**InterPACK '07**
**July 8-12, 2007**
**Vancouver, BC, CA**

# Paper InterPACK-1176

# ANALYSIS OF ALTERNATIVE DATA CENTER COOLING APPROACHES

**Dr. Robert Hannemann**
Thermal Form & Function Inc.
Manchester. MA USA

**Herman Chu**
Cisco Systems, Inc.
San Jose, CA USA

## ABSTRACT

The average equipment rack in most of today's high-performance data centers is limited to 1-3 kW with a typical overall heat load density of less than 100 W/ft$^2$. Near-future racks, however, will dissipate up to 15 kW; in 2-4 years, computer and communications rack heat loads are projected to balloon to 30 kW with heat load densities exceeding 500W/ft$^2$. Handling these heat loads is becoming increasingly difficult and expensive using traditional rack and data center cooling approaches.

Based on an analysis of a realistic data center expansion plan, the current paper compares capital and operating costs associated with three alternative cooling approaches: (1) a business-as-usual approach, (2) employment of cooling augmentation systems based on chilled water and refrigerant-based heat exchangers, and (3) deployment of water- and refrigerant-based device-level cooling for some of the heat load. A major conclusion of the work is that challenging current industry norms can result in significant energy savings while allowing the benefits associated with increased functional density.

## INTRODUCTION

Electronics cooling continues to be significantly important for the computer, networking, and telecommunications industries as a limiting factor in equipment performance and customer acceptance for new products. While device-level cooling (for example, for high performance microprocessors) has been a topic of interest and innovation for some time, data-center-level cooling has emerged as an increasingly critical technical, economic, and legislative policy issue [1]. From the end-user perspective, improvements in device-level performance and integration now allow unprecedented levels of functional integration. Cooling issues, however, present a stumbling block to achieving the equipment densities allowed by racked servers and communications gear. Furthermore, the skyrocketing electric power requirements for such gear present a growing operating expense (OpEx) burden – and cooling systems represent 35 – 50% of the total electric power required. The business imperatives for IT and operating managers are (a) "densification" and consolidation of functions to limit the need for new data center facilities construction, and (b) lower OpEx.

The average equipment rack in most of today's high-performance data centers is limited to 1 – 3 kW with a typical overall heat load density of less than 100 W/ft$^2$. Near-future racks, however, will dissipate up to 15 kW; in 2 – 4 years, computer and communications rack heat loads are projected to balloon to 30 kW with heat load densities exceeding 500 W/ft$^2$. In order to handle these cooling loads, it is becoming increasingly difficult and expensive to continue to employ the traditional rack and data center cooling approach – the circulation of ever-greater amounts of conditioned air through racks using a raised floor air distribution system and computer room air conditioners to remove the heat from the circulating air. Several approaches to augment this business-as-usual approach, such as the deployment of liquid-cooled rack heat exchangers or refrigerant-based auxiliary cooling, have recently been introduced.

Based on an analysis of a realistic data center expansion plan, the current paper compares capital and operating costs associated with three alternative cooling approaches: (1) a business-as-usual approach, (2) employment of cooling augmentation systems based on chilled water and refrigerant-based heat exchangers, and (3) deployment of water- and refrigerant-based device cooling for some of the heat load.

## SCENARIO DESCRIPTION

In order to bound the problem in a meaningful way, we posed the following scenario: given a baseline data center in operation in 2008 with 15 kW per rack, what would the impact be if the rack heat load were to increase to 30 kW per rack? We assumed a “model data center” based on an existing equipment room, upgraded it to a plausible 15 kW per rack capability cooled with conventional technology, then analyzed various approaches to cooling the data center if each rack were raised to 30 kW. This would mimic the experience of a customer who has a current data center (generation “n”), upgrades to 15 kW racks (generation “n+1”), and finally attempts to configure the same data center with all racks at 30 kW (generation “n+2”).

The assumed data center parameters are shown in Table 1. Floor areas and airflow requirements were specified using leading-edge rules of thumb: 1800 CFM of airflow per 15 kW heat load in both generation “n+1” and generation “n+2” and 500 CFM airflow per standard perforated floor tile.

| | Model Data Center - Generation "n" | Baseline Data Center - Generation "n+1" | Advanced Data Center Generation "n+2" |
|---|---|---|---|
| Floor area (ft$^2$) | 8,687 | 12,000 | 12,000 |
| Racks | 239 | 240 | 240 |
| Rack power (kW) | 3.4 | 15.0 | 30.0 |
| W/ft$^2$ | 93.2 | 300.0 | 600.0 |
| Rack airflow (CFM) | 360 | 1800 | 3600 |
| CRAC units | 9 | 40 | 80 |
| Total airflow (CFM) | 86,040 | 432,000 | 864,000 |
| Perf tiles / rack | 1 | 4 | 8 |
| Area / rack (ft$^2$) | 8 | 20 | 36 |
| Total rack area (ft$^2$) | 1,912 | 4,800 | 8,640 |
| CRAC floor area (ft$^2$) | 450 | 2,000 | 4,000 |
| Aisle / other area (ft$^2$) | 6,325 | 5,200 | -640 |

**Table 1: Model data center parameters**

The leftmost column represents the existing equipment room. In this case, the average power per rack was estimated using the fact that 9 commercial Computer Room Air Conditioner (CRAC) units were located in the room. Given the capacity of this equipment (90 kW per CRAC unit exclusive of the CRAC power draw), the racks were estimated to dissipate 3.4 kW each. This is quite consistent as a “typical” data center in 2005; 90% of existing data centers in that year averaged between 1.5 and 4.5 kW/rack [2]. (Note that no provision for redundancy was assumed for this overview analysis.)

The so-called “baseline” data center could reasonably be provisioned in 12,000 ft$^2$, as opposed to the 8687 ft$^2$ for the “model” data center. In fact, the distribution of area between equipment, cooling, and open space approaches that of the canonical data center used in some prominent studies [3].

The key question, then, is whether or not alternative cooling approaches would allow the “advanced” (30 kW rack) data center to fit in the same space. Note that the right-hand column represents floor areas for a traditional air-cooling approach. Clearly, there is an inconsistency: more floor area is needed for racks and CRAC units than is available (ignoring aisle space). This is an illustration of the basic problem: in such a situation, the customer would need to add data center floor space through an expensive construction effort.

The relative floor area requirements are shown schematically in Figure 1 below.

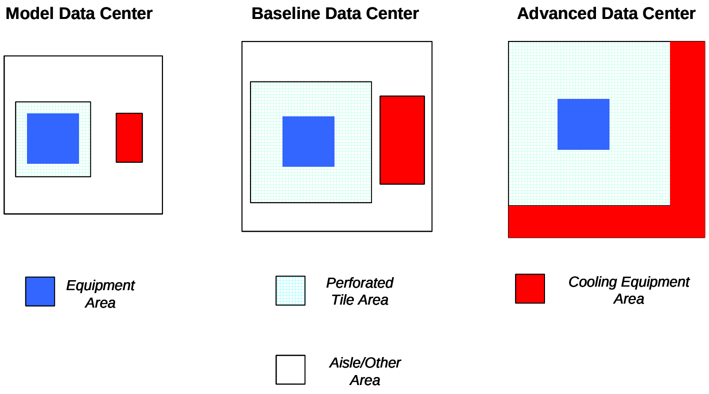

**Figure 1: Model, baseline, and advanced data center floor space allocation**

Finally, in moving from “baseline” to “advanced” racks, it was further assumed that:

- 15 kW of the heat load is due to 160 high-power devices (microprocessors or network processors) at 93.8 W each;
- 15 kW of the load is due to distributed heat dissipation requiring typical forced air cooling.

## SCOPE

Five cooling approaches were considered as part of the study:

1. Standard air cooling, as the baseline approach;
2. Water augmentation via rack heat exchangers;
3. R134a augmentation via rack heat exchangers;
4. Water “touch” cooling (individual water cold plates for the high-power devices);
5. R134a “touch” cooling (individual evaporating-refrigerant cold plates for the high-power devices.

Additional details for each approach are provided in the next section of this paper.

A significant number of combinations and permutations of cooling equipment could be employed. To simplify the analysis, we assumed that chilled water was not available a priori in the building in question and that the ultimate heat rejection would occur on the roof of the building. (Note, however, that the water-based approaches do require the addition of a chilled water unit in utility space outside the clean room.)

An exhaustive optimization exercise for the designs was not within the scope of the study. Also, while two geographic locations were studied (Minnesota and Texas) to gauge the effects of climate, complete modeling of yearly electricity use and costs were not performed given the complexity of taking into account daily, monthly, and yearly average temperature and humidity variations. Thus, the electricity use in the conclusions of this report represent peak loads, and associated costs will be overstated. We believe that the comparative results, however, accurately represent the relative performance of the alternatives and could be converted into estimates of actual use and cost using ASHRAE climate data.

## ALTERNATIVE COOLING APPROACHES

### Standard Air Cooling

In this approach, air is circulated through the underfloor space, exits through perforated tiles, is drawn through the equipment racks, flows to a number of CRAC units, where it is cooled and appropriately humidified, and returns to the underfloor plenum This approach is illustrated schematically in Figure 2.

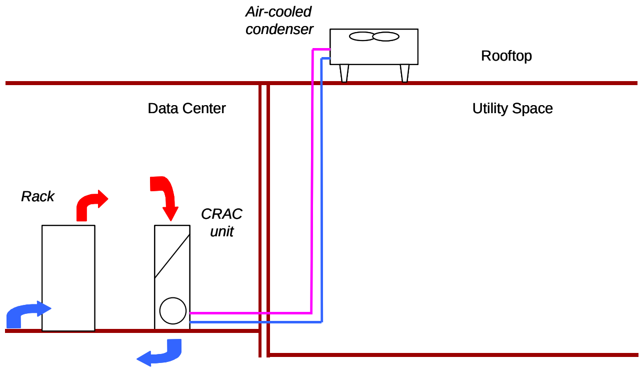

**Figure 2: Standard air cooling**

Note that we assume compressor-based CRAC units with heat rejection on the rooftop via air-cooled condenser units. In order to estimate capital costs, equipment floor space requirements, and power use, we selected appropriate products from the dominant data center air conditioner supplier. In this case, the CRAC units selected were commonly used commercial units, as were the condensers (fan-cooled rooftop units).

It should further be noted that the use of "standard" air-cooling for the 30 kW case will likely result in (a) capital expenditures to upgrade the raised floor with greater plenum height (30" – 36" or more) and (b) acoustic noise emissions may become dominant and unacceptable. These issues are beyond the scope of this overview study.

### Water Augmentation Cooling

For this approach, 15kW of the 30 kW load is assumed to be handled by existing CRAC units (as in Figure 2 above). The other 15 kW in each rack is assumed to be removed by water-cooled door heat exchangers, as in Figure 3.

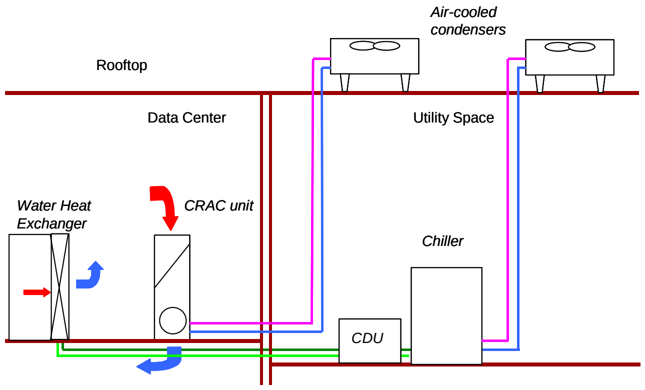

The water-cooled rack door heat exchangers were modeled on a recently announced commercial product, for which detailed performance data were available.

**Figure3: Water augmentation cooling**

CRAC units and rooftop condensers were identical to those used for the previous case; the water chiller assumed was a commercially available 37-ton product. (Note that this unit has an integrated CDU – coolant distribution and pumping unit).

### Refrigerant-based Augmentation

As previously noted, this approach uses a pumped refrigerant (typically R134a) that evaporates in the heat absorbing heat exchanger and then flows to a heat rejection location. Because the exit air temperature is constrained to be at or close to the desired room air temperature, the primary refrigerant loop is cooled via a vapor compression refrigeration device in the utility space and is then pumped back to the door heat exchangers. The working fluid (again, a refrigerant, likely R134a) condenses at the ultimate heat rejection location, the rooftop, as in Figure 4.

In this case, note the "chiller" shown is an R134a chiller, not a chilled water unit. As in the previous case, 15 kW is handled by existing CRAC units and 15 kW is removed via the door heat exchangers. These heat exchangers were modeled after an internal TF&F design, and the chiller was assumed to be a commercial refrigerant chilling and pumping unit (integrated CDU).

Note that although this approach will provide secondary benefits (such as the absence of chilled water piping within the data center, and lighter and less expensive heat exchangers), because it is a 2-loop system with a vapor compression refrigerator, only marginal differences between this approach and the water augmentation approach were expected.

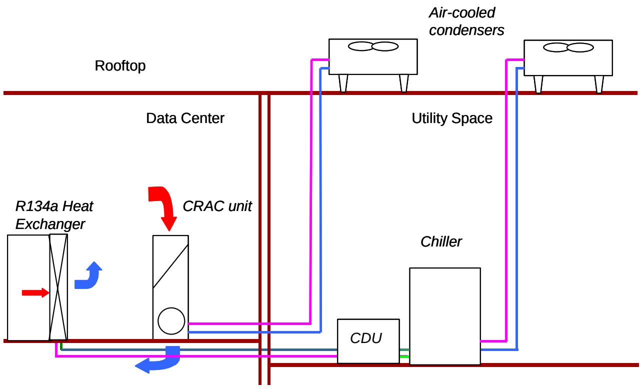

Figure 4: Refrigerant-based augmentation

## Water Touch Cooling

In this approach, 15 kW is removed via airflow and CRAC units, while 15 kW is absorbed by individual device cold plates cooled by chilled water (160 devices per rack at 93.8W each). Because significantly improved thermal resistances (over air heat exchangers) are possible (~0.4 K/W for the device-to-fluid resistance), chilled water temperatures can be fairly high (38 $^{o}$C in this case) while still maintaining all devices below 100 $^{o}$C junction temperatures. In turn, this allows the nominal capacity of the chilled water units to be higher, with fewer units requiring less power.

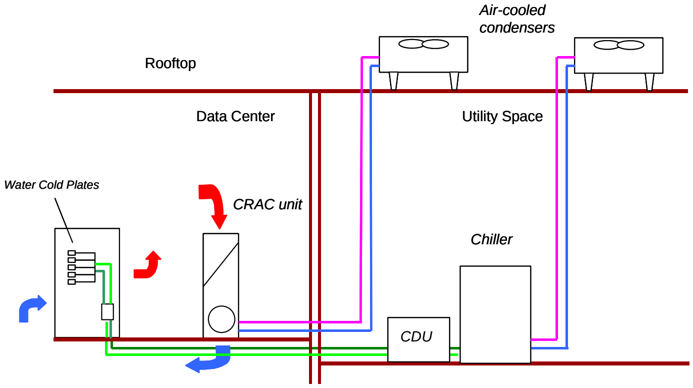

Figure 5: Water touch cooling

In this cooling approach, the chiller was assumed to be the same 37-ton commercial unit as in the water augmentation approach. Rough assumptions were also made regarding the costs associated with the individual cold plates and the rack-level piping and controls.

It should be pointed out that a significant volume of water must be pumped through the rack piping and miniature cold plates – at least 2400 gallons/minute for the ensemble, 10 gpm per rack. The introduction of water into the electronic racks poses significant issues that must be resolved. (For a comparison of water and R134a as heat transfer media, see the Appendix.)

## Refrigerant Touch Cooling

This approach (Figure 6) uses high performance evaporative cold plates to cool the individual high power devices [4]. Unlike the coolant in water touch cooling, the refrigerant can be condensed at the rooftop. Thus, a chiller is not needed, only a CDU. This is possible due to the very low device-to-refrigerant resistance (~0.2 K/W). The saturation temperature in this case was 70 $^{o}$C, which maintains junction temperatures below 100 $^{o}$C even on the hottest day in the Texas location. As in the previous cases, 15 kW was absorbed in the cold plates and 15 kW was assumed to be handled by CRAC units.

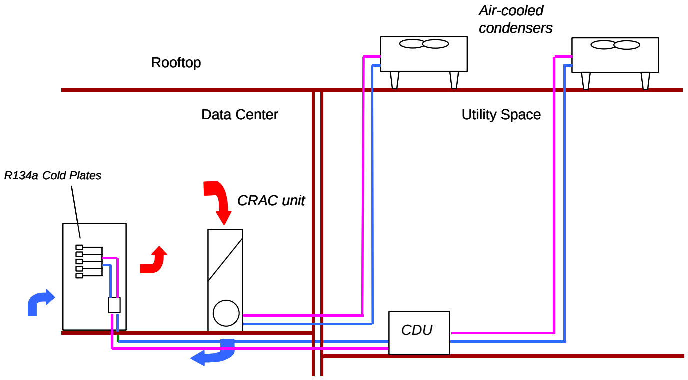

Figure 6: R134a touch cooling

An estimate of the cost of the CDU pumping unit, the cold plates, and the rack piping and controls was made based on internal TF&F designs. The use of a two-phase cooling approach with R134a allows fluid flow rates to be at least a factor of 4 lower than the use of single-phase water.

For each of the cooling options, consistency between equipment choices and required performance – airflows, liquid pumping capacities, and operating temperatures – was checked.

# ANALYSIS RESULTS

## Cooling Equipment Floor Space

A primary driver in the analysis was to determine if the advanced (30 kW rack) data center could be configured in the existing data room's 12,000 ft$^{2}$. As previously noted, this was NOT possible using standard air cooling with internal CRAC units. In the other cases, the major cooling equipment was located either in utility space outside the data room or on the rooftop. Note that each of the three classes of space will have significantly different construction and maintenance costs;

given the highly variable nature of these costs depending on location and climate, no attempt was made to convert the floor areas calculated to corresponding costs.

In each case, real equipment specifications and sizes were used, and access space accounted for. Summary floor space data for the alternatives is shown in Figure 7. This represents floor area for cooling equipment alone.

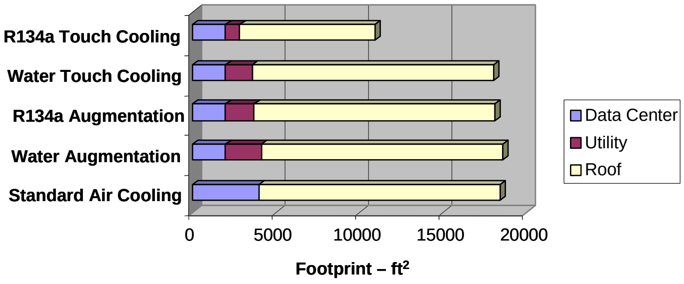

**Figure 7: Summary floor area data for alternatives (Minnesota case)**

With the addition of CRAC units to cool the added 15 kW per rack, the standard approach will require additional data center space to be constructed. The other approaches will not require additional data center space but will require the addition of utility and rooftop space to varying degrees. Refrigerant touch cooling will need about 50% of the next-best utility room space and a significantly lower amount of rooftop space.

## Cooling Power

With the soaring cost of energy, the amount of power needed for cooling is a critical consideration in the deployment of IT equipment. Each of the cooling alternatives in this study was analyzed for power use. In each case, manufacturer's data on power draw was used. Three caveats are important in examining the results for power consumption: first, the difference between nameplate power and actual or typical power was not accounted for; second, it is well known that server workload significantly impacts actual power draw; and perhaps most importantly, as previously noted, the power use was assumed to be for a worst-case weather situation (maximum dry bulb and wet bulb temperatures). The data thus represent the maximum power draw for the cooling equipment – which the power system must be sized for. This design power data is shown in Figure 8.

The cooling power figures in green do NOT account for fans, etc. within the racks, which would push the "cooling" power requirements towards 40% of the total power draw in the facility. The distribution of power between electronics and cooling equipment is similar to that reported for actual, installed data centers [5,6].

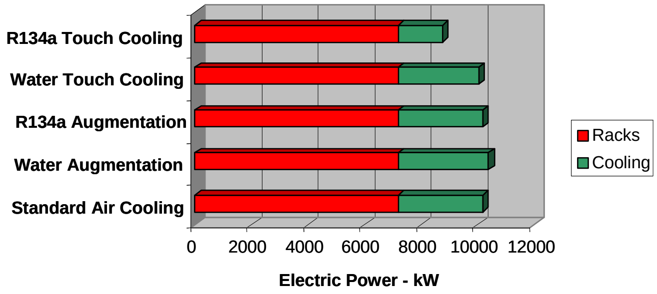

**Figure 8: Power requirements (Minnesota case)**

Except for refrigerant touch cooling, none of the alternatives provide a significant benefit over the standard approach. Note that this alternative uses very little power for the added 15 kW per rack; the cooling power bars include the original 15 kW cooled via conventional CRAC units. Thus the power savings for this alternative are substantial.

## Capital Expenditures

Cooling equipment is a not-insignificant portion of the overall capital expenditure required to equip a data center. For each of the alternatives, the major costs were determined by obtaining quotes from the manufacturer. A 10% adder was assessed for installation, piping, etc. The cost for the electric power supply (which in many cases will require expensive uninterruptible supplies) was not included. Door heat exchangers, cold plates, and rack-level piping and controls were included in the CapEx estimates.

The CapEx results are shown in Figure 9. It is interesting that, except for the refrigerant touch cooling option, the alternatives are all more expensive than the standard air cooling option. The likely reason is that water and refrigerant chillers are newer, lower-volume products than the mainstream CRAC units, as opposed to any inherent thermodynamic inferiority.

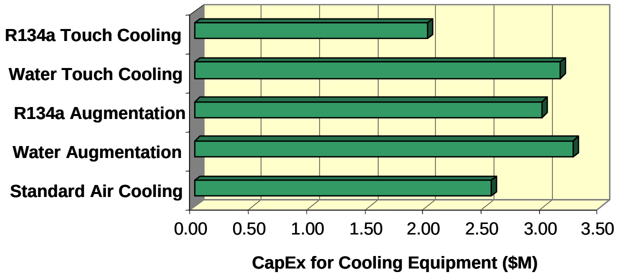

**Figure 9: Capital expenditures for cooling alternatives (Minnesota case)**

Once again, the refrigerant touch cooling approach has a significant advantage over all the alternatives – primarily because the high thermal performance of the evaporative heat exchangers and the ability to run high saturation temperatures in the refrigerant cycle eliminates the need for a refrigerant cooler and a large number of condenser units.

## Operating Expenses

The true operating expenses for a data center must include power, maintenance, routine monitoring, and electricity costs. In this study, our focus is on the cooling alternatives, and electricity costs are the OpEx elements calculated. Given operating data for the equipment and the data room, utility space, and rooftop areas, a better notion of the true operating expenses could be calculated. As has been previously mentioned, a more sophisticated climate and equipment-operating model would be required for accurate absolute OpEx calculations. The power data in Figure 8 above do, however, accurately provide relative performance data for a key element of OpEx.

At $0.10 / kW-hr, a reasonable country-wide electricity cost average, the costs associated with cooling range from $8.9M per year for standard air cooling to a low of $7.7M for refrigerant touch cooling. In relative terms, this approach saves 48% on cooling equipment electricity costs (15% overall electricity savings). This is for yearlong operation at climate extremes.

## Impact of Location

Calculations were performed for all options assuming two potential data center locations, Minnesota and Texas. The data in the figures so far is all for the Minnesota location. Since dry-bulb temperatures in Texas can be significantly higher than in Minnesota, the Texas case may be expected to require more heat transfer equipment. In a practical sense, for this study, the Texas case requires a significantly higher number of rooftop condenser units.

Figure 10 illustrates the impact on the key parameters for the R134a touch cooling case.

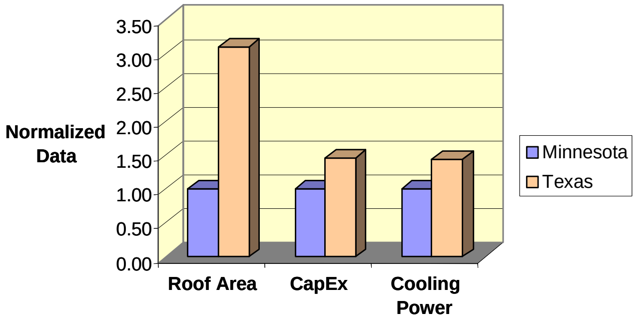

**Figure 10: The impact of data center location (R134a touch cooling case)**

# Conclusion

One obvious conclusion from this work is that standard air cooling will be a significant hindrance for customer acceptance of very-high-powered racks. The standard approach, with ever-greater underfloor heights and more powerful air movers will likely continue to be deployed, especially for greenfield installations. There would be no way for the baseline data center to be modified for the advanced (30 kW) racks without construction of significant new space.

The results for each of the alternative cooling schemes are shown in Table 2.

| | Standard Air | Water Augmentation | R134a Augmentation | Water Touch | R134a Touch |
|---|---|---|---|---|---|
| ***Area (cooling eq't , ft$^2$ )*** | | | | | |
| *Data Center* | 4000 | 2000 | 2000 | 2000 | 2000 |
| *Utility* | 0 | 2160 | 1728 | 1620 | 864 |
| *Rooftop* | 14400 | 14400 | 14400 | 14400 | 8065 |
| ***Power (kW)*** | | | | | |
| *Racks* | 7200 | 7200 | 7200 | 7200 | 7200 |
| *Cooling* | 2976 | 3198 | 2976 | 2843 | 1556 |
| ***CapEx ($M)*** | | | | | |
| *Cooling Equipment* | 2.55 | 3.26 | 2.99 | 3.15 | 2.00 |
| ***OpEx ($M/year)*** | | | | | |
| *Electricity* | 8.89 | 9.08 | 8.89 | 8.77 | 7.65 |

**Table 2: Summary results (Minnesota case)**

Each of the "advanced" cooling approaches preserves the existing data room area. The refrigerant touch cooling option uses the least utility space and rooftop area, with cooling power requirements about 50% of the other options. CapEx for this option is about 33% lower. (Note the OpEx figures are for peak power assuming electricity costs of $0.10 / kW-hr.)

These results are actually not surprising. The augmentation solutions ultimately are stopgap measures that allow for increased server / communications density in a given data room area. Thermodynamically, though, they are similar to standard air cooling in that they require vapor compression refrigeration units. Because the augmentation solutions represent the least change to current practice, and do not require the cooperation of the electronics OEMs, they could be expected to be the first solutions implemented – and they are.

Table 3 is a summary of the pros and cons of each approach studied. R134a touch cooling is clearly the superior *technical* option because of the very high performance it has in both absorbing heat and in rejecting heat at the rooftop. Water touch cooling offers some advantages over the augmentation solutions, but it is clear that if OEM cooperation in deploying advanced cooling technology is obtained, going all the way to refrigerant touch cooling is by far the most beneficial approach.

It should be noted that the touch cooling solutions also offer the opportunity for a significant boost in rack density, which has not been analyzed in this study.

| *Cooling Approach* | *Pros* | *Cons* |
|---|---|---|
| **Standard Air Cooling** | ◆ **Known, somewhat comfortable technology** | ◆ **Limits densification and IT consolidation**<br>◆ **High energy cost**<br>◆ **Severe acoustic noise** |
| **Water Augmentation** | ◆ **No increase in data center floor area**<br>◆ **Modest extension to today's cooling approach**<br>◆ **Available, but vendor-specific** | ◆ **Introduces water to the data center**<br>◆ **No energy or CapEx savings** |
| **R134a Augmentation** | ◆ **No increase in data center floor area**<br>◆ **Modest extension to today's cooling approach**<br>◆ **Significantly lower flow rates than water** | ◆ **No significant energy or CapEx savings** |
| **Water Touch Cooling** | ◆ **Some (slight) energy savings**<br>◆ **Lower acoustic noise than standard air** | ◆ **Introduces water into electronic enclosures**<br>◆ **Complex plumbing, high flow rates** |
| **R134a Touch Cooling** | ◆ **Saves significant energy**<br>◆ **Lowest CapEx** | ◆ **New technology (for IT applications)** |

**Table 3: Cooling approach pros and cons**

The work reported on in this paper provides only a high-level analysis of various data center cooling alternatives (as opposed to a detailed design analysis). Other combinations of cooling subsystems may prove to be optimal for particular real-life circumstances. Also, important details have been ignored – such as connection schemes and the number of fluid loops used that touch cooling might require. We do hope that we have provided a context for further exploration in line with our belief that advanced cooling is best viewed from the end customer standpoint – densification and total capital and operating expenses – rather than the conventional "our microprocessors are getting too hot" approach.

## ACKNOWLEDGEMENTS

This study was sponsored by Cisco Systems, Inc. and carried out by Thermal Form & Function, Inc.

## REFERENCES

[1] Mullins, R., "Bush Signs Law to Study Data Center Energy Usage," Network World (IDG), December 2006.

[2] Rasmussen, Neil, APC White Paper #46, Rev. 2005-4, APC, Inc., 2005.

[3] Brill, K., "2005-2010 Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment," Uptime Institute, 2006.

[4] Hannemann, R., J. Marsala, and M. Pitasi, "Thermal Design and Performance of Two-Phase Meso-Scale Heat Exchangers," Paper HT2005-72743, 2005 National Heat Transfer Conference, July 17-22, 2005.

[5] "Power struggle: how IT managers cope with data center power demands," Computerworld, 3 April 2006.

[6] "NY Data Center No. 2: Energy Benchmarking and Case Study," Lawrence Berkeley National Laboratory, July 2003.

## APPENDIX: HEAT TRANSFER FLUID COMPARISON

In considering advanced cooling options, conventional wisdom has gravitated to the use of water as a coolant (for example, mainframes in the bipolar era were often water cooled). This is understandable, given the excellent single-phase heat transfer characteristics of water and its use as a coolant in other applications.

Table 4 below provides a pro-con summary for water vs. R134a.

| *Coolant* | *Pros* | *Cons* |
|---|---|---|
| **Water** | ◆ Known, comfortable technology<br>◆ Good transport properties for single-phase cooling<br>◆ Non-toxic (before chemical additives)<br>◆ Non-flammable<br>◆ Readily available and inexpensive | ◆ Electrically conductibve (leaks can be disastrous)<br>◆ Additives required for freeze protection and microbial control<br>◆ Corrosion is an issue<br>◆ High fluid velocities can cause erosion in micro-heat-exchangers |
| **R134a** | ◆ Non-conductive (no leak issues; gas at room conditions)<br>◆ Widely used as a coolant<br>◆ Superior thermal properties in two-phase operation<br>◆ Non-toxic<br>◆ Non-flammable<br>◆ Inert to most engineering materials<br>◆ Readily available and inexpensive<br>◆ Low flow rates required | ◆ Operates at higher-than-atmospheric pressure<br>◆ Not yet widely used for IT applications (products exist) |

**Table 4: Liquid coolant alternatives**

Water has many positive attributes, including low cost, availability, and non-toxicity (although with required anti-freezing and microbial control additives, it becomes a toxic fluid). There are a number of negatives associated with the water cooling of electronics, including the fact that it is a conductive liquid (so cooling system leaks can be electrically disastrous) and the fact that it is highly corrosive and requires careful engineering of the materials used in system construction. Since water-cooling is single-phase, high flow rates are required (2400 gpm in the case of the advanced data center of this study).

By contrast, R134a is non-conductive and is a vapor at room conditions (Refrigerant systems of the type studied operate at a moderate pressure of ~130 psi; the fluid is in liquid form as it is pumped to the heat exchangers, partially vaporizes to absorb heat, and recondenses at the heat rejection location). It is non-toxic, non-flammable, and fully approved for use as a coolant (it is used in myriad air conditioning applications, including automobiles). As a single-phase coolant, its transport properties are not as good as water, but in two-phase operation provides very high performance heat transfer. Flow rates tend to be a factor of 4 – 8 lower than single-phase water, and pressure drops in the cooling system are significantly lower than for water systems. R134a is compatible with a very wide range of materials.

R134a costs about $8 per pound; overall, the advanced data center of this study would require perhaps $10K worth of coolant. Note that this would be offset by the chemical additive requirements for water systems. Finally, although R134a is not widely used in electronics cooling applications (yet), there is a very solid technology infrastructure to support its use (e.g., supermarkets use R134a-based cooling systems extensively).