DATA CENTER Thermal Management

A holistic approach to energy efficiency.

BY IAN SEATON

How important is data center thermal management energy efficiency? How about cashing in a 50% reduction in cooling energy on a one megawatt data center for $1.9 million savings over the life of the data center? (Savings based on fifteen year life of data center, $0.08 per kilowatt hour electricity cost and power usage effectiveness improvement from 2.0 to 1.63.)

There are many different paths to improved data center thermal management efficiency, and both the technical literature and vendor claims can be confusing and contradictory. Rather than attempting to make a recommendation for a one-size-fits-all solution, there are surface issues that should be considered in any plan for maximizing the efficiency of data center cooling.

Figure 1: Average Data Center Energy Allocation. Source: Lawrence Berkeley National Laboratory

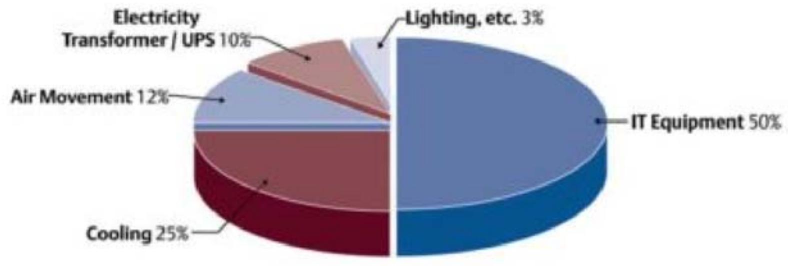

Power Usage Effectiveness

Power usage effectiveness (PUE) has become a near de-facto metric for describing data center efficiency. PUE is simply the ratio of information technology (IT) equipment energy consumption to the total data center consumption (i.e., total power ÷ IT critical load power). As shown in Figure 1, an average data center would have a PUE of 2.0. Based on the allocation of power in Figure 1, it is readily apparent that the greatest opportunity for improving PUE resides with reducing the energy for cooling the data center.

Among the several flaws with PUE is the fact that a data center with an overall high operating cost due to various types of server inefficiency could actually have a better PUE than a data center operating at higher density with fewer, more efficient servers and a lower total operating cost—reducing the divisor actually drives up the metric. Nevertheless, beyond the importance of achieving better utilization and efficiency in the server hardware itself, improving thermal management represents a significant opportunity for either reducing the total cost of operating a data center or extending the life of legacy data centers by re-allocating significant portions of the power budget from mechanical to critical IT equipment.

Harvesting Energy Savings

The two most important considerations during the development of a plan for harvesting significant energy savings from the data center thermal management architecture are:

- Maintaining recognition of the entire mechanical cooling process
- Defining the facility continuity objectives

Figure 2: Data Center Mechanical Cooling Generic Process

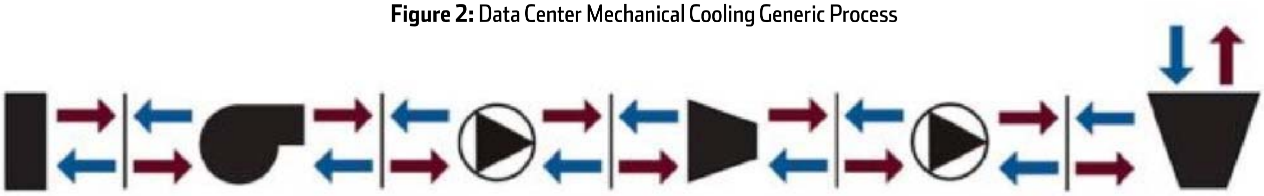

The entire mechanical cooling process includes all the heat exchanges beginning at the heat sinks or cold plates inside the servers and extending to the water tower or condensers where the thermal energy is finally removed, and all the sources of energy load (e.g., pumps and fans) moving that thermal energy from one heat exchange to the next.

Vali Sorrel and Terry Rogers of Syska Hennessy Group developed a simple "thermal sentence diagram" to describe the elements of this process as shown in Figure 2. Since two-thirds or more of the total cooling energy consumption is outside the data center in the chiller plant, it is important to understand the impact of all design options on the efficiency of the chiller plant. In addition, every heat exchange requires some level of energy input to move the thermal load to that exchange. Reducing the number of heat exchanges will also contribute to achieving a lower total thermal management operating cost.

Data center continuity requirements are typically described in terms of the Tier 1-4 levels developed by the Uptime Institute. At the risk of great over-simplification, this article assumes the following continuity levels for planning thermal management deployment:

- Neither fault tolerant nor maintainable without downtime
- Fault tolerant or maintainable
- Concurrently maintainable and fault tolerant
- Concurrently maintainable at original design-intent fault tolerance

Close-Coupled Cooling & Isolated Heat Removal

The importance of tracing the impact of all design decisions back through the entire mechanical process, as well as some hint of the issues to be considered, can be revealed through a comparison of a close-coupled cooling solution to a system of isolated heat removal (e.g., containment aisles or ducted exhaust cabinets).

Assuming a baseline of a close-coupled cooling solution with variable air volume (VAV) fans and computer room air handlers with VAV fans versus completely isolated supply and return air streams, the room with air handlers will have whatever slight efficiency advantage comes with the coefficient of performance (COP) advantage generally associated with the larger fans. In both situations, efficiencies are realized by raising the supply air temperature, raising change in temperature across the evaporator coils, raising chilled water loop temperatures and increasing available water-side economization hours. That baseline will be modulated by any number of the following factors:

- Computer room air handlers without VAV fans would mean that cooling capacity could only be adjusted in increments of a cooling unit, and the close-coupled solutions would therefore gain any affinity law advantage (see Figure 3), particularly when server cabinet loads were below maximum capacity.
- Close-coupled solutions typically add one extra step of heat exchange at some type of consolidation unit between the data center chilled water loop and the chiller plant loop. There is an approach loss associated with each heat exchange. Therefore, that higher temperature will translate into a slight chiller plant efficiency and economizer hours advantage for the computer room air handlers.
- According to the D'Arcy-Weisbach equation, pressure drop increases as a factor of the ratio of duct length to duct diameter (see Figure 4), so long overhead ducted delivery systems can reduce or even eliminate the efficiency advantages of computer room air handlers or air-side economizers during mechanical cooling hours.
- While close-coupled cooling solutions provide nearly similar benefits as computer room air handlers associated with isolated supply and return air streams with respect to available hours for water-side economization, most close-coupled solutions will not

The Affinity Law

Also known as the cube effect of fan speed, whereby the required fan horsepower varies as the cube of the speed:

$$hp2 = (rpm2 / rpm1)^3 \times hp1;$$

While fan air delivery varies directly as the speed varies (linear relationship):

$$cfm2 = (rpm2 / rpm1) \times cfm1;$$

80% RPM = 80% air volume = 51.2% energy required (0.8^3)

Figure 3: The Affinity Law

Pressure Loss in a Duct

$$\Delta p = \lambda \, (l / dh) \, (pv^2 / 2) \quad (1)$$

Where:

Δp = pressure loss (Pa, N/m^2)
λ = D'Arcy-Weisbach friction coefficient
l = length of duct of pipe (m)
dh = hydraulic diameter (m)
p = density (kg/m^3)

Figure 4: Pressure drop increases as a factor of the ratio of duct length to duct diameter

be compatible with air-side economization or energy recovery wheel heat exchange. Both usually provide significantly more hours of cooling without running the chiller plant. In a data center with complete supply and return isolation, the computer room air handlers can be replaced by the economization air handlers and cooling coils.

Project continuity requirements likewise have a significant impact on data center thermal management operating efficiency and associated design decisions.

- Close-coupled solutions have VAV fans directly coupled to the air volume demand of the associated server load and will therefore gain the affinity law advantage particularly in Tier 1 and, to some lesser degree in Tier 2 data centers. For example, if one server cabinet is at 60% of capacity, another at 80% of capacity, and another at 30% of capacity, the total energy requirement for the close-couple fans would be around 25% of full capacity (.60^3, .80^3 and .30^3).
- Where continuity must be simultaneously maintainable and fault tolerant, there is an advantage to separating the cooling source from the heat load to minimize the serial levels of redundancy. For example, in a space where 120 tons of cooling is required, and it has been determined that N+2 will meet the continuity requirements, six 30-ton air handlers equipped with VAV variable frequency drive (VFD) fan motors or electronically commutated (EC) fans would operate normally at 67% air movement capacity or 30% energy consumption per unit, or 45% total by virtue of there being six units running at 30% versus four units running at 100%. The energy savings actually increase during the phase of IT equipment deployment ramp-up at lower loads due to the cubing effect of the affinity law.
- Since close-coupled cooling solutions directly couple cooling capacity to load, the fail-over for maintenance or fault-mitigation will typically need to be either a totally separate reserve depending on reliable actuation or a separate running capacity running in the background and on top of the directly scaled cooling capacity.
- When isolation between supply air and return air is achieved through cold aisle containment and fault-tolerance is required, air handler fans will need to be on uninterrupted power supplies (UPS) and will be subject to efficiency losses associated with UPS operation.

Final Considerations

Close-coupled cooling solutions and isolated supply/return air solutions represent opportunities to significantly improve the energy efficiency of data centers by eliminating waste, maximizing change in temperature between supply and return, increasing chiller efficiency with higher temperatures, and exploiting affinity law efficiencies of VFD and EC fans. In addition, the greatest chiller efficiency is achieved when the chiller compressor is not running. Both approaches provide access to increased hours of water-side economization, and the data center design based on isolation further provides access to increased hours of air-side economization and energy recovery wheel heat exchange cooling. Since the greatest saving benefits derive from the most extreme temperature differences, plumbing, chillers and evaporator coils all need to be specified for flow rates appropriate for the total design. Finally, design decisions need to exploit continuity requirements and the impact on the total system of each discrete element. ■

Ian Seaton has over thirty years of product development experience in electromechanical systems. He is currently the technology marketing manager for Chatsworth Products, Inc. and can be reached at iseaton@chatsworth.com.