Better Cooling, Hot Savings

Fluid dynamics keeps data center server racks cooler using less electricity.

By Marco Lanfrit, Continental Europe Consulting Manager, ANSYS, Inc.

Heat is the enemy of electronic circuits. So engineers at the ANSYS data center in Otterfing, Germany, cranked up the air conditioning full blast when peak temperatures near the servers rose above acceptable limits. Hot spots remained, however. Engineers used fluid dynamics to identify airflow problems, and the insight it provided led to a quick and inexpensive way to even out room temperatures and significantly improve cooling efficiency.

Hot spots were concentrated near three of six racks that house data center servers and other hardware, such as backup power supplies and data storage units. Peak temperatures averaged 41 C (106 F), far above the 34 C (93 F) temperature needed for optimal operation and dangerously close to 45 C (113 F), at which heat-sensitive electronic components start to fail.

The three hot racks in question house servers used to run the center's solutions for extremely large problems — including highly detailed analyses that exceed customer computing capacity, complex simulations for engineering applications contracted by the ANSYS engineering team, and internal test cases run by ANSYS for physics-related research. These and other complex simulations use high-performance computing (HPC) in which large problems are divided into smaller segments and then solved in parallel, rather than in conventional serial fashion.

Parallelization allows the solution of extremely large problems by using multi-core CPUs that each contain more than one microprocessor. At data centers such as the one in Otterfing, racks of these HPC multi-core CPUs may contain hundreds or thousands of microprocessors, giving them the ability to complete complex simulations in a few hours.

A major related issue is dissipating heat generated by so many microprocessors running simultaneously. Overhead space limitations ruled out using the standard-practice procedure of applying direct cooling through perforated panels on a raised floor.

In search of an alternate solution, engineers created a simulation model representing the major factors affecting thermal distribution in the data center. The room is relatively small, with 32 square meters of floor space and a height of 2.5 meters. Server racks are approximately 2 meters high, leaving about a half-meter clearance to the ceiling. There is approximately 2 meters of clearance to the adjacent side wall at either end of the rack row. High-power fans blow out hot air, approximately 26 KW of thermal energy, from the three HPC cluster racks containing approximately 450 microprocessors. Another 9 KW of energy exits the other three racks, which contain several mid-range conventional servers, a backup server, data storage units, a file server and an uninterruptible power supply. At the rear of the room (behind the racks, an ideal position for sucking in cool air) are two cooling units mounted horizontally near the ceiling.

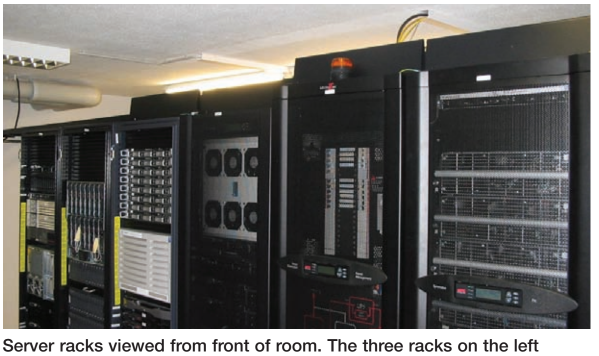

Server racks viewed from front of room. The three racks on the left house HPC servers. Ideally, cool air from wall-mounted air conditioners on the rear wall (behind the racks) gets drawn into the racks, and hot air from the servers is expelled out the front. Racks to the right of the HPC servers contain mid-range conventional servers, backup server, data storage units, file server and an uninterruptible power supply.

Once thermal partitions were installed in the data center, the room temperature could be turned up to save electricity and still maintain proper cooling for the equipment.

Each unit is rated at 12.7 KW cooling with an air-handling capacity of 8,000 m^3/hr — more than enough to sufficiently handle the 35 KW of thermal energy given off by all equipment in the room.

Based on these parameters, fluid dynamics simulation using ANSYS FLUENT software produced heat contours for the walls, indicating a hot spot in the corner where hot air exited the HPC server racks. This hot area continued along the top of the adjacent side wall and was apparent at the air inlet for the HPC server. Cooler temperatures existed along the side wall farthest from the HPC servers. Temperature measurements in these areas agreed closely with the simulation prediction.

A streamline output showing airflow and related temperatures helped to explain the hot spots and uneven temperature distribution in the room. Essentially, hot air from the HPC racks (especially the rack at the end of the row) was blowing toward the front of the room and was sucked over and around the racks directly back into the inlet side of the racks, instead of being drawn into the cooling unit as intended. In addition, hot airflow from the HPC server was diverting airflow from the cooling units, so cool air was only partially reaching the HPC racks, thus further raising the temperature of air being drawn into the HPC server racks by fans in the units.

With this insight, engineers reasoned that adding two sections of thermal partition (one extending from the top of the racks to the ceiling and another from the side of the HPC racks to the adjacent wall) would resolve these issues. The team quickly evaluated four other configurations using the same room model. Though some configurations were marginally more efficient than the one selected, they were ruled out because of greater complexity.

Simulation of the chosen configuration (with the thermal partitions added) showed dramatic improvement in delivering temperature uniformity in the data center room and cooling the racks. According to the analysis, HPC rack peak temperature was lowered from 41 C (106 F) to 34 C (93 F). Additionally, increased cooling efficiency in the range of 10 percent to 20 percent enabled engineers to actually raise average room temperature 8 C (approx 15 F) by adjusting the air conditioning control to a warmer setting.

Temperatures predicted by the simulation proved to be correct when measurements were taken following partition construction. This improvement, the result of adding two simple thermal partitions, translates into an average annual cost savings of 10,000 euros — an outstanding return on an investment for a project that took about two days of work in engineering time and just a few more days to renovate. ■

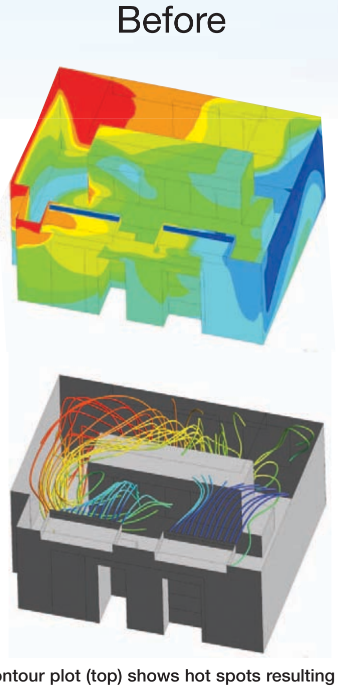

Contour plot (top) shows hot spots resulting from hot airflow traveling over and around the rack units directly back into the rack inlet. This hot airflow pushes cooling air partially away from the HPC racks to further increase air temperature drawn into the racks (bottom).

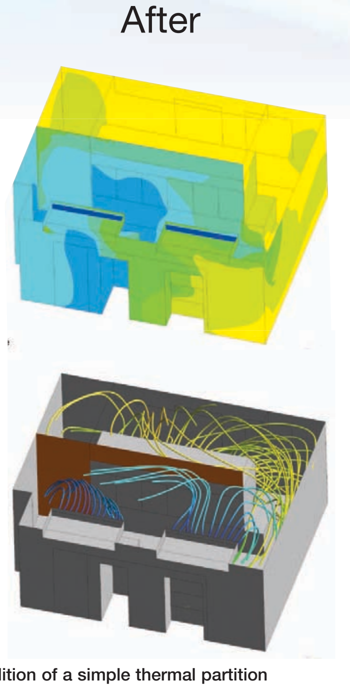

Addition of a simple thermal partition significantly improved temperature uniformity and airflow.

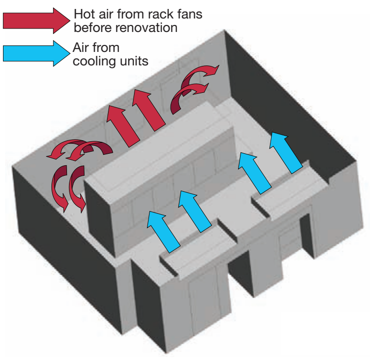

Original data center room configuration

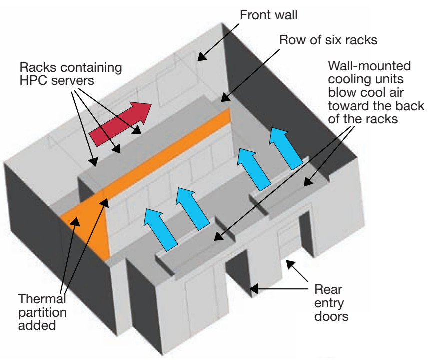

Constructing thermal partitions (highlighted) at the top and one side of the racks was a quick and inexpensive way to modify airflow for correcting the overheating problem.