Guidelines for

Assessing Power and Cooling Requirements in the Data Center

As computing deployments continue to increase in density, cooling dynamics have become a paramount consideration in many data center environments. This article discusses steps that administrators can take to become familiar with intensifying power and cooling requirements, particularly for rack-dense blade servers—and explores tactics to help optimize the deployment of computing components throughout the data center.

BY DAVID MOSS

Related Categories:

Blade servers

Data center density

Data center technology

Planning

Services

Visit www.dell.com/powersolutions for the complete category index.

The advent of blade server technology per se did not create a new challenge in data center power and cooling requirements, but it did exacerbate existing conditions. Administrators no longer can rely simply on HVAC (heating, ventilation, and air-conditioning) systems to cool data centers. They must move past the mindset of cooling the room into a mindset of cooling each rack.

Servers generate heat—in fact, 100 percent of the input power is released as heat, usually measured in kilowatts (kW). Most servers are cooled with internally generated airflow, commonly expressed as cubic feet per minute (CFM). In general, the more heat a server produces, the more airflow it is likely to consume to maintain the temperature requirements for its internal components. Unless data center administrators are satisfied with racks operating at a 2 to 3 kW load, they should pay special attention to airflow consumption within each rack. For example, in the late 1990s a typical 4U, two-processor server consumed about 40 to 50 CFM to cool a maximum load of approximately 400 watts (W)—or a nominal load of around 200 W. A full 42U rack of these servers would typically operate on 2 to 4 kW of power and require 400 to 500 CFM for cooling. In contrast, eighth-generation Dell™ PowerEdge™ 1855 blade servers consume approximately 2,000 CFM to cool 15 to 25 kW.

Moreover, as computing density increases in the data center, a corresponding increase in equipment exhaust temperatures occurs. Inside each rack, an inadequate supply of chilled air may result in the consumption of used air; some equipment may consume air from its own exhaust traveling over the top of the rack (see Figure 1). Exhaust temperatures can severely limit deployment potential if the warm air is allowed to recirculate to the intake of any servers. This is a fundamental challenge in

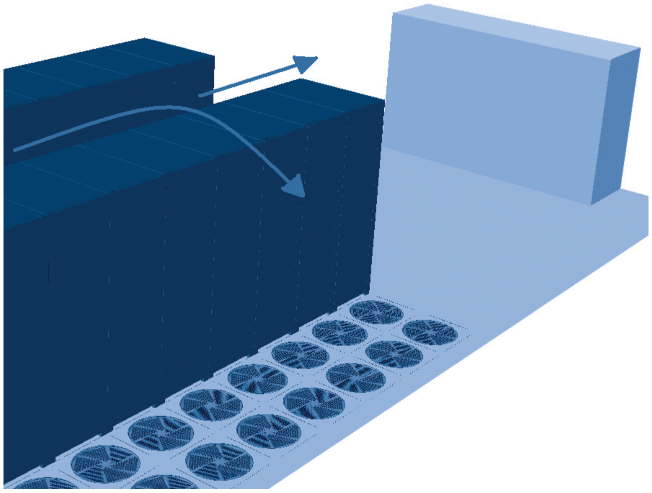

Figure 1. Equipment exhaust recirculated through the rack

today's data centers. The server intake aisle must fill up with chilled air about as quickly as the servers consume the air. Otherwise, hot exhaust air may recirculate through the equipment and result in overheating.

Because of the open nature of a raised-floor environment, administrators cannot be certain that the air they supply directly in front of each server actually gets to that server. Administrators must understand and plan around the actual airflow dynamics of the data center. The airflow typically includes areas of weakness, which can result in hot spots, as well as areas of strength, which can be cool enough to support more equipment. The challenge is to understand and take advantage of the strengths and weaknesses in data center locations or to intelligently manipulate those strengths and weaknesses to normalize the airflow.

Following standard best practices

The first step for data center administrators is to understand basic best practices. Wherever possible, administrators should ensure strict adherence to the *hot aisle/cold aisle* approach of alternating intake aisles and exhaust aisles—that is, rack intakes should face each other and rack exhausts should face other. This approach accomplishes two important goals: It helps keep the exhaust from one row feeding the next, and it provides for a concentration of chilled air at the server's intake. Administrators should avoid introducing chilled air to the exhaust aisle—both intentional introduction through hot-aisle vent tiles and seepage from floor cable openings. Even a small cable opening can equal or exceed the amount of air that flows through a cold-aisle vent tile. Cable grommets can be used to help block the intrusion of chilled air into the hot aisle, as shown in Figure 2.

Although it may be tempting to cool a hot aisle so the temperature is more comfortable for administrators, the exhaust aisle should remain hot to maximize the HVAC return temperature, which helps the HVAC to operate as efficiently as possible. Gaps between equipment should be avoided—including gaps between racks and gaps within racks. Blanking panels can be used within racks. Such basic best practices lead to a common goal: creating the coldest possible intake aisle with the highest possible volume of chilled-air delivery and the warmest possible exhaust aisle.

Estimating airflow rates and cooling capacity

To move beyond the basics, administrators should develop an understanding of data center dynamics, including airflow delivery capability as well as the equipment's airflow consumption rate. The delivery capability is simple—administrators simply measure tile-flow rates. The facilities department within an organization may have a flow-hood, such as the one shown in Figure 3, to help take such measurements. In addition, consulting services such as Dell™ Data Center Environment Assessment (DCEA) are available not only to measure airflow rates but also to analyze that data and make specific recommendations.

The first step for data center administrators is to understand basic best practices.

The other side of the equation is consumption, which may be trickier to calculate because hardware vendors do not always provide airflow consumption rates. Dell provides consumption rates for most of its rack-mount equipment through the Dell Product Configuration Calculator, which is available at www.dell.com/calc.

Dell servers use variable-speed fans controlled by algorithms that use ambient and component temperature sensors. Airflow rate has a high dependence on inlet ambient temperature. Cooler inlet temperatures equate to lower airflow rates. In the absence of measured values, administrators can estimate equipment airflow rate by

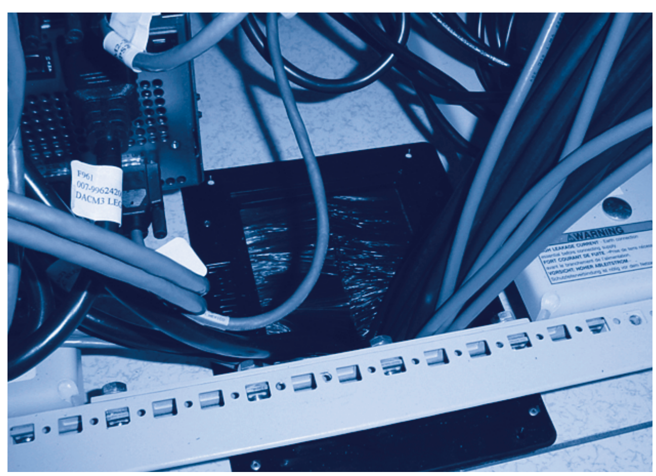

Figure 2. Cable grommets to minimize loss of chilled air

using the temperature difference between the intake and exhaust as well as the equipment's power draw. The airflow rate can be approximated using one of the following formulas:

$$\text{CFM} = 1.78 \times \frac{\text{equipment power (W)}}{\text{temperature difference (°C)}}$$

$$\text{CFM} = 3.2 \times \frac{\text{equipment power (W)}}{\text{temperature difference (°F)}}$$

If the equipment power is known but the temperature difference is not, administrators can approximate the airflow rate at 9 CFM per 100 W, which assumes a temperature difference of about 20°C or 36°F, depending on the formula used. Administrators should be careful when estimating airflow rates with either method. Because the airflow rate is directly proportional to the equipment power, administrators can easily overestimate airflow rates if using conservative power numbers. Best practices recommend using measured values over estimates for equipment power whenever possible.

This information can then be used to calculate cooling capacity. At minimum, airflow rate balances should be studied. Administrators should compare cumulative rack consumption rates versus cumulative delivery rates, which are best compared on a row-by-row basis. Also, the total airflow rate supplied to an aisle should be examined in comparison to the consumption of the aisle's two flanking equipment rows. This process should then be applied to the rest of the data center. For example, if vent tiles are allowing some areas to be overprovisioned with cooling in comparison to consumption, administrators may ameliorate the situation simply by moving vent tiles away from areas that are overcooled to areas where more cooling is needed.

Administrators can determine the effect on an area in which the chilled-air delivery is reduced by studying the exhaust temperatures and top intake temperatures in that area. If the intake of the highest mounted systems is still significantly below equipment specifications—which is generally around 35°C (95°F) but may be de-rated at higher altitudes—then administrators may be able to move some of the cooling from that area. If the exhaust temperatures in the aisles flanking that cold aisle are high and approach the equipment operating (intake) temperature limits, then the removal of the chilled-air supply in that area could result in a deficiency and subsequent recirculation of exhaust into the intake aisle. A simple test of blocking off the airflow from a tile or two may help administrators understand whether chilled air can be reprovisioned from one area to another. If no significant change in the inlet temperatures of several surrounding top-mounted servers ensues, the air could be reprovisioned to an area of greater need. *Note:* The methods described in this section involve a substantial amount of trial and error, which can affect mission-critical equipment. A more scientific approach is explained in the next section.

Predicting data center cooling needs with CFD

Numerical methods for computer modeling have been used effectively to predict the behavior of many types of dynamic systems. Whether for the prediction of automobile aerodynamics, the strength of a building structure, or the weather patterns for a five-day forecast, computer modeling helps solve a variety of engineering problems, and the data center is an excellent candidate for this approach. *Computational fluid dynamics* (CFD) is the term generally applied to the analytical modeling of fluid systems or air systems, such as in the case of a data center. Several software products can help administrators create a thermal/airflow model of their data centers to perform virtual trial-and-error scenarios. *Note:* Although this type of analysis is best performed by engineers with the proper expertise, data center administrators can also learn to perform this type of environmental study effectively.

Organizations also have the option of hiring CFD consultants, including those from Dell. Dell plans to offer this type of analysis as an option of the Dell DCEA service. Through the DCEA service, an organization can have its data center measured and analyzed, including a CFD

Figure 3. Flow-hood to measure the airflow rate of an entire tile

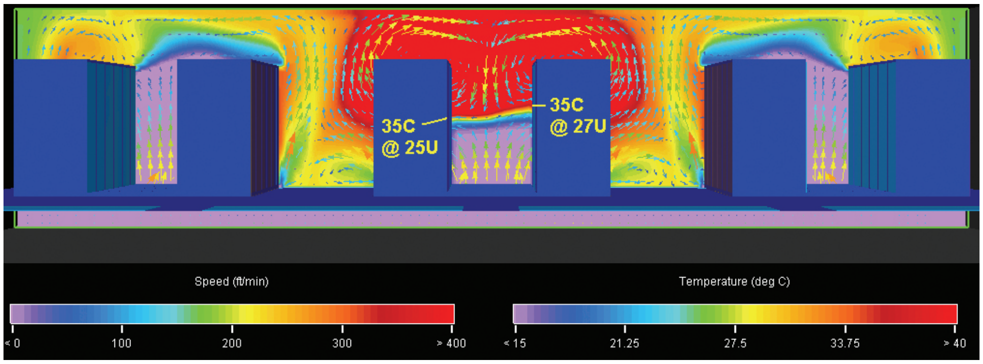

Figure 4. Temperature and air vector plot for example CFD analysis

analysis, with a resulting confidence level that the infrastructure will support future needs.

To understand how a CFD analysis can help, administrators can consider the example CFD model depicted in Figure 4, which was created using Flotherm, a leading thermal analysis application from Flomerics. This cross-section through six rows of racks shows airflow patterns superimposed onto a temperature profile. The rows are arranged according to the hot aisle/cold aisle approach. In this particular example, the CFD analysis proposed that equipment with a much higher air consumption and heat load be deployed in the two middle racks. Airflow is distributed uniformly to each of the three cold aisles. However, the cool temperature (depicted in lavender in Figure 4) completely fills the two outside cold aisles but does not fill the middle aisle. In fact, the cool temperature billows out of these two outside aisles. In the middle aisle, the cool air travels only about halfway up the rack before it is depleted by system consumption. The air pattern shows hot air being re-circulated over the top of the rack at temperatures that exceed equipment specifications.

The plot in Figure 4 is set up to represent in red any air temperatures greater than 40°C (104°F), which is five degrees higher than most equipment's maximum inlet temperatures. The analysis ultimately shows that, without changes, this middle aisle cannot handle the proposed deployment of new equipment because equipment deployment above the 25U rack location will receive unacceptably high inlet temperatures. However, the overabundance of cold air in the outer two aisles indicates that a significant amount of air could be diverted from these aisles to the middle aisle to help satisfy its cooling needs. Using CFD analysis, the administrator in this example scenario could model various configurations to help determine the best way to adequately divert the excess cold air—without incurring the risk of physically reconfiguring the actual data center equipment for numerous trial-and-error comparisons. In this particular example, subsequent numeric iterations suggested removal of 20 percent of the vent tiles from the outer cold aisles and the insertion of several grates in the center cold aisle. The CFD analysis indicated that grates, at 60 percent open, could offer substantially more airflow than the standard 25 percent open, perforated tiles that had previously been in use.

By taking advantage of CFD analysis, Dell has helped many organizations optimize their data center cooling capabilities—often without the need for additional HVAC infrastructure. While such optimizations can be determined through the trial-and-error approach, using CFD analysis software as a predictive, iterative tool enables data center administrators to find a quick, efficient path to resolution.

Once the delivery of chilled air has been optimized, data center administrators who need the capability to increase the deployment should consider alternate sources of chilled air. For example, supplemental cooling systems can offer chilled-air generation at the rack level, and refrigerant-based systems are designed to blow air down into the cold aisle to supplement the raised-floor delivery. Other types of systems can fit in line with racks to deliver supplemental or primary chilled air. In addition, some self-contained racks can generate chilled air inside the rack for the enclosed equipment.

Balancing power and cooling needs in the data center

In today's world of high-density data center equipment and chilled-air delivery challenges, administrators must understand the supply limitations of the data center and the demand of the equipment. In short, power and cooling needs are a supply and demand problem. By learning how to properly assess data center power and cooling requirements and procuring the appropriate tools to help resolve such environmental considerations, administrators can be enabled to design the optimal enterprise computing environment—and to deploy rack-dense servers with confidence.

David Moss is a senior thermal/mechanical architect with Dell and has over 20 years of experience in electronics packaging design. In his role at Dell, he holds 18 U.S. patents. David has a B.S. in Mechanical Engineering from Texas Tech University.

FOR MORE INFORMATION

Dell Product Configuration Calculator:
www.dell.com/calc