PROCESSOR

Cover Focus Articles

Click To Print

General Information
November 6, 2009 • Vol.31 Issue 27
Page(s) 1 in print issue

Power & Cooling Capacity Planning

Introspective Strategy Drives The Process Of Gauging Power & Cooling Levels

Pull the wraps off the best-laid plans for any modern data center and you'll discover an intricate roadmap for power and cooling that forms the backbone of that space. Because these physical infrastructure elements are critical to the overall success of the data center, it's essential to determine the correct amount of power and cooling for both present and future needs.

Due to the wide range of technologies now populating data centers, that task can be monumental in scope. Whether building new, renovating, or creating an addition, data center managers not only must gauge the power and cooling required to provide basic support for their equipment, but they are also being increasingly pushed to consider energy-efficient options that can help save their organization money in both the short term and long run. Balancing these factors to find the perfect outlay of power and cooling requires a targeted strategy.

Key Points

• Determining specific power and cooling needs depends not only on the data center's hardware but also on the utilization levels of those devices.

• Power consumption levels provided in manufacturer specifications can refer to maximum levels, so keep this in mind when gauging the amount of overall power required for your space.

• The most critical step for ensuring future power and cooling needs will be covered is power capacity, because unplanned cooling needs often can be addressed on an individual level.

Predict Your Power

A primary issue to consider is the level of reliability you need from your data center and the layers of redundancy needed to achieve that reliability, says Matt Kightlinger, director of solutions marketing for Emerson Network Power's Liebert Products (www.liebert.com). Another consideration is high-density equipment, whether at the row level or at the rack level, because most of today's data centers are going to have a high-density zone.

Arriving at a determination for specific power and cooling also starts with the application of servers and other hardware using those elements. For example, if an organization deploys server virtualization or cloud computing, the servers could be operating at a higher load level and therefore approaching wattage use that's closer to the maximum usage, says Jim Scherr, director of sales and operations at PDUs Direct (www.pdusdirect.com).

"Equipment manufacturers provide estimated power consumption in the specs of their products," Scherr says. "These vary widely depending on the types of products and application. The watts described in the specs are usually top-end, so in normal usage—dependent on application—they can run significantly less."

Like the specifications, manufacturer-provided programs designed to determine power and cooling requirements also tend to deliver estimations on the high end of the scale. Dennis Julian, a principal with Integrated Design Group (www.idgroupae.com), says that when considering the electrical service load to the space, you can assume that the percentage of actual load is lower than when considering the power to an individual rack.

“Any rack can operate at full load, but the likelihood of all racks operating at maximum load is less probable due to the diversity of equipment types and loads of that equipment,” Julian says. “If the design is for a very homogenous environment—such as research computing or bulk processing—then the average rack load will approach the maximum work load. For a new buildout, a maximum and average rack load should be determined based on the expected equipment to be installed.”

In turn, that maximum rack load will not only help to determine the typical circuit capacity for each rack but also the type of cooling system required. According to Scherr, the standard for cooling traditionally was one watt of cooling for every watt of equipment power, but that has changed as the heat differential between the front and back of racks has significantly increased. Today, hot/cold aisle designs are generally favored over entire-room cooling, but Scherr says even that approach can run into problems when trying to cool highly dense racks. In such cases, supplemental cooling might be needed to address hot spots that materialize on the data center floor.

Plan Ahead

Figuring the correct amount of power and cooling for your current infrastructure is only one piece of the puzzle, because inevitable equipment variations over time require that your data center can roll with any changes. In particular, power is the most important element when it comes to ensuring that a data center can accommodate future changes.

“There are many ways to populate a data center with high- or low-density equipment and local supplemental cooling [such as in-row, above-row, rear-door heat exchanger, or cabinets with supplemental cooling],” says Julian. “The one overriding limiting factor is power capacity. With adequate power, space and cooling limitations can most likely be overcome. High-density equipment may be more power-efficient but will need appropriate cooling—if these technologies use less space, then additional equipment can be installed, thus raising the total power used.”

Kightlinger adds that although it’s necessary to plan for growth, that growth might not occur in the way you expect. He recommends anticipating higher power and heat densities at the rack, row, and zone levels and ensuring that the infrastructure plan includes everything from the incoming utility to switchgear to the size of the generator, chiller plants, and UPS systems. According to Scherr, power companies can limit available power in certain locations due to local infrastructure constraints, so he advises checking with the utility provider to assess capacity and possibly devise an agreement for future power needs.

Whereas future changes in power requirements can be addressed before those changes occur by implementing power systems that can accommodate change, cooling can be trickier. However, it’s still possible to build a cooling infrastructure with future change in mind. For example, Scherr notes that it can be wise to pre-plumb a data center with water (for potential water cooling) and leave space for a future water chiller.

“Research, research, research,” Scherr says. “Know your applications and where you plan to go. Know your power consumption—don’t just use specs; test the product. Plan for future growth, then double [the capacity].” ■

by Christian Perry

Watch For Pitfalls

Establishing the correct amount of power and cooling for a data center is certainly challenging, but it's also fraught with potential peril. Jim Scherr, director of sales and operations at PDUs Direct (www.pdusdirect.com), identifies common pitfalls encountered in the planning process.

Insufficient circuit size. If you run low-voltage, low-amperage power to the racks, you could quickly run out of capacity. Envision your maximum size and place circuits accordingly. You may pay a little more up front in capitalized cost, but it will be less than adding them after the fact.

Basing the usage on manufacturer specs. Specs are published at maximum usage and load, but equipment usually operates at a much lower level. If you plan for 6,000 watts in a cabinet, the specs may indicate you can operate only 10 servers per rack. However, if you measure and monitor your actual consumption, you may find you can operate 15 servers per rack.

Insufficient outlets. Ensure your PDUs have enough outlets to run the circuit close to the derated maximum load. You do not want to have a circuit at half capacity and no more outlets.

Limited cooling limits growth. Install more cooling than you initially need. No one wants to limit their growth because they can't cool the room. Leave space for future additional coolers if possible or build big from the onset.

Copyright © 2009 Sandhills Publishing Company U.S.A. All rights reserved.