

GREEN FIELD DATA CENTER DESIGN – WATER COOLING FOR MAXIMUM EFFICIENCY

Shlomo Novotny, Vice President and Chief Technology Officer, Vette Corp.

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.
Source: ASHRAE

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).

Annual Amortized Costs in the Data Center for a 1U Server

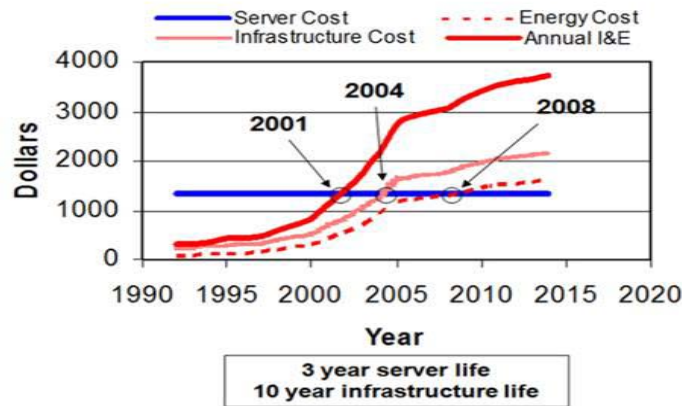


Figure 2. Server Energy Costs versus Acquisition Costs
Source: Electronics Cooling Magazine

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 W/ft², which translates to 1–3 KW/rack. Newer, more expensive data centers are designed to cool an average 200 W/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

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

Rear Door Heat Exchanger (RDHx)

Heat removal from a rack is most efficiently done by locally coupled heat removal devices. Ideally, these heat removal devices are passive, thus utilizing the airflow from the IT equipment mounted within the rack. The device shown in Figure 5 is a passive rear door heat exchanger. The RDHx replaces the existing perforated rear door commonly used in IT racks for protection. The RDHx consists of a fin and copper tube air-to-water heat exchanger that absorbs the heat flowing through it, utilizing pumped water to transport the heat away from the “hot” aisle. The RDHx introduces a very minimal increase in pressure drop over the existing rear door, thus the rack airflow is minimally impacted, resulting in negligible internal temperature rise. The additional pressure impact on the internal fans’ power consumption has been calculated to be in the 5% range per typical rack of 8–15 KW dissipation. Also, being passive, the RDHx does not increase the acoustic noise in the data center.

Coolant Distribution Unit (CDU)

The coolant distribution unit (CDU) shown in Figure 5 is the device that separates the outside water, which typically includes glycol for antifreezing, thus reducing the liquid heat transfer capacity, from the secondary liquid loop leading to the racks. This loop could carry either refrigerant or water. Water in this loop would have a high heat-transfer capacity, as it does not need glycol. In addition, in order for the components to achieve long life, the water should include a mixture of antibacteria and anticorrosion materials. In a very large data center, this CDU could be replaced with larger centralized CDUs, saving on equipment cost. In areas where freezing is not an issue, the CDU could be eliminated and the secondary conditioned water loop could come directly from the chiller.

Thus, by extracting the heat at the source, new data centers could eliminate traditional air cooling by using only liquid-circulating pumps to transport the heat to a chiller or to a water-side economizer. A few traditional units may be needed for humidity control. The water cooling would permit densification of the racks and the data centers while reducing the energy consumption significantly.

Modularity

A traditional data center requires, as mentioned above, that all cooling and air distribution systems be installed in new construction prior to populating the room with IT equipment. However, very few data centers are fully populated with IT equipment upon operational start-up. Therefore, the data center owner is committing significant up-front funding to enable future IT growth. For example, an owner of a 5 MW data center with water-cooled CRAC units would be spending approximately \$1.8M for CRAC units before any IT equipment has been installed. Assuming that initially only 30% of the data center would be populated, the owner could see the initial investment reduced by \$1.2M if a modular solution is implemented. Localized liquid cooling solutions would enable this savings. The liquid cooling solution could be implemented locally where and as required by the IT equipment and is scalable as the data center populates more of the IT equipment. In addition to energy savings from water cooling efficiency, additional operational savings will be realized by not having to supply cooling to the unpopulated IT areas.

Redundancy and Availability

Maximum availability and redundancy of the cooling system in an air-cooled data center is accomplished by adding 20% or more CRAC units for extra cooling capacity in case of failure or when there is routine servicing of a CRAC unit. Liquid-cooled data centers accomplish redundancy and availability more efficiently. The liquid-cooled data center uses just the right amount of cooling needed for the IT racks. If a failure occurs at one or more RDHx units, the racks continue to operate normally with respect to airflow because the additional resistance to the flow that an RDHx unit introduces is minimal. What happens in the unlikely event of a failure in the secondary liquid loop is that the ambient temperature will rise because the heat generated by the rack would not be removed by the water. When the ambient temperature rises, the other RDHx units in the room will absorb the excess heat, as depicted in Figure 5. The water-cooled data center would be designed with multiple interleaved secondary loops, thus ensuring that any single failure or major service event would not affect IT performance and availability.

Roger Schmidt and Madhusudan Iyengar, of IBM, demonstrated in their July 2009 paper presented at InterPACK 2009 that a failure of one loop of five RDHx units did not cause the temperature in the data center to exceed the ASHRAE-recommended limits, thus showing the robustness of the water cooling system as described in the previous section.

The New Green Field Data Center — Taking Energy Efficiency to the Next Level

A power consumption comparison of both air and water cooling techniques is shown in Figure 6, generated by comparative testing performed by the Silicon Valley Leadership Group a year ago (Reference 6). It clearly demonstrates that passive rack-mounted liquid-cooled rear doors are significantly more efficient than any other means of heat removal.

Cooling Solution Comparison From Sun Microsystems "Energy-Efficient Modular Cooling Systems" Case Study presented at the SLVG Data Center Summit June 26, 2008

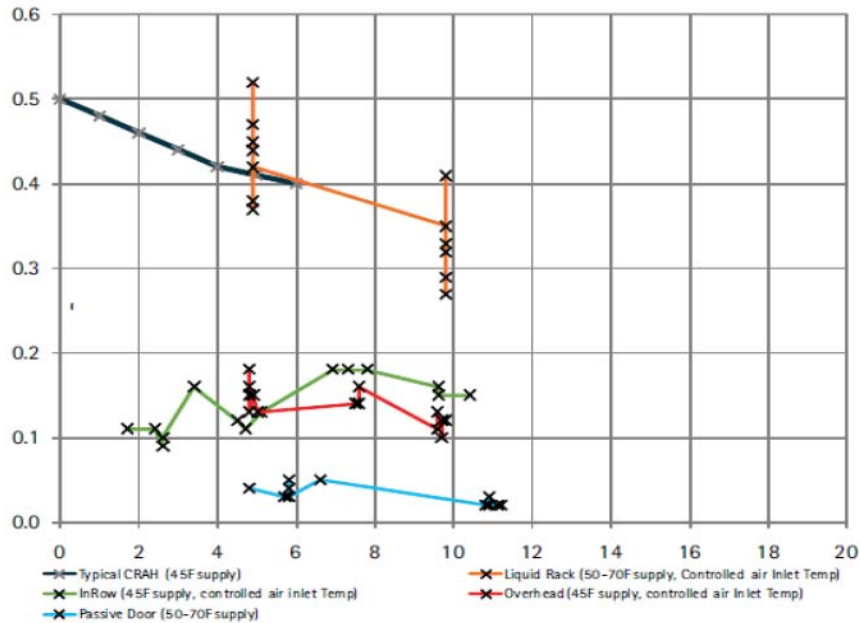


Figure 6. Sun Microsystems' "Energy-Efficient Modular Cooling Systems" Case Study as presented at the SVLG Data Center Summit

Two cooling systems in the data center consume most of its cooling energy: the chiller supplying the chilled water and the air-movers (high-pressure blowers) supplying the air for heat removal from the IT equipment. As the previous paragraph has indicated, localized liquid cooling, direct and indirect, could replace air cooling in the data center. Figures 7 and 8 show two performance examples of passive rear door heat exchangers.

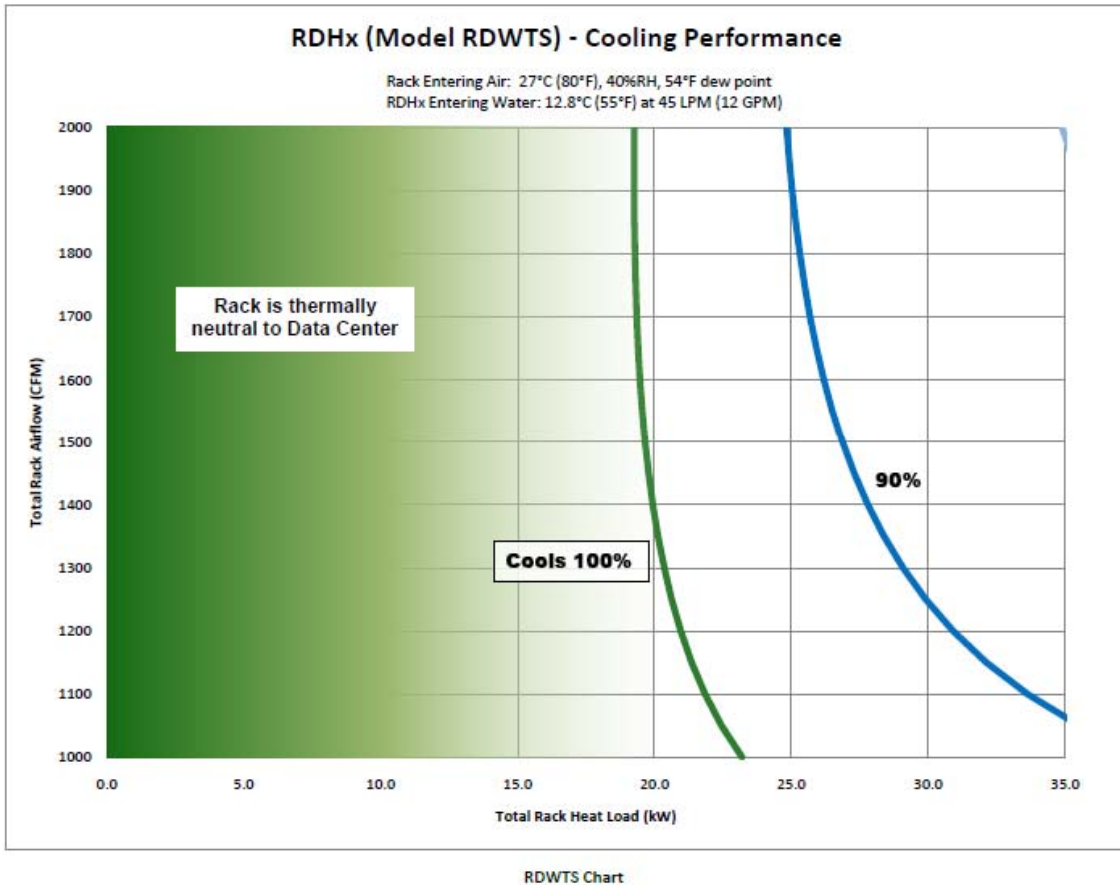


Figure 7. RDHx Cooling Performance Example 1

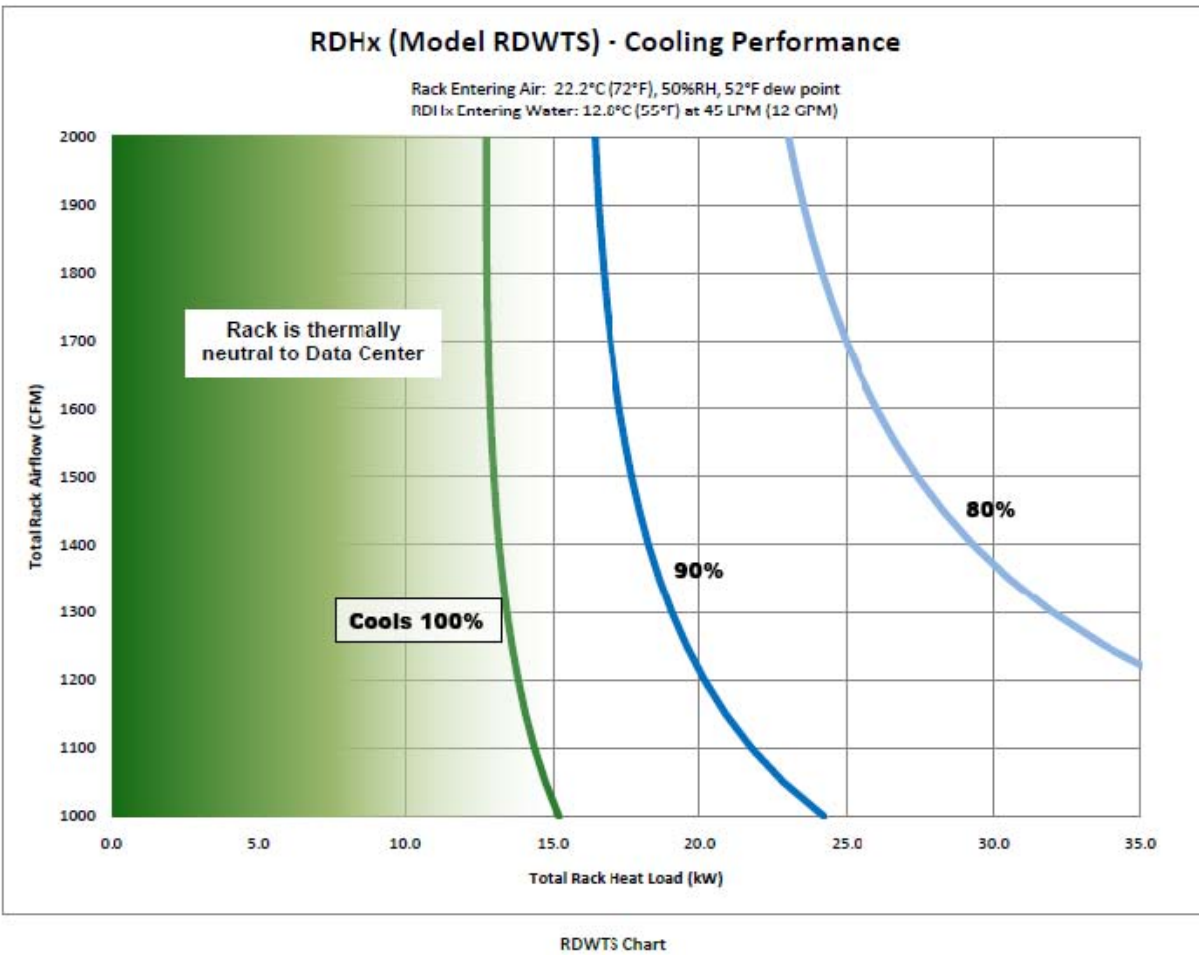


Figure 8. RDHx Cooling Performance Example 2

What these results indicate is that most racks used in data centers today could have their power dissipation neutralized with an RDHx, as most dissipate less than 18–20 kW per rack. Higher-power-dissipation racks could use a higher-performance rear door for neutralizing their load.

One of the most important cooling efficiency trends in data centers is the raising of the data center ambient temperature. As Figures 7 and 8 indicate, the efficiency observed from the use of an RDHx is improved by higher ambient temperatures. Higher temperature (the new ASHRAE guidelines limit) enables the maximum removal of heat from the rack. At the same time, higher temperatures significantly improve the chiller efficiency and performance, thus saving additional energy. According to Chris Kurkjian (Reference 4), a centrifugal chiller would save 5% in energy consumption for every 2°F increase in chiller water temperature and would produce 19% improvement for every 10°F decrease in condenser water temperature. Furthermore, by using a water-side economizer that connects through a heat exchanger to the RDHx when outdoor temperature is appropriate, energy usage is further minimized.

The question, of course, that would be asked focuses on the RDHx heat removal capability at higher water temperatures from chillers or from water-side economizers. Figure 9 shows the RDHx performance sensitivity to water temperature at 80°F ambient air temperature, 30%RH, 46°F dew point, 2400 CFM per rack, and 10 GPM water. As Figure 9 shows, the difference in performance between 48°F and 66°F is 30%. Referring back to Figure 8, it is apparent that the neutralization capability declines from approximately 20 KW to 14 KW. This heat-extraction capability still covers most of the high-density racks used in data centers today. Therefore, a data center operator could increase the air temperature and water temperature without affecting reliability, but with significant benefits to operating cost resulting from significant savings in energy expenses.

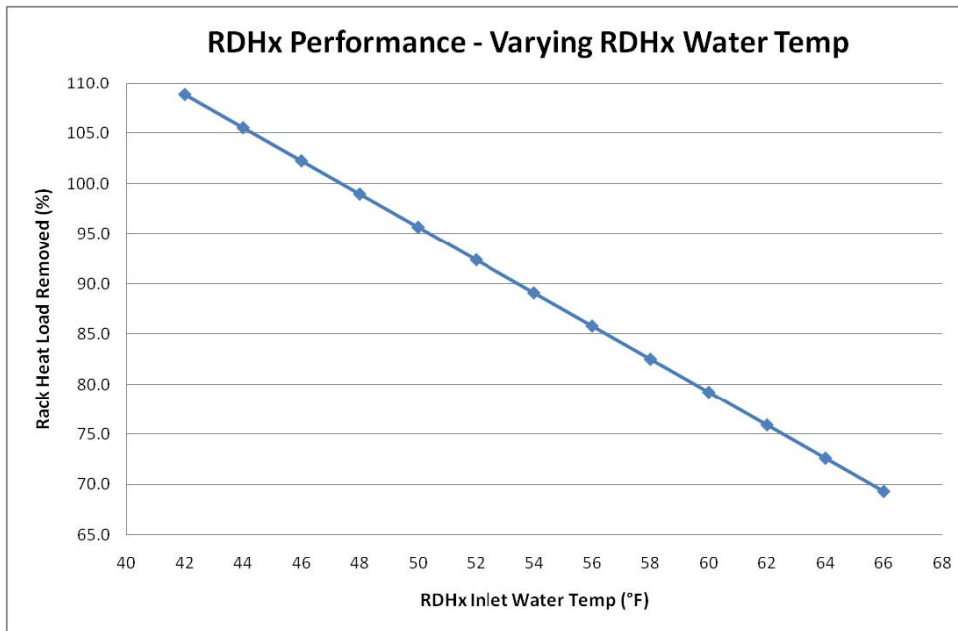


Figure 9. RDHx Performance Sensitivity to Water Temperature

As the RDHx utilizes the rack internal airflow for cooling, the concern may arise regarding the variation in rack internal airflow impact on the RDHx performance. Figure 10 clearly indicates that the RDHx performance sensitivity to internal airflow variations is minimal.

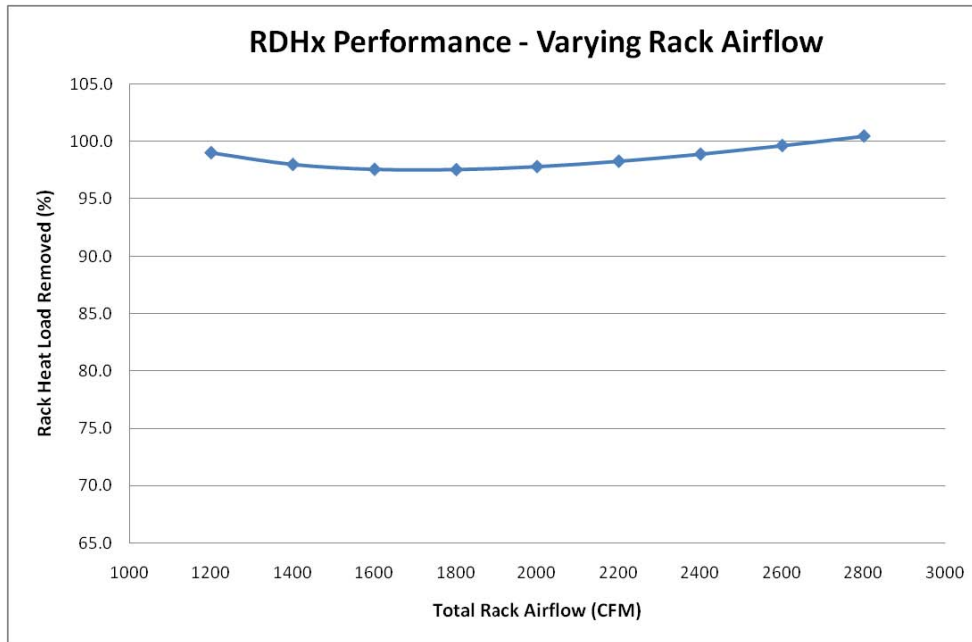


Figure 10. RDHx Performance Sensitivity to Airflow

Conclusion

Data center operations have historically focused on addressing the accelerating demands of data center customers. Because of their primary focus on growth, data centers have become highly inefficient operations. Today, with other factors such as rising energy costs, limited space, and power and green initiatives, management’s focus is squarely on the cost and operational inefficiencies that have been building up in data centers.

With infrastructure costs now exceeding the costs of the IT equipment itself, management must focus on gaining significant efficiencies in the data centers, which are becoming the factories of the 21st century. With cooling being a major portion of data center infrastructure costs, sizable efficiencies can be found by rethinking how data centers are cooled.

Localized liquid cooling at rack level and rack proximity eliminates one of the most inefficient elements of data center infrastructure cooling, substantially reduces energy consumption, and allows for higher rack-level compute density. In addition, liquid cooling enables cooling modularity with a “pay as you grow” investment and ensures efficient infrastructure redundancy and availability. Localized liquid cooling is no longer simply a hot-spot solution, but is now a critical component in the basis of design for new, sustainable data centers.

References

1. American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE), 1 August 2008, *2008 ASHRAE Environmental Guidelines for Datacom Equipment*, available at http://tc99.ashraetcs.org/documents/ASHRAE_Extended_Environmental_Envelope_Final_Aug_1_2008.pdf.
2. American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE), 2005, *Datacom Equipment Power Trends and Cooling Applications*.
3. Koomey, J.G., Belady, C., Patterson, M., Santos, A., Lange, K-D., 17 August 2009, *Assessing Trends over Time in Performance, Costs, and Energy Use for Servers*.
4. Kurkjian, C., PE, 18 May 2009, *Data Center Energy Savings — by the Numbers*.
5. Schmidt, R., and Iyengar, M., July 2009, *IBM's Rear Door Heat Exchanger and the New ASHRAE Recommended Environmental Guidelines*, Proceedings of IPACK 2009 and ASME InterPACK 2009.
6. Sun Microsystems, 26 June 2008, *Energy-Efficient Modular Cooling Systems*, Case Study presented at the SVLG Data Center Summit.
7. United States Environmental Protection Agency, 2 August 2007, *Report to Congress on Server and Data Center Efficiency*, available at http://energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.

About the Author

Mr. Novotny is the Vice President and Chief Technology Officer of [Vette Corp.](#) As a world-recognized expert in electronic cooling, he is one of the few industry specialists in electronic cooling and packaging to achieve the rank of A.S.M.E. Fellow. Mr. Novotny is a co-founder of the ASHRAE TC 9.9 cooling consortium to create an industry standard for data center cooling, as well as being an influential member of the PCI and InfiniBand industry committees that defined new interconnect bus architecture standards. He holds 16 patents to date, with several patents pending. Mr. Novotny is an A.S.M.E. Fellow, an I.E.E.E. senior member, and a member of I.M.A.P.S., and ASHRAE-TC 9.9.