Ernst-Mach Institute uses cool high-end Linux cluster

Fraunhofer pushes the boundaries

Numerical simulation is a rather dry, unemotional term for complex computer calculations that the Ernst-Mach Institute, based in Freiburg, performs for customers using a new high-performance computer cluster. In addition to live tests, these simulations allow experts to check the deformation dynamics of vehicles, materials and structures without destroying them. Supercomputing produces heat. Spatial conditions, future-safety and manufacturer independence prompted Rittal to develop a climate control solution with cutting-edge, rack-based air/water heat exchangers.

The Ernst-Mach Institute (EMI), part of the Fraunhofer-Gesellschaft, has an outstanding international reputation when it comes to testing the road and crash safety of vehicles, the stability of buildings or the behaviour of specific materials in aerospace applications. In addition to a wide range of specialized research projects, extending through to the military sector, numerical simulations from areas of compressible flows and non-linear finite element methods are part of the daily routine in Freiburg. The best known tests for anyone unfamiliar with the industry are the finite element methods from the automotive industry. Material structures are represented by a large number of small but finite polygons, whose properties are easy to calculate. The solved system of equations is then used to determine the results for the polygon network.

This system is complicated in as much that it requires extreme computer power to resolve the complex equations for this type of numerical simulation. In fact, the standard PC would take weeks or even months to work out these equations, whereas the new Linux cluster developed by the Institute in Freiburg can provide the results in a few hours or days.

Simulating car crashes on a computer

“In our computer network, for example, we crash cars and subassemblies on behalf of industry. We use parallelized programs to perform numerical simulations in a massive parallel calculation,” explains Stephan Engemann, head of the computer centre. “Four years ago, a 32-bit Linux cluster with 700 Gflops/s was in operation. This cluster was augmented in 2006 by a 64-bit system with 2000 Gflop/s, which delivered a much more cost-effective computing power than the earlier monolithic systems.” Flop/s is the number of floating-point number operations performed per second.

The legacy system is a proprietary solution that utilizes the entire cooling capacity of the climate control system in the computer centre. Therefore, creating an efficient climate control system was a real challenge. The new cluster includes 64 quad opteron computer nodes, which corresponds to a total of 256 Intel CPUs with 2.6 GHz each and a total of 1 TB RAM. It was the 64-bit technology that made it possible to address an even greater memory range from a single CPU. These CPUs are relatively inexpensive, but they require a permanent exchange of information. Using appropriate tools, the calculations run in parallel and the results are exchanged via an Infiniband high-speed network. This network has less ‘administration overhead’ than conventional networks, making it much faster at 10 Gbit.

Keeping cool has never been so easy

"The new cluster has a total heat loss of around 35 kW," describes Udo Ziegenhagel, system administrator in the Institute's computer centre. "The entire computer hardware is at risk if the cooling system does not work efficiently. It would overheat and malfunction very quickly. Due to the limited space and low ceilings in the computer centre, it was no longer economically viable to use a traditional climate control system in the room. We realised that only a water cooling system would provide the required level of safety and performance."

Not only is the thermal capacity of water higher than that of air, it also offers a more pleasant climate. According to Udo Ziegenhagel, it was very unpleasant in the room even before the expansion of the computer centre. To avoid heat building up at unfavourable locations in the server racks, the room temperature would have had to be kept even lower and the air speeds so high that employees would have had to wear coats to work in the room. An alternative would have been to equip the racks with fewer components. However, this option would have required much more space, and in this case the office's city centre location, combined with the necessary technical infrastructure, would have resulted in very high operating costs.

In their search for a scalable and hardware-independent solution, the Ernst-Mach Institute chose a rack climate control system from Rittal. The Institute decided on the modular Liquid Cooling Package (LCP), which is mounted on the side of the server rack as a "climate control enclosure." In this solution, heat loss is absorbed and dissipated close to the point of origin so that the heat extraction process does not affect the room temperature. As safety is of the utmost importance, six LCPs are installed among the row of five server racks. This means that each rack is provided with cooling air from both sides, thereby providing additional design redundancy for maximising availability.

Cooling on demand

The cooling system is essential for safe operation of the servers. The Rittal solution is also equipped to meet future demands because each LCP already has three cooling modules, which generate a cooling capacity of 15 kW per rack at the water inlet temperature. The server racks are enclosed and fitted with inspection doors at the front and back. The air is circulated horizontally, which means there is no danger of heat building up in the upper section. In the LCP climate control enclosure, the warmed air from the computers is cooled by the air/water heat exchangers and then fed into the servers again ahead of the front 19" level. This microclimate is easier to control and size than a room air conditioning system, which bases the cooling capacity on the worst possible conditions.

The cooling medium is provided by recooling systems – also from Rittal – which are located on the roof of the building. "We wanted to be prepared for whatever the future may bring, so we sized the recooling systems accordingly," explains Stephan Engemann. "Due to reasons of redundancy, we installed three recooling systems, two of which always operate alternately. In the same way, the pumps are switched regularly to prevent excessive cooling and to stop the cooling medium from freezing in winter during dwell phases in the recooling system." A water tank with a volume of 300 l is installed in each recooling system as an additional buffer for the cooling unit. In total, the three recooling systems provide a cooling capacity of 108 kW, which means they are fully equipped for the future.

Electricity is not an issue in Freiburg

Systems generating so much waste heat inevitably need energy in the form of electricity. For this purpose, and for its many test setups, the Institute has its own transformer station from the energy provider, which gives it very high reliability of supply. In contrast to commercial servers, the simulations are not critical company processes, so there is no need for an uninterruptible power supply. Udo Ziegenhagel: “The software that we use – partly developed in-house and partly bought in – stores the calculation results in a non-volatile storage device and marks restart points if the system crashes. If the system does crash, the calculation is started again from where it left off. Safeguarding the processes in this way means that the server can be switched off without any risk of data loss if, for example, the cooling system were to fail. If a condition arises where the servers have to be disconnected, they are switched off automatically using a CMC-TC (Computer Multi Control – Top Concept) via the switchable PSM (Power System Module) socket modules from Rittal.

Udo Ziegenhagel goes on to explain, “We use the CMC-TC monitoring system from Rittal to monitor the ambient conditions of the cluster. Monitoring the temperature is the most important factor, but we also monitor whether the doors are shut as this is essential to ensure reliable cooling. We also check whether the LCP is working and receive information from the Lampertz rack fire extinguishing system.”

The EMI also takes advantage of another benefit of Rittal’s integrated system – every server is connected to a Rittal SSC KVM switch. This solution can be used to manage each server separately via a hardware connection using either the integrated 1 U monitor/keyboard drawer or remotely via the network. Furthermore, each individual PSM module in this system can be disconnected and restarted, for example when a server is not responding.

Conclusion

The 2 Tflop/s capacity of the new Linux cluster at the Ernst-Mach Institute puts it in the realm of supercomputing. To dissipate the incredible amount of heat produced by the system, the experts at the Institute rely on the integral infrastructure solution RimatriX5 from Rittal, which integrates all important areas. Stephan Engemann concludes, “Rittal was a very capable provider, which meant that we were able to hand over the installation work for the climate control system with its very precise specifications to a specialist company. The hardware-independent infrastructure was installed in autumn 2006, and it is so flexible that any type of server can be fitted in the system in the future. The cooling capacity of the LCP solution can grow in line with our requirements thanks to the capacities of the recooling systems installed. What’s more, we can monitor the entire system from a central console. After installing the recooling systems, we now benefit from a system that is far superior to a traditional air cooling unit.” The future at the Institute in Freiburg may be exciting, but the servers will always keep their cool.

Conclusion

HUK-COBURG's new RIO data centre includes 40 server racks and 35 network racks from the RimatriX5 programme. Thanks to its flexible and innovative solutions, Rittal has replaced the racks of the old computing centre and has also scored points over the simple racks provided by the server manufacturers. Another decisive point: apart from the greater depth and individual adjustments made to meet customer requests (air inlet and outlet, power supply) it is also possible to retrofit a liquid cooling system. "Our contacts were always open for specific requests and suggestions", Horst Sonnenberg explained with satisfaction. "If advances in servers mean that we need to assemble our racks more closely in the future, we will still be able to continue using the available Rittal infrastructure. That is a good protection of investments for such an extensive installation, and this has also helped convince us - along with the overall RimatriX5 concept."

Components: Rack, LCP, PSM , CMC