved. Intel, the Intel logo, Intel. Leap ahead., and the Intel. Leap ahead. logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2005 Booz Allen Hamilton 0205/CEG/ESP/XX/PDF Part Number: 306623-001 EN