High-end climate control for a high-end cluster

Taking the heat off servers at the Max Planck Institute

The Max Planck Institute (MPI) for biophysical chemistry in Göttingen analyses the tiniest elements of life. In the process, it also uses imaging techniques such as transmission electron microscopy. High-performance computers use a large number of projections to create 3D models. These improve scientists' understanding of these extremely complex elements. The servers' extreme computing power, together with high packing density and 20 kW heat loss per rack, made air-conditioning a critical factor in the small server room. This is where air/water heat exchangers from Rittal came in.

Proteins are the very elements of life. Movement, metabolism and reproduction – none of these are possible without the aid of proteins. Their forms and functions are as varied as life itself. The creation of a protein from genetic information requires the interaction of a number of different cellular factors, none more important than the ribosome. This universal cell motor interprets the genetic code and establishes the connections between the amino acids. The aim of the MPI's research is to find new ways of producing medicines for treating existing genetic defects. Even chemical laymen will appreciate that these elements are extremely tiny. State-of-the-art high-tech is therefore required to view and catalogue the processes and conditions in this microcosm.

The transmission electron microscope used at the Max Planck Institute in Göttingen provides resolutions of down to 1.8 Á (1 angstrom = 1x10-10 m). In the MPI, 2-dimensional projection images are taken at different angles and – in a similar manner to medical computer tomography – 3-dimensional structures are created from these. In a complex computing process, the images have to be aligned and overlapped at precisely the right location and the image noise has to be removed in order to generate a 3D image on the screen.

Maximising performance while staying cool

"The task of the IT team is to provide the required computing power to ensure that the servers require a minimum of space for their racks and that their performance is not restricted by heat problems," explains physicist Mario Lüttich, IT Manager and member of the 3D Kyro electron microscopy working group. "Our latest cluster produces a 20 kW heat loss in just a single rack. Because of the high packing density of the servers, this heat cannot be dissipated using conventional room air-conditioning." Had it not found a climate control solution for the racks, MPI would have had to distribute the servers across several racks in order to avoid thermal problems and the risk of failure, since room air-conditioning cannot reliably dissipate more than 5 kW heat loss from one rack. This would have led to greater space requirements and higher operating costs.

With this in mind, the IT managers sounded out potential solutions. They opted for the air/water heat exchangers – the Liquid Cooling Package (LCP) – from Rittal. These are flanged onto the rack in place of a side panel and produce a defined microclimate in the enclosed casing by each cooling just a single enclosure. The cooled air from the LCPs is blown in front of the 482.6 mm (19") server mounting angles. This reliably prevents the formation of hot spots. Moreover, increasing the power of the computers does not affect the temperature in the room. This fact enabled MPI to install several high-performance server racks in its IT room despite the huge heat losses. The LCPs are scalable and, depending on the cooling requirement, accommodate up to three cooling modules per rack which thereby allows them to attain a cooling output of up to 20 kW. The fully configured LCP systems at MPI draw the cooling water in their micro circuit from a Rittal water/water heat exchanger, which is connected to an existing building cooling water circuit. The circulation pump in the heat exchanger is redundant. The system switches between the pumps each day to ensure that both are continuously operational. A long downtime might otherwise cause the pumps to seize up and thus restrict availability in the event of an incident.

Computer-controlled monitoring safeguards operation

“Intact pumps are vital for cooling the servers, so we monitor their performance using Rittal’s CMC-TC monitoring system,” says Mario Lüttich. “If one of the pumps fails, we immediately receive a message via the integrated Web interface or by e-mail and can organize a replacement while the other pump maintains operation.” However, using the CMC for monitoring also has an added advantage for Lüttich. He uses CMC to switch off the Rittal PSM socket strips for the servers in the event of a pump defect, thus ensuring that the computers do not overheat.

The complex calculations often take several days and are made using programs developed in-house that run on the Windows cluster. The results are stored successively on an RAID drive system that is secured against power loss using a mini-UPS. “It therefore makes no difference if the calculation is interrupted by a server emergency stop,” says Lüttich, describing the special features of the application. “Our tasks are not time-critical, so we can restart the systems from one of the previously defined and stored starting points. This is also the reason why no UPS is used for the servers."

The server rack contains the latest generation of 1 U servers, which are each equipped with two AMD Dual Core Opteron processors. That equates to four processors per unit, each of which provides more than 2 GB of RAM. In other words, just 14 units accommodate 48 CPUs plus two systems that supply the RAID system with 8 TB of storage capacity. This is the reason for the huge heat generation – even with a rack that has only been half equipped up to now. Consequently, the doors of the racks must not be kept open for long periods because the risk of failure increases greatly within a short space of time if they are not exposed to cooling air from the Rittal LCPs. The temperature of the LCP is controlled using the adjustable temperature hysteresis via the CMC. In a second server rack, MPI is still using the previous generation of Dual Athlon processors which, while supporting the same computing power, fill a complete rack because of the lower packing density.

Conclusion

The new cooling solution from Rittal enabled the MPI in Göttingen to place its IT services for calculating ribosome models on a much more secure footing. The servers now offer high availability, while outages due to thermal problems are a thing of the past. “Rittal is the technology leader and greatly impressed us with its system of water cooling using LCP air/water heat exchangers. We can now house all our computing power in just two racks – and in a small IT room – without taking risks with the service life of the server.”

Components: Rack, LCP, Water / Water Heat Exchanger, CMC-TC