Data Center 2020: Delivering high density in the Data Center; efficiently and reliably

Energy consumption at data centers can be reduced using methods that are easy to effect.

Review: Results of the Optimization Phase

The researchers at T-Systems and Intel were able to show reductions in energy consumption in the first optimization phase of the data center. The measures were straightforward and at a reasonable cost. The improved energy efficiency is mainly due to two effects:

1. The airflow was optimized thru a strict separation of hot and cold air using a cold-aisle containment and reduced leakage from under the raised floor. This allows a reduction of fan speed (with an even greater reduction in fan energy) in the room-level air handling units.

2. The increase of supply air temperature to the raised floor (T_1) and cold aisle containment; allowing a simultaneous increase in chilled water temperature reduced the amount of time requiring chiller operation and increased the time the data center could operate in a free-cooling mode. The PUE was optimized by the team while keeping the server inlet air temperature well below the ASHRAE recommended IT inlet temperature (T_R) of 27˚C.

With these measures, the Data Center 2020 researchers were able to reduce PUE from 1.8 to 1.4 and still maintain a computer inlet temperature of 22˚C. Power Usage Effectiveness (PUE) measures the efficiency of the energy used in the data center. It is a ratio of total energy used in the data center to the energy used by the IT equipment.

The following two figures illustrate these changes.

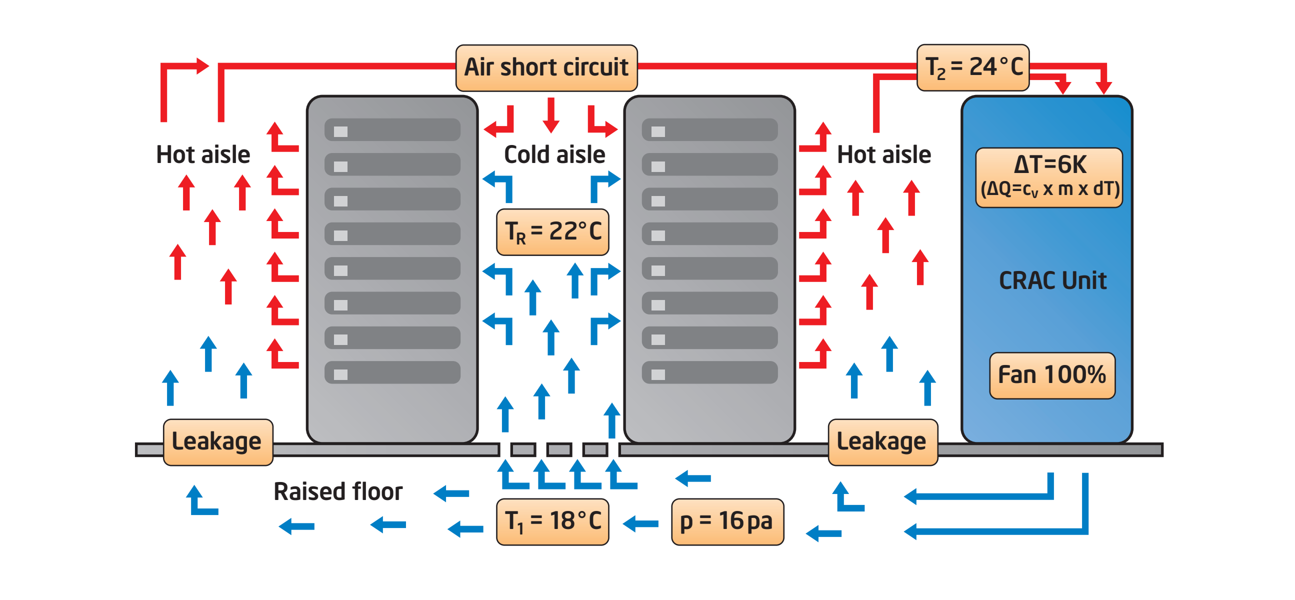

Figure 1 shows the initial temperature values (typical in standard data centers today) with air leakage and no strict separation of cold and hot air. The three key temperatures are:

T_1 = Supply to the raised floor (18˚C in the initial configuration)

T_R = IT inlet temperature (initially 22˚C)

T_2 = CRAC return air temperature (initially 24˚C)

The fan speed in the room cooling unit (CRAC in figure 1) is at 100% and it's ΔT (return minus supply temperature) is at 6˚C.

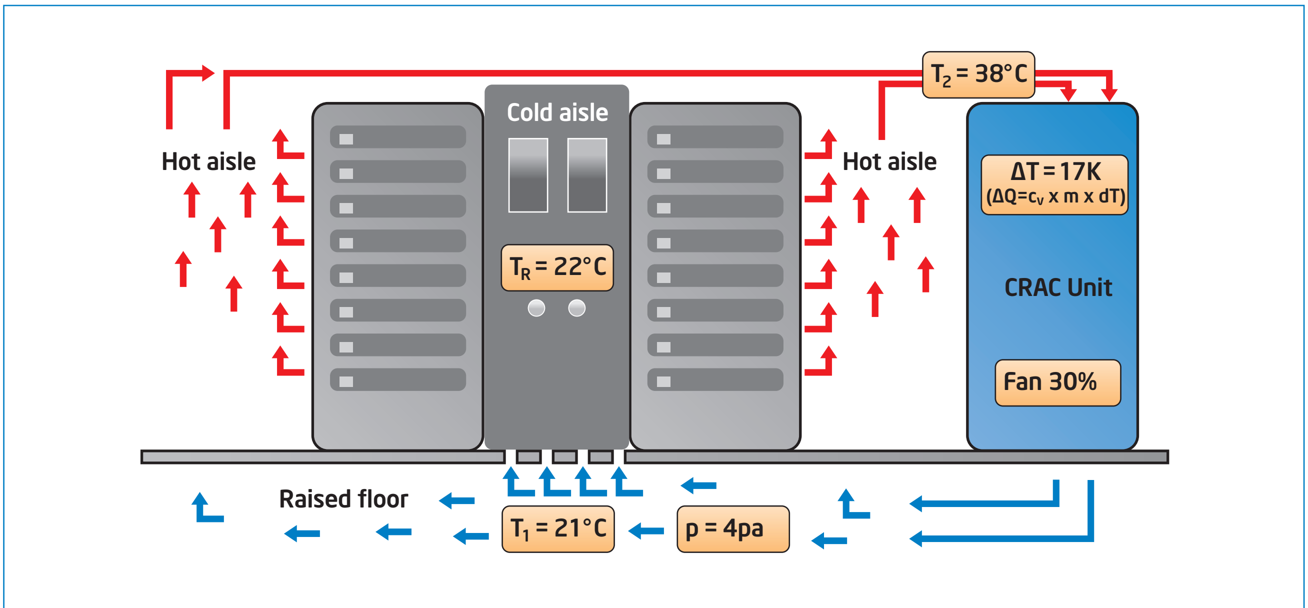

Figure 2 shows impressive improvements when the raised floor leakage is corrected and the cold aisle containment is added. The IT inlet temperature (T_R) remains constant at 22˚C.

The sealing of the raised floor and the clear separation of cold and warm air results in an increase in pressure under the raised floor as well as an increased ΔT. This allows a reduced airflow volume and reduced fan speeds. The reduced airflow absorbs more heat, increasing the return airflow temperature (T_2).

The return air temperature went from 24˚C to 38˚C. In addition, the supply air temperature under the raised floor (T_1) can be increased from 18˚C to 21˚C.

The room air handler (CRAC in Fig 2) consumes far less energy with the fan speed of 30% (initially 100%) After this initial optimization phase the fan motor uses roughly 90% less energy. Combined with the CRAC ΔT being raised from 6 to 17˚C, the fan cooling unit operates much more efficiently than before the changes.

Next Steps: Increasing the energy density to 22 kW/rack

In the first optimization phase, rack density began at 5 kW/rack and was increased to 10 kW/rack. This caused the total data center IT load to go from 40kW to 80kW, doubling the total connected load. This also allowed the PUE to be improved. This increased rack density contributed to the overall PUE improvement of 30%.

To understand the scalability of the data center infrastructure to handle higher rack densities and to verify that the partial load range retained the scalable same level of efficiency, the team increased densities to 22 kW/rack. The IT inlet temperatures (T_R) remained constant at 22˚C.

The team chose two scenarios:

- In the first scenario a single room cooling unit was operated with a chilled water temperature of 8˚C. With the 22kW rack density the PUE was reduced to 1.32.
- In the second scenario two room cooling units were operated with a chilled water temperature of 16˚C and correspondingly reduced fan speed to cool the data center with a rack density of 10kW. Reducing the airflow to half (50%) from each unit will also reduce the energy consumption. From the fan affinity laws, running two fans at ½ speed will only require ¼ of the total power to run one fan at full speed. The warmer chilled water temperature also reduces the use of the chiller, allowing more hours of the indirect free cooling, further reducing the total energy consumption. This allowed the DC2020 team to reduce the PUE even further to 1.23.

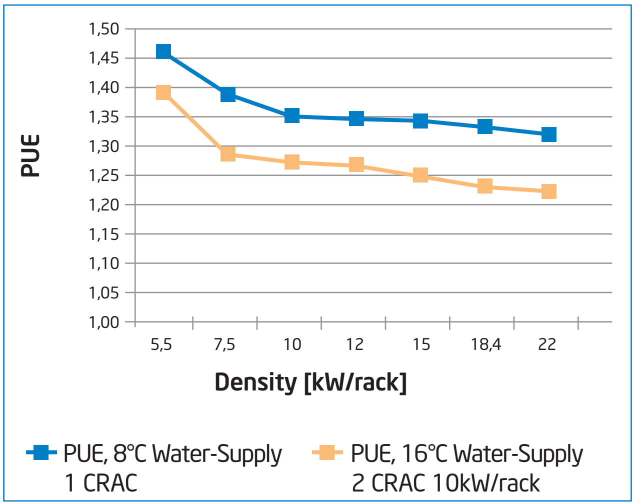

Figure 3

In conclusion, energy densities of 22 kW/rack can be reached with standard techniques. As the density increases the overall result is a flattening of the PUE curve, asymptotically approaching a limit. The measured densities of 22 kW/rack are rarely achieved in reality and only in some specific cases (HPC or high density bladed applications). Therefore, the more common lower thermal loads routinely dealt with in existing data centers will result in slightly higher PUEs.

However, in the construction or design of a new data center high densities and optimum conditions can be built in.

Energy Density and Availability

The optimization of the data center moving towards higher density must also include considerations of high availability and reliability for the server. The data center operator must always balance the benefits of higher density with the risk of a shorter response time in the event of a cooling system failure. For the most critical data centers; energy efficiency must not take away from high availability. To gain information on the optimal balance the DataCenter 2020 test lab was operated in various failure modes (loss of chiller, pumps, and room cooling unit fans). The tests included different energy densities and spatial configurations and the impact of different space temperatures during the failure.

Method: The server inlet temperature (T_R) in the DataCenter 2020 was controlled to 22˚C. Today servers generally have a maximum inlet temperature (T_R) of 35˚C. Below this temperature the manufactures will warranty the reliability of the server. The team then simulated a power failure with the result of the entire cooling system being offline. The servers continued to run on the UPS, releasing heat to the space. As a result the server inlet temperatures (T_R) cycled up.

Generally, the cooling equipment will be on emergency generators which will supply electricity to the data center. After a shutdown chillers take some time to reset and then time for the entire data center to be able to return to appropriate operation temperatures. The red line in Figure 4 shows the critical period (about 5 minutes), after which time it may be expected that servers may fail or shutdown due to overheating (>35˚C inlet temperatures).

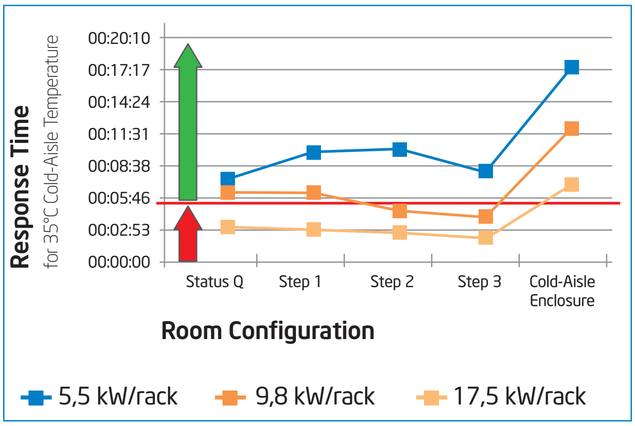

Figure 4

The measurements showed that with reduced leakage and cold aisle containment (to ensure separation of the warm and cold air), the energy density or IT load per rack can be roughly three times that as the standard data center and have the same reliability. In this specific example, the data center operations at 17.5 kW/rack could run longer without the cooling system returning than a standard data center with 5.5 kW racks.

Thus, a relatively inexpensive cold aisle containment supporting the 17.5 kW/rack could replace the need for the comparatively expensive UPS for some cooling outages (obviously the enclosure and UPS are not an even trade-off in all aspects). In this case the enclosure allowed three times the duration more than the hot-aisle/ cold-aisle configuration before the 35˚C limit was reached.

Influence of Ceiling Height

T-Systems and Intel have also discussed in the literature the effect ceiling height may have on room temperatures in the event of cooling system failure. The popular theory is that the higher ceiling height increases the cold air volume, slowing the rate of the temperature increase. This would allow more time before IT system shutdowns due to high temperatures.

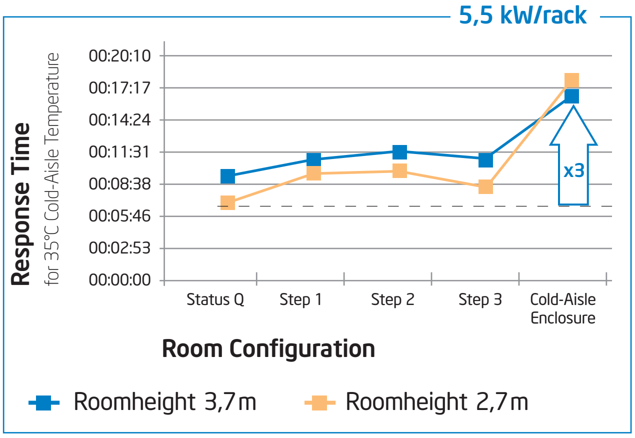

As **Figure 5** shows the higher ceiling had some beneficial affect, but not as significant as was expected. This can only be confirmed as a secod order effect, with a 40% larger volume not providing a correspondingly longer reaction time. At higher rack densities the effect is even less, with the curves closer together. In the case of cold aisle containment, the ceiling height has no significant impact. It remains to be seen what the results are when the DataCenter2020 ir reconfigured as hot-aisle containment.

Conclusion and Outlook

Our second phase of research in the DataCenter 2020 has shown that energy densities of greater than 20 kW/rack can be obtained reliably with available standard air cooling technology. Methodical application of the first White Paper provided a means to achieve significant optimizations. Now, with the cold aisle containment we see the ability to reach higher densities and have a better response in various failure scenarios.

In the Datacenter 2020, the researchers will repeat the above measurements in the next few weeks for a hot-aisle containment to see how the results compare with the cold-aisle containment. Furthermore, they will focus on the IT capability / energy density and its effect on data center total energy consumption. Total energy consumption is just as important as PUE. PUE is not the only indicator of improved energy efficiency in data centers. Interestingly the PUE value will rise when the IT is more efficient in itself, and consumes less power. For this reason the researchers are also considering the IT performance as a measure of efficiency.

0311/KLO/RLC/XX/PDF 325232-001DE