

Green field data center design
11 Jan 2010
by Shlomo Novotny



Shlomo Novotny, Vice President and Chief Technology Officer, Vette Corp.
explores water cooling for maximum efficiency - Part 1

Overview

Data centers are an ever-growing part of our economy. The IT community is experiencing constantly increasing demands in areas such as Internet media and video, banking and finance, research, and government, just to name a few. Consequently, data centers continue to grow in both size and numbers. In addition, the performance of IT servers continues to follow Moore's law; this improves the performance per dollar spent on IT equipment, but also increases the heat density inside the data center equipment. Driven by these factors, the cost of running a data center continues to grow, even as the cost of the IT performance continues to decline.

Historically, data center management has been focused on growth, performance, uptime, and reliability. IT organizations generally had responsibility for data center management. Facility organizations were responsible for much of the infrastructure surrounding the IT equipment. As a result of the significant growth and the lack of a holistic view of the economics of data centers, data center facilities today can be compared to factories of the 1970s and 1980s: ripe for significant cost, design, and process improvements. With attention from C-level executives today, data centers are driven to be lean, scalable, cost effective, and green.

One of the most inefficient issues with many data centers today revolves around cooling. Typically, data center operators have managed heat density issues by spreading the load and under populating racks. This has come about because of the limitations of air cooling, which creates a very costly, complex, and inefficient infrastructure. Localized, passive, low-power dissipation liquid cooling devices at either rack level or rack proximity, when compared to traditional air-cooled methods, have the capability of reducing the power consumption of in-room cooling devices by as much as 90%. In addition, by allowing data centers to operate with higher-density racks, rack-level liquid cooling can reduce the data center IT footprint by as much as 80%.

Liquid-cooling technology also permits a "pay as you go" cooling implementation, saving significant CAPEX when constructing a new data center. Present air-cooled data center construction requires an implementation from "day 1" of all the cooling hardware for proper operation of the IT room. Localized liquid cooling would be implemented with IT expansion. Cooling hardware is modular and is purchased at the same rate as the IT hardware.

This white paper addresses the emergence of localized liquid cooling, rack level or rack proximity, as the basis of data-center design, which will provide the most effective, efficient, and sustainable method to cool today's high-performance data centers.

Data Center Equipment Power Trends

By the year 2000, the IT industry had accelerated its transition from CPU speed-based performance to high-density system-level performance. The competitive focus moved to performance per cubic foot of packaging. Two popular server packaging form factors have evolved rapidly: the 1U and the blade. These two form factors populate rack structures in order to obtain maximum packaging density and maximum performance at lower cost per computer operation. Roger Schmidt, Christian Belady, and the author of this paper, Shlomo Novotny, identified the major challenges that the need to obtain maximum performance per volume packaging introduces to system cooling in data centers.

Fully populated rack power dissipation had grown dramatically from an average of 1–1.5 KW per rack to an average rack-level design from IT companies of 12–25KW per rack by 2000; however, these designs were not implementable because of the high power and cooling requirements. Typical data centers at that time, with a cooling capacity of 100 W/ft², were limited to cooling 1–3 KW per rack without major investments. Many, if not most, data centers still operate in that range today. To create awareness of this challenge, an industry consortium initially led by Schmidt, Belady, and Novotny was created and developed the power dissipation charts that were then published by the Uptime Institute.

To further bring this problem to the industry's attention and promote the introduction of new technologies that would mitigate this problem in data centers, the consortium's efforts were moved to a new group, TC 9.9 (Technical Committee), within ASHRAE. Because it represents a larger cross section of industries, including IT, construction, and HVAC, all of which relate to data centers, ASHRAE was the appropriate institution to carry the charter to create guidelines for data center designs. To date, ASHRAE has published multiple books dealing with the densification and energy savings in data centers, including liquid cooling and "free cooling."

In 2005, ASHRAE published *Datacom Equipment Power Trends and Cooling Applications* (Reference 2) with the power curves first published by the Uptime Institute and updates that included blade servers, which were not in the original plot. Figure 1 depicts the extent of the power density problem. It shows that from a maximum of approximately 5 KW/server-rack in the year 2000, the maximum power was projected to jump to 15–30 KW/server-rack in 2010. Of course, the extreme high-performance computers (mainframes) are shown as dissipating even higher power levels.




Figure 1. Power Density Problem over Time

For a bigger version of the image [click here](#)

Part II: Key Drivers affecting energy consumption, the impact of power density on cooling. See our [Thermal Management Knowledge](#) bank for more cooling news, features and analysis

About the author: Shlomo Novotny, is Vice President and Chief Technology Officer, Vette Corp.



Keywords: Data centers, energy, consumption, EPA, US, liquid-cooling, power, Koomey, ASHRAE,

Green field data center design – water cooling for maximum efficiency - Part 2

14 Jan 2010

by Shlomo Novotny

Key drivers for why water cooling should be considered - Shlomo Novotny,

Key Drivers

J.G. Koomey in the United States Environmental Protection Agency's Report to Congress on Server and Data Center Efficiency, August 2007 (Reference 7) identified that data centers

consume 1%–2% of the total U.S. energy consumption. Energy consumption in data centers has been growing, according to the EPA, at 12% a year (Reference 7). Increasing heat density, which demands more cooling, and the cost of electricity are pushing data center operating costs to exceed the costs of the IT equipment itself. The energy cost per server operation is now larger than the server acquisition cost, as shown in Figure 2, generated by Christian Belady and included in the EPA's report to Congress (Reference 7).

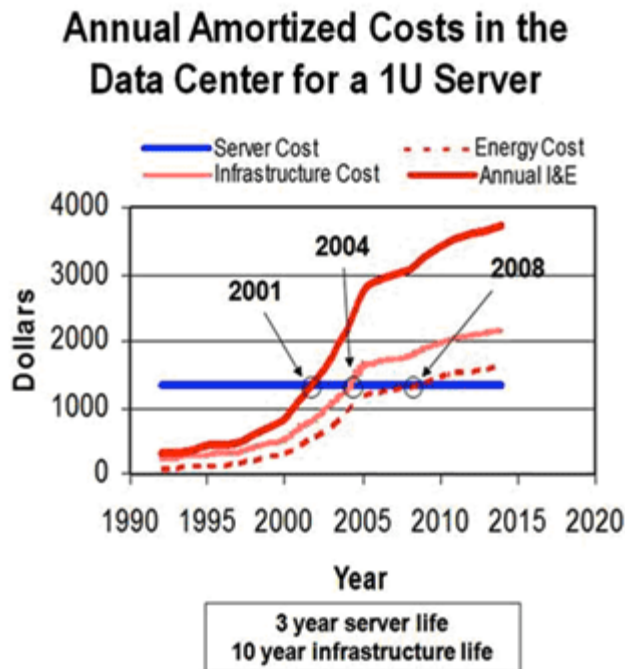


Figure 2. Server Energy Costs versus Acquisition Costs

Looking at it another way, in a report published in September 2009 by Koomey et al. (Reference 3), as performance per dollar spent on IT hardware increases, the percentage of infrastructure and related costs increases relative to the cost of the infrastructure plus equipment. This trend is presented in Figure 3.

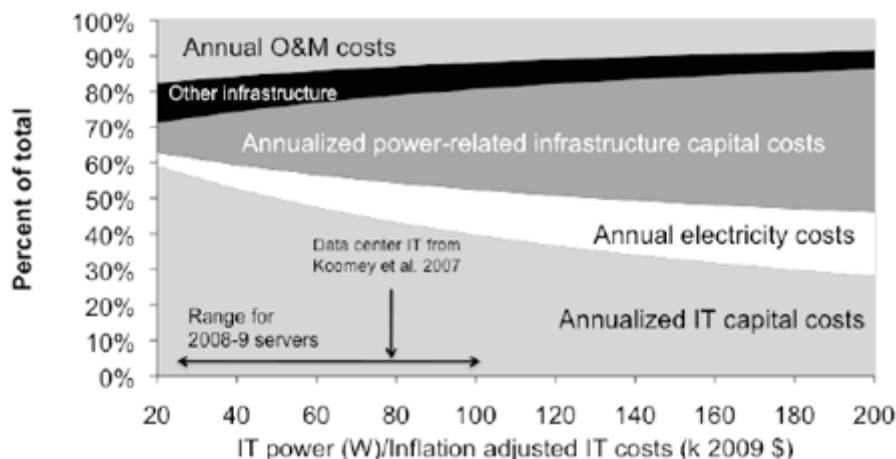


Figure 3. Annual Infrastructure and Equipment Costs

In addition to the fact that infrastructure needs are increasing, one of the most significant costs, energy, will, by most predictions, continue to increase in price. Thus, the operational cost as related to energy consumption has become a major part of data center OPEX. Power infrastructure has been a major CAPEX in data center construction, according to the Uptime Institute. As a result, and in response to concerns regarding inefficient energy usage in the data center, the IT industry created and is heavily funding the Green Grid organization. This organization, jointly with ASHRAE, is promoting energy conservation and efficiency in IT equipment and in data centers.

Data Center Cooling Today

The traditional way to cool electronics in data centers has been with air, using computer room air conditioning (CRAC). A traditional data center has been designed to cool an average 75–100 KW/ft², which translates to 1–3 KW/rack. Newer, more expensive data centers are designed to cool an average 200KW/ft², which still limits the power density per rack to 4–5 KW (recall that full rack capacity is 25 KW/rack). The traditional architecture employs CRAC units at the periphery of the data room, utilizing chilled water from an outdoor chiller as described below in Figure 4.




Figure 4. CRAC Unit Deployment

For a bigger version of the image [click here](#)

The air is distributed through an air-pressurized raised floor plenum where the air arrives at the racks through perforated tiles in the floor in front of them. Most of the energy consumed inside the data center in this process is consumed by very powerful air blowers. The air itself is not an efficient heat carrier, and the airflow from the servers must be matched to the airflow coming from the perforated tiles in order to minimize air circulation and thus prevent overheating of the servers mounted on the upper part of the rack. Multiple papers have been written on these subjects, and the computational fluid dynamics companies modified their programs to enable efficient analyses of the airflow and optimization of server location and airflow distribution.

As data centers install high-power-dissipation racks (up to 25 KW), air cooling is becoming challenging, if not impossible. The traditional way of cooling such high-power devices is through low-populated racks (low power per rack), rack spreading (getting more air from the floor), hot and cold aisle containment, and creation of small dedicated areas for high-power-density racks. The main problem with all the solutions optimizing and utilizing airflow and CRAC units is the thermal and energy inefficiency of heat transport by air, not to mention the overall complexity of the solutions. Thus, the solutions are typically costly from CAPEX and OPEX points of view and

do not minimize the thermal energy usage and cost. Furthermore, new data center construction requires that all cooling systems be purchased and installed prior to populating the room with IT equipment, thus requiring a large CAPEX expense prior to the actual cooling need. And in order for the cooling to operate properly and deliver the flow and pressure required by the racks at any location in the data center, all cooling equipment must be not just installed but fully operational from “day 1,” with resulting inefficient energy consumption until the data center is fully populated. Some hardware suppliers use water- or liquid-cooled cabinets or in-row units for cooling “hot spots” (localized high-power-density racks). Although more efficient than CRAC units, these liquid-cooling units still utilize very powerful air movers that dissipate a significant amount of power. Most present solutions are described in *Datacom Equipment Power Trends and Cooling Applications*, pages 32–38.

Water — an Efficient Cooling Alternative

The most energy-efficient and cost-effective way to remove the heat from the rack is by extracting the heat at the source (the rack), utilizing the airflow built in the rack through its servers and transporting the heat with liquid, which is significantly more efficient than air. Water is 3400 times more efficient than air in removing heat. This heat extraction at the rack level could be done by a passive rear door heat exchanger (RDHx) consisting of a liquid coil or by liquid-cooled cold plates mounted inside the servers for removal of heat from the power-dissipating components. Alternatively, a hybrid of passive rear doors and cold plates for high-power components inside the servers can be used. In any of these designs, one could eliminate or minimize the use of air cooling generated by the CRAC units. In addition, as much as 80% of the IT area could be reduced by densification utilizing liquid cooling. Furthermore, by fully populating the racks, one could obtain savings in CAPEX by eliminating excess racks and other ancillary equipment.

Description of the Water Cooling System

Figure 5 is an illustration of a rack-level water cooling system. Although the system is shown here with hoses running underneath the floor, a raised floor is not needed for its implementation, thus offering more installation flexibility and saving in CAPEX. In addition, this system could be used as a retrofit or as a basis for the design of new data centers. The chiller (“chilled water”) shown in the illustration could be eliminated or supplemented with a water-side economizer, thus obtaining “free” cooling.




Figure 5. Water-cooled Data Center

Part III:

Keywords: IT, Data Center, cooling, water cooling