White Paper
Digital Enterprise Group

The State of Data Center Cooling

A review of current air and liquid cooling solutions.

The challenges of increased data center density elevate the importance of the right cooling technology. This paper discusses liquid and air cooling solutions in data centers and when each makes the most sense.

Michael K. Patterson and Dave Fenwick
Intel Corporation

March 2008

Contents

Executive Summary

This paper provides basic guidance for selecting the right cooling methodology for multiple types of data center projects, as well as varying densities of new servers. We focus on proven cooling technologies you can use right now, rather than emerging technologies.

Data center provisioning is a multidimensional problem with a variety of constraints. Those constraints drive the choices, and choices need to be compared based upon a thorough analysis of each option's total cost of ownership (TCO).

In this paper we describe liquid and air cooling solutions and discuss where each is best applied. Liquid cooling generally tends to be more expensive and not necessarily more efficient. Air cooling faces limitations from rack density, though some of these can be surmounted by various airflow segregation strategies.

The paper also includes several case studies. Of particular note are two recent new installations on the same site that resulted in one using a water cooling system and the other air, each solution appropriately applied based upon the constraints of each different project.

The results in the paper are based on Intel's direct experience, as well as detailed analysis of some of the potential options.

Business Challenge

Increasing compute capabilities in data centers has resulted in corresponding increases in rack and room power densities. How to cool these new higher-powered racks is a question that challenges all data center managers. In the past it was sufficient to roll in new servers and just add computer room air conditioners (CRAC units) around the perimeter. Not a lot of engineering was required. When data centers were in the range of 75-100 watts per square foot (nominally square foot of total raised floor area), this method generally worked fine and the energy costs (and losses from inefficiencies) were small enough that they did not receive much attention. However, today's IT equipment can push data centers to 750 watts per square foot and the cooling challenges and costs are much more obvious. It is no wonder that at a 2007 Data Center Users Group meeting sponsored by Liebert (a provider of power and cooling technologies), the 107 participants listed as their top three facility/network concerns: heat density (cooling), power density, and energy efficiency (energy costs and equipment efficiency).[1]

Looking at power density, today's increase is driven primarily by the ability to pack an ever greater amount of performance into today's servers. Intel recently completed an internal study that showed an HPC machine (capable of 3.7 TFlops) that was in the top 20 on the Top 500 list in 2002 consisted of 25 racks and consumed 128 kW (~5.1 kW/rack). That same 3.7 TFlops can now be had in a single rack. That's the good news. The challenge is that the power needed by that rack is roughly 21 kW. If an appropriate cooling solution can be put in place, all the benefits of this increased performance—fewer servers, one fifth the total power, and 1/25 the floor space—can be realized. Intel believes this trade off is well worth the investment in the more advanced cooling systems. The other advantages of high density and a lower TCO associated with it are discussed in detail by Patterson.[2]

The power density trend will continue to challenge cooling technology. ASHRAE has published a trend chart and it is included in this paper (see Figure 1). We believe this chart is fairly accurate through 2014. For instance, it suggests that a full rack of the densest servers in 2008 will require about 30 kW of power in operation and cooling (heat removal). This matches our experience with the peak power draw of a full rack of dual processor (DP) peak-performance quad-core blade servers.[3]

From now until 2014, the trend should linearly increase to roughly 37 kW per rack when fully configured with the latest products. Interestingly enough, the increase does not come primarily from CPU power increases, but comes instead from additional memory in the platform. CPU power levels have been fairly well controlled or capped and will remain inside their nominal power envelope for the next several years. The memory (and associated power) increases are driven by the increase in core counts in the CPU. Optimum platform architecture generally needs a number of DIMMs or amount of memory per core, and as the number of cores continues to increase, the installed memory will too. Memory cooling thus becomes the primary thermal challenge for the years ahead.

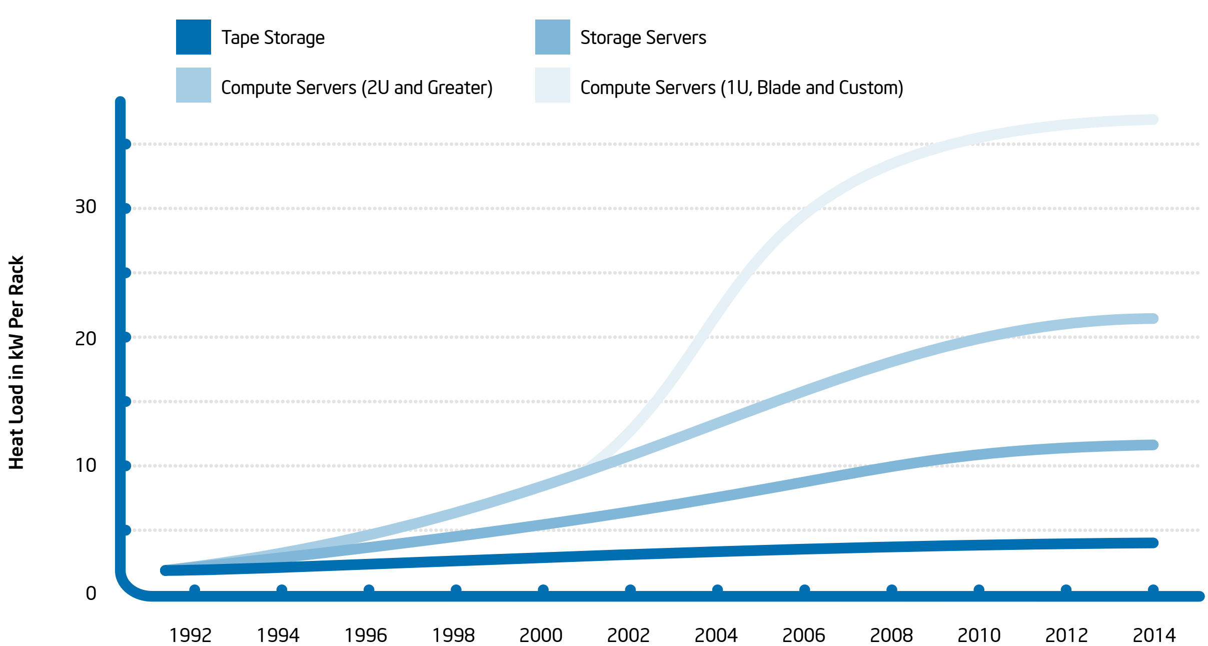

Figure 1: Adapted from ASHRAE's projected density loads (kW per Rack) for IT equipment.[4]

Technical Challenge

Data center cooling is primarily a thermodynamics problem. There is a heat source (the server) and a heat sink (typically the outdoor environment). Naturally, there are other components in a data center responsible for heat as well, but for simplicity's sake in this section of the paper we will focus on the server.

Depending on the temperatures and proximities of both the heat source and the heat sink, the movement of heat energy from the server to the outdoors will have a range of efficiency, particularly in how much additional energy is needed to accomplish adequate transfer. It is important to consider these additional energy uses in the overall system. The first is the energy involved in moving the fluid (air, water, or liquid) to carry the heat between the source and the sink. This is typically a fan or pump or both. It is important to understand that there may be multiple stages of fluid movement and that the energy use can be significant.

Another energy user is that of the chiller system. In the very common situation where the heat sink temperature is too high for efficient energy transfer, a chiller system can create a low temperature intermediate sink (such as chilled water) for the heat to move to from the source. The chiller plant then must "pump" the heat from the low temperature intermediate sink into the final sink (again typically the outdoor ambient). This is most often done with a standard refrigeration cycle, where a refrigerant is compressed and heats up. This heat can then be ejected to the sink. The cooler fluid is then expanded and is at a much lower temperature. This intermediate sink is then used to remove the heat from the source or a cooling fluid stream. This cycle repeats itself to produce the needed cooling effect, but at an added energy cost.

The intent here is not to provide a treatise on thermodynamics but instead to point out the complexity involved and that all three expenditures of energy must be evaluated in determining the efficiency and subsequent cost of the cooling system.

Cooling Solutions

We will now consider different types of cooling as solutions to the heat density question. We begin with definitions of air and liquid cooling, followed by discussions of their advantages and capabilities. We close by discussing how each would fit into three different scenarios: 1) addition to a legacy data center; 2) expansion into an existing non-data center space; and 3) the design and construction of a new data center from the ground up (frequently referred to as "greenfield" deployment).

Liquid Cooling

The use of liquid cooling is certainly not new and, in fact, has an extensive history. In this paper, we discuss it primarily in the context of liquid-cooled racks. Currently most data centers use liquid cooling to the CRACs or air handlers and associated cooling coils. The use of liquid-cooled racks is primarily an extension of the existing liquid-cooling loop closer to the IT equipment.

Liquid cooling in or to the server itself (i.e., crossing the boundary between the server and rack or fully contained in the server) is not included in this discussion. In the case of liquid cooling inside the server, this is already in place in one sense, with regard to the CPU. The use of heat pipes or vapor chambers is growing and represents a variant of liquid cooling, but still must be considered air cooling as the heat-carrying medium out of the server is generally still the server airflow. Potential future pumped-liquid loops inside the server also fall into this category.

Liquid crossing the server/rack interface holds potential for higher heat dissipations, but also adds significant cost to the overall data center and complexity to the operational use model of the servers. Because of the cost and complexity, Intel does not believe this is an appropriate mainstream solution and expects this type of liquid cooling to be limited to special high-density situations.

Liquid cooling to the rack, which we do consider here, can take several forms. One is a liquid-cooled door, where the door, typically on the back of the rack, cools the air leaving the rack to or near ambient room temperature, essentially making the rack thermally neutral to the room.

Another variation is a closed-liquid rack where the rack is sealed and the airflow is fully contained within the rack, passing through a heat exchanger typically at the bottom. Heat is removed in a connected liquid-cooling system. The advantage to this configuration is that it is thermally and airflow neutral to the room, and typically quieter as well. (See the Solution Brief on this topic listed at the end of this paper to learn about the successful implementation of this cooling strategy in an Intel facility.) The potential downside is, in the rare event of a failure in the cooling system, the rack would need to be manually opened to the room to preclude overheating.

Another related rack-cooling strategy is in-row liquid coolers and overhead liquid coolers. In both cases, these solutions act similarly to the liquid-cooled door. Local room airflow exists (potentially augmented by the added cooling unit), but the heat is removed very near the heat load and does not stress the room-level air-cooling system. In-row liquid coolers and overhead liquid coolers can take up valuable rack or overhead real estate. Consequently, this must figure into any analysis of their use. They do have the benefit of being rack-independent. In other words, their use is not tied to a particular server or rack manufacturer.

Another potential downside of these solutions is the risk of fluids close to the IT equipment. The fluid may evaporate or be inert, but in all cases the risk of leaks creates a potential problem that must be weighed within the overall decision process. Choosing one of these solutions often carries with it the cost penalty of a leak detection system.

Liquid Cooling Efficiency

There are many claims that the efficiency of liquid cooling is the answer to solving the data center power problem. Again, it is important to ask what exactly is meant by liquid cooling. If it implies liquid all the way from the CPU to the cooling tower, then yes, the efficiency of this sub-system is very high. Unfortunately, there is much more in the platform to cool than the CPU, so air cooling of the remainder of the server components is still a requirement.

If we consider liquid cooling as we have in this document as an air-to-liquid heat exchange process in or near the rack, then the "liquid cooling" efficiency argument does not hold up. Whether the air-to-liquid heat exchange happens there or at the CRAC or cooling coil, they both have that additional inefficiency of the additional heat transfer step. Moving that heat exchange closer to the load has some advantages, but also some energy penalties and TCO trade-offs. We discuss this in the context of a new construction study on liquid cooling, later in this paper in the section titled Intel Studies in Efficiency. Suffice it to say that liquid cooling is not a silver bullet in the data center energy challenge.

Air Cooling

On the other hand, if we need to cool parts of the platform with air, and we can cool the entire platform with air, perhaps the most efficient thing to do is to skip the liquid step altogether and exhaust this hot air directly outdoors using an air-side economizer cooling system. Unfortunately, these are not simple to retrofit and in many cases do not have a positive return on investment (ROI) in the retrofit arena. On the other hand, they can be very efficient and provide a low TCO if designed into the data center initially in a location that maximizes hours per year of economizer use.

For our purposes here, air cooling occurs any time the prime cooling medium used from the server to outside the rack is air. Generally, air will convey server heat to perimeter CRAC units or to air handlers with cooling coils. There may be a number of these CRAC units or air handlers serving an entire room or large zones, but they are generally associated with facility-level cooling rather than at the rack level.

In-row coolers and overhead coolers are essentially a hybrid of these solutions. In these cases, air is the heat removal fluid from the rack, but liquid takes over to remove heat out of the room. For the sake of these discussions, these hybrid solutions are considered liquid cooling as they most closely resemble liquid-cooled rack doors.

Limits of Air Cooling

The question of the limits of air cooling always generates a healthy debate, but it must be broken down to a theoretical limit and a practical limit. The practical limit is constrained by the existing design of a particular data center. If the space has a limited raised floor height with significant underfloor obstructions and limited fan capacity, these constraints will obviously set an upper bound on the air cooling. Data center operators need to determine this upper bound for themselves based on their specific operational experience, as well as direct measurement. The key to determining this value is to compare a direct measurement of local room airflow with the "required for cooling" vendor-supplied airflow rates. Unfortunately, most data centers do not have real-time airflow values and these may need to be obtained manually. As for the servers, the majority of server suppliers are now providing this required airflow data in accordance with the ASHRAE thermal report.[5] Far too often "industry tribal knowledge" is relied upon to establish the practical limit of air cooling and, as pointed out, it is different for every data center. Data center owners do themselves a disservice by not doing their own work to determine their specific limits.

The theoretical limit for air cooling is based more upon the server than the room. The server supplier builds the thermal management system to cool the CPU and all components in the server with the airflow pulled through the server by the server's fans. Therefore, if the room can meet the airflow need of the server without recirculation in the data center, the limit is based on the server and not the room. Currently Intel is successfully air-cooling racks of 30 kW density and expects to be able to do the same through 2014 and the 37 kW density as shown in Figure 1.

Airflow Segregation

The ability to air cool up to 30 kW successfully is fully dependent on the airflow segregation scheme. The first and most often used technique is as simple as hot aisle/cold aisle segregation. Configuring racks in this simple way where the inlets of the racks are in a common colder aisle supplied by the cooling system and the exhausts blow into a common hot-aisle return that goes back to the cooling system. Simple adherence to this strategy, as well as employing other best practices – such as 1) computational fluid

To "Raised Floor" or Not to "Raised Floor"

One design consideration for expansions into existing buildings and certainly in a "greenfield project" is whether to have a raised floor. The decision should be made only after weighing many factors, but for new projects, the decision need not be made on a perceived requirement for cooling. A review of figures 3 and 4 reveals that neither of these installations had a raised floor. They are not a requirement for cooling in all cases. In high-density applications, proper airflow segregation techniques may render the raised floor obsolete.

Pluses

+ Flexibility for routing liquids, power, and networking under raised floor for future needs
+ Shared supply air plenum for redundancy
+ Many options for grate opening percentage for regulating airflow in front of server racks
+ Grounding of racks

Minuses

- Costly to replace if new racks exceed floor rating
- Potential risk of rack falling through the floor
- Sometimes difficult to balance airflow
- Increased fan power to maintain positive pressure in supply plenum
- Increased capital cost
- Floor height increases with airflow requirements

dynamics (CFD) modeling to verify no recirculation; 2) proper selection of perforated tiles, blanking plates in racks, and 3) plugging holes, gaps, and electrical cut-outs—should be sufficient for 6-8 kW racks. Again, this is predicated on the room being able to deliver the requisite airflow required by the servers. Depending on the servers, the flow rate for the entire rack should probably be in the 800-to-1200 cubic feet per minute (CFM) range.

Air cooling beyond this is still reasonably straightforward, but does require additional cost and complication. As before the key is airflow segregation, keeping the cool supply air from mixing with the warm return air, and then delivering the requisite volume of airflow. Above the ~12 kW/rack range, this segregation must include physical barriers (Plexiglas*, low-cost/nonstructural walls and other architectural features) between the cold aisle and hot aisle. Intel has successfully applied this concept with cold aisle containment (cold aisle becomes an enclosed space), hot aisle containment (hot aisle becomes an enclosed space) and hot-side exhausts or chimney cabinets (where the back side of the server or cabinet is ducted directly to the return). All three work. Cold aisle containment has some advantages in retrofit situations with raised floors. Hot-aisle containment or chimney cabinets generally are more forgiving and efficient in a greenfield design. Each has further advantages and disadvantages, but the ultimate decision should be made based on the specifics of the data center and building under consideration. (See the white paper on air-cooled data centers that is listed at the end of this paper to learn more.)

Density-Driven TCO Has an Upper Bound

Intel does believe that there is an upper bound for what makes sense for density. While the final analysis is not yet complete, we believe that the additional specialized cooling to take racks, either liquid- or water-cooled, much beyond today's 30 kW or 2014's 37 kW may actually add undue cost to a project, while simply spreading the load out in 20-30 kW racks to maintain a more manageable density may provide the lowest TCO. Of course, this analysis will also depend on the cost of space. There is a different tipping point for Manhattan than for a sparsely populated location in central Washington state.

There may be reasons to go beyond these density levels, particularly in the high performance computing (HPC) space where density can drive performance based on node interconnectivity, but these are special cases outside the scope of this work. For these HPC top-end racks, we expect we may see 60 kW and above.

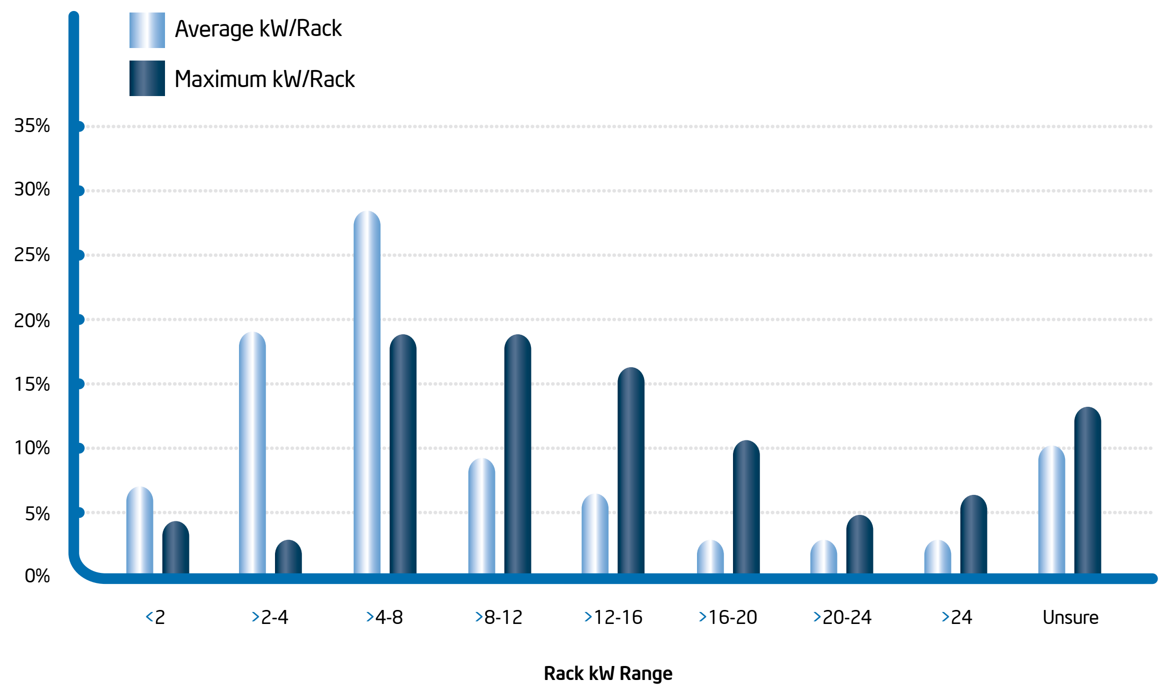

Figure 2: Frequency of occurrence of rack power levels for average rack and maximum rack densities.[6]

Legacy Data Center Discussions

The already cited Liebert Data Center Users group meeting survey provides excellent insight into where today's data centers are with respect to density.[6] A total of 107 businesses responded to a wide range of issues. More than half of the users were in the 2-4 and 4-8 kW/rack density ranges.

Imagine if 2-8 kW is the average rack power draw for your data center and you learn the new full rack of blade servers will be roughly 30 kW. Would you have some initial hesitation on whether it could be made to work?

The first step in finding out is a detailed analysis of ways you might handle the additional load. For instance, you might be able to handle it with an extension of the existing air cooling design. Would your central cooling plant be able to handle the additional load? If the answer to this question is yes, then you need to determine if the room can handle the required airflow for the rack. You should also consider whether you've applied all the best practices possible for efficient cooling in your data center and whether a more rigorous airflow segregation scheme might be a viable solution. For example, for a few high-density racks in a room, direct exhaust ducting (chimney cabinets) might be the ideal solution.

If all those options are exhausted and found wanting, localized cooling and liquid racks could be considered. Intel recently went through this exact scenario and determined that for Intel's specific situation, enclosed liquid-cooled racks were the best solution. (See the Solution Brief on this topic listed at the end of this paper to learn about the successful implementation of this cooling strategy in an Intel facility.)

Case Study of an Expansion in an Existing Facility

Frequently the reuse of existing industrial or commercial space as data center space is considered. In these instances the building invariably comes with design aspects that cannot be modified without significant costs. In these cases, a detailed study of the building's capabilities will lead to the most appropriate cooling solution.

An interesting installation of two new high-density systems at Intel serves to illustrate this point. Intel had excess industrial space and the IT organization had space needs. Consequently, the unused space was reconfigured as IT equipment or data center space. Figure 3 shows a new supercomputer being hosted in one of Intel's buildings in Rio Rancho, New Mexico (NMSC). This computer, named "Encanto" is ranked as the third fastest in the world on the Top500 list (Top500, 2007).[7] Each rack draws approximately 30 kW. Figure 4 shows an engineering compute cluster of roughly the same density racks.

The cooling systems on these two systems could not be more different. The NMSC is liquid-cooled through liquid in a rack rear door. The engineering compute servers are air-cooled with direct exhaust out of the back of the racks (these are often called chimney cabinets). One interesting common feature is that neither has a raised floor.

Figure 3: NMSC, the world's third fastest supercomputer hosted by Intel in Rio Rancho, NM.

Figure 4: Engineering compute resources in Rio Rancho, NM.

Both work very well and show good efficiency. Perhaps what's most surprising is that they are in the same building, roughly 50 meters apart. The section of the building with the NMSC had no air handling equipment but an extensive process cooling water system. The section of the building with the engineering compute system had extensive air handlers but no cooling water near the racks. Together, these two implementations show that neither water nor air cooling is always best in retrofit situations. Instead a thorough engineering analysis of the building's capabilities will lead to the correct choice.

Greenfield Solutions

Starting fresh with bare ground is always one of the best situations if the goal is to optimize efficiency for the data center. Intel's current direction for cooling any new data center is air cooling. We continue to evaluate air- and water-side economizers, as these features in a greenfield site are, at this point, almost a mandatory feature from an economic and often energy code perspective. These data centers will be able to handle 30 kW per rack and more as we move forward.

Intel Studies in Efficiency

Intel believes that high-density data centers are the most efficient and that the major cost of any cooling system is the central plant—for both capital and expense costs.

Intel has recently completed two different internal studies on efficiency. Both studies came to the same conclusion for large, new data centers: Air cooling with hot aisle/cold aisle segregation and room-level air handling units costs less to build and costs less to operate. The cost to move the airflow (per kW of IT equipment) is particularly less with large, open airflow paths, and in the case of raised floors, especially with floor tiles providing 50 percent or more open area. These features can greatly reduce the energy used in moving room air. We have found that liquid-cooling systems often have numerous small fans and/or additional pumping power requirements that make them actually consume more power than a well designed, high-density air-cooling system. In addition, liquid cooling brought to the rack through local cooling hardware is often capital cost intensive and in both our internal studies drove the cost of these systems higher than the equivalent capacity air-cooling systems.

Homogeneity versus Heterogeneity

The one caveat to the above studies was that they applied to largely homogenous data centers where the kW/rack was going to be essentially the same. Examples of this include a new virtualized data center or scientific computing. In those cases a simple cooling system design could be applied across the entire data center. Often this is not the case and the data center may contain a diverse load from a kW/rack perspective. In these cases, localized liquid cooling may be a useful tool in the design of the data center. One can easily imagine an air-cooled data center of 8 kW racks with a few 30 kW racks inside liquid-cooled cabinets. The base 8 kW design can be readily handled by a standard air-cooled design and the hybrid solution of liquid cooling used only where needed.

Conclusion

Should liquid cooling be a requirement for high density racks? In a word, no. But is liquid cooling sometimes a good idea? Absolutely. Particularly in an existing data center with a cooling capacity limited by the existing infrastructure. However, in all cases, a detailed engineering analysis should be completed that particularly examines the possibility to extend air cooling.

That said, we discussed a case study where, based on the existing local conditions, liquid cooling was the ideal choice and represented the lowest cost and quickest installation. What's more, liquid cooling may be applicable in a greenfield data center design in the case where there is a wide range of power densities per rack.

So what's a CIO to do? First think of good engineering as an investment. Each data center project will be different and only by evaluating what you have to work with, along with understanding the constraints involved, will you achieve the optimum high efficiency and low TCO solution you seek.

kW/rack and $/kW Better Metrics than Watts/Square Foot or Dollars/Square Foot

- Watts per square foot can be a misleading and confusing metric. Is it square feet of the entire campus? The sum of the raised floor area? The pitch of the racks? Does it include the utility (power and cooling) space? Because of all these uncertainties, it is often better to discuss densities in kW/rack, particularly when discussing cooling technologies.
- Consider a low-density space (~75 W/sf). This space could certainly have a new rack of top-of-the-line IT equipment drawing 30 kW and the W/sf may not change significantly. However, the cooling system that was ideal for the low-density racks will likely fail miserably cooling the new high-density rack.
- Similarly, dollar per square foot is often a misleading metric. Consider that 60-75 percent of the costs in a new data center are the power and cooling infrastructure which are largely independent of square footage. What's more, the chosen density of the data center can significantly alter the dollar/sf, but will have a much smaller effect on the total capital dollars needed to build the data center. A far better metric is $/kW (kW = power provisioned to the rack). This metric allows better comparison, both design-to-design and data center-to-data center.

Here are some guidelines for the various scenarios we've considered.

Existing extension

1) Understand the limits of the existing data center
2) Implement all best practices
3) Consider enhanced airflow segregation
4) Add liquid/local cooling

Major retrofit or existing building conversion

1) Understand the baseline capabilities of the building
2) Perform detailed engineering analysis and TCO evaluation of various options, air and liquid

Greenfield site

1) Understand rack power values peak and homogeneity
2) Plan for high density if possible
3) Apply air- or water-side economizers
4) Implement air cooling where possible for lowest cost to build and lowest cost to operate

The most important thing is to recognize that there is no single answer and a detailed engineering energy analysis and TCO analysis must be undertaken for each decision.

Note: Specifically excluded from this paper was technology for cooling servers in containers. These have recently become more prevalent and it is a rapidly changing market. The complexity of the topic and the pace of development preclude us from giving it adequate treatment in this work. However, the bottom line is the same: A detailed study of the efficiencies and capabilities of the configuration, combined with a detailed TCO analysis will lead to the right decision.

Acknowledgements

The authors would like to acknowledge the support, review, and input provided by Robin Steinbrecher and Charles W. Rego (Digital Enterprise Group), Michael Meakins (Corporate Services), and Doug Garday (IT@Intel).

Learn More

Solution Brief: Expanding Data Center Capacity with Water-Cooled Cabinets ***(www.intel.com/it/pdf/expanding-data-center-capacity-with-water-cooled-cabinets.pdf)***

White Paper: Reducing Data Center Energy Consumption with Wet Side Economizers ***(www.intel.com/it/pdf/reducing-dc-energy-consumption-with-wet-side-economizers.pdf)***

White Paper: Air-Cooled High-Performance Data Centers: Case Studies and Best Methods ***(www.intel.com/it/pdf/air-cooled-data-centers.pdf)***

IT@Intel Brief: Energy-Efficient Performance for the Data Center ***(www.intel.com/it/pdf/energy-efficient-perf-for-the-data-center.pdf)***

White Paper: Data center TCO: A Comparison of High-Density and Low-Density Spaces ***(www.intel.com/technology/eep/datacenter.pdf)***

White Paper: Increasing Data Center Density While Driving Down Power and Cooling Costs ***(www.intel.com/business/bss/infrastructure/enterprise/power_thermal.pdf)***

White Paper: Turning Challenges into Opportunities in the Data Center ***(www.intel.com/technology/eep/data-center-efficiency/dcwhpaper.pdf)***

Visit The Green Grid: ***www.thegreengrid.org***

Visit Climate Savers Computing Initiative: ***www.climatesaverscomputing.org/***

Visit ASHRAE: ***http://tc99.ashraetcs.org/***

Authors

Michael K. Patterson is a power/thermal architect with the Eco-Technology Program Office at Intel.

Dave Fenwick is the chief server architect in Digital Enterprise Group Architecture and Planning at Intel.

www.intel.com

1. "Emerson Network Power Survey Shows Data Center Energy Efficiency Initiatives are gaining Ground Despite Challenges," press release, Emerson Network Power, September 20, 2007.
2. Patterson, M.K., Costello, D., Grimm P. Loeffler, M. (2007) "Data Center TCO; A Comparison of High-density and Low-density Spaces," THERMES 2007.
3. Datacom Equipment Power Trends and Cooling Applications, ASHRAE, 2005.
4. Ibid.
5. Thermal Guidelines for Data Processing Environments, ASHRAE, 2004.
6. "Emerson Network Power Survey Shows Data Center Energy Efficiency Initiatives are Gaining Ground Despite Challenges," press release, Emerson Network Power, September 20, 2007.
7. This computer, named "Encanto" is ranked as the third fastest in the world on the Top500 list (Top500, 2007): http://www.top500.org/list/2007/11/100.

 0408/LB/HBD/PDF 319808-001US