

Case Study

Crash Simulation: Linux Cluster at VW

Summary

NEC has been able to assert itself over the competition at the Wolfsburg automaker in two regards. With the installation of a Linux cluster for crash simulation, NEC was able to continue its presence at VW as well as being able to simultaneously make the move to dedicated high-performance computers on a Linux basis now necessary in the sector.

By Konrad Buck (freelance journalist in Düsseldorf)

The Selection Process

What is a supercomputer? Amongst experts, the answer has long since ceased to be a vector computer. Instead, Linux clusters are now some of the most coveted number crunchers. Such a cluster with 128 computers and 258 IA32-series CPUs has been located in the R&D computing centre in Wolfsburg since the beginning of the year. It allows VW to calculate its crash tests up to five times more economically than before.

The selection process for machines of this type is both simple and straightforward: the contract is awarded to the bidder who offers the highest performance at the best price. In selecting their new high-performance computer, the decision-makers at VW also benefited from the new trend in supercomputing: computers assembled with standard components but which have been optimised for one single application purpose and which hence offer an exceptionally good price-value ratio.

Jörg Warnecke, responsible for High-Performance Computer Systems at VW, says: “As part of the selection procedure, we provided a benchmark data set and demanded a certain throughput per day. The contract was simply to be awarded to the bidder who supplied the best results. The supplier also had to agree to our terms of acceptance so that the machine really lives up to its promises – and if not the manufacturer must retrofit it”. That is a tough selection procedure – but those who want to play in the top leagues of number crunchers can only meet customer demands with the perfect mix of state-of-the-art technology and sophisticated technological concepts.

According to Warnecke, NEC came into play in particular because it was the only company which could supply clusters with 2.8 GHz Xeon CPUs from Intel at the end of last year. Gunther Mayer, one of the users of the IA32 system and R&D co-ordinator for CAE hard- and software at VW, underscores the importance of the co-operation with the manufacturer, which extended over a period of years, and the resulting relationship based on mutual trust: “We place great value in service from a single source. We had previously purchased SX-5 and SX-6 vector computers from NEC. We could thus assume that a solution from NEC would function and that the supplier would support us in hardware and software issues so as to assure performance”.

The Installation

The NEC cluster was supplied in December 2002. Warnecke describes its introduction as smooth: “After deciding on the Linux cluster from NEC, we began test operations at the beginning of this year and were productive two weeks later. The only initial problem was the heat transfer from the

rack. Because of the cabling, the air collected on the back side and couldn't really be moved out of the rack. NEC then had to install additional fans on the top".

The supplier's overall concept was otherwise fully functional, which made it easier for the team in Wolfsburg to decide to become better acquainted with the Linux cluster technique. "At first, we weren't certain how well Linux would fit in our environment in terms of stability and performance", reminisces Warnecke, "in our invitation to bid we therefore requested that an operating concept be co-ordinated with us. All manufacturers submitted good concepts, but the overall concept from NEC – paired with the professional co-operation and the good price-value ratio – finally convinced us".

NEC relies on a simple yet precise approach, in particular when it comes to maintenance, such as mirrored SCSI plates in each computer which can simply be exchanged in the event one is not functioning correctly. If a computer is not functioning for a different reason, it is replaced by a reserve computer which is always kept on-site. This reserve computer receives the fixed discs from the defective computer and thus starts up with the same IP address. A simple and very convincing concept..."

The PAM-Crash Simulation

Another yardstick was the functioning of the cluster with the PAM-CRASH simulation programme from ESI Group used by VW. Originally developed for use on SMP and vector computers, PAM-CRASH has been further developed for the special requirements of Linux clusters for many years now with the help of the MPI library. In order to achieve optimal utilisation of processor performance, MPI software must be extremely well co-ordinated in regard to hardware and software. Thorsten Queckbörner, Project Manager "PAM-CRASH Solver" at ESI Group: "We have spent years optimising PAM-CRASH for clusters. The weak points of clusters, in particular the latency and throughput of the network, should not influence performance at all. After all, when 8, 16 or more CPUs are being used for a job, the network can limit the data exchange and computing capacity may go unused while the CPUs are waiting for data. In order to keep this effect in check, we have now succeeded in optimising PAM-CRASH so that performances losses are kept to a minimum".

In order to meet the high architecture-related requirements of the simulation programme used by VW, the computers brought together for a computer job communicate over a gigabit Ethernet. Each computer has two interfaces, one in the direction of the gigabit switches for co-operating with its "colleagues" and a 100-megabit interface for communication with the intranet over which the computing results are sent to the engineers' workstations. There was initially a problem here as well: the computer was unable to copy back files of more than two gigabytes. NEC had a ready solution, however. A patch that resolved the problem was supplied within just a few days.

The Performance

Andreas Gitt-Gehrke, a system analyst with the VW service provider GNS-Systems in Braunschweig, has this to say of the computer's performance: "We had compiled the LINPACK performance test with the GCC Compiler and then obtained 519 gigaflops in the benchmark. We tried out 2-3 different Executables, with and without HyperThreading. In the process it became

evident that the limiting factor was not individual CPU performance, it was obviously the network itself". In productive operations, the cluster is therefore used not for one job across all CPUs, but for several jobs simultaneously which, for instance, each use 8 or 16 CPUs respectively.

According to Warnecke, the fact that this computing capacity could be made available to the crash calculators is a result of a quantum leap in price-value ratio associated with the cluster technology: "The same performance with vector or RISC computers would have cost far more – if it could even be attained at all. In other words: these 519 gigaflops simply wouldn't have been economically sensible with vector technology".

This quantum leap in price-value ratio now, for the first time, makes it possible to implement techniques such as stochastic methods which require a great deal of computing power. Stochastics are used to take into account deviations present in reality, e.g. in material and geometric parameters, within an accepted tolerance framework. A large number of calculations are needed to ascertain the ranges in the parameters of a vehicle and can be generated automatically with stochastic tools.

Another application is the optimisation of plate thickness. Weight must kept low here without having to cut corners in the crash results. A more economically priced Linux cluster makes such mass calculations more feasible. The engineers can now simply run 100 instead of just 10 calculations in order to see how thin the plate can be made while still yielding optimal crash test values. That can all be calculated with the cluster as often as needed, both quickly and economically.

The Benefits

Added to this are the multitude of statutory tests required by the respective countries to which goods are exported as well as here in Germany. Volker Tietz, head of the R&D department CAE Methods: "If we used conventional methods, we would be unable to afford to perform the majority of tests which we currently need to cover the legal guidelines and consumer tests in Europe and the USA. We would have to build far more extremely expensive prototypes. Of course we still have to perform hardware tests today. But we can rely entirely on simulations in order to help ensure that we will, if possible, pass these extremely expensive tests the first time around".

The goal is also to use the cluster technique in order to represent component failure. In design jargon, "failure" means that a component or a link tears or breaks. Until now, it was primarily the three-dimensional deformation of the supporting structures that was calculated. Now the VW designers also want to represent the failure. To do that, they must be able to define elements at a much finer level – and they can do just that with clusters. In order to be able to already simulate the breaking of components or the tearing of structures in the calculation phase, one has to network in great detail. This is the only way "worst case scenarios" can be virtually played through and hence prevented in reality. "The cluster technique", says Tietz, "is the prerequisite for this".

An Outlook

According to Warnecke, there is currently a clear trend in supercomputing to the Linux cluster. For high-performance computing specialists, both vector computers and RISC CPUs are meanwhile passé. The latter, he says, are now slower than the new Intel chips in many applications. “We, too, are very interested to see in which direction Itanium64 and AMD Opteron will now develop”, says Warnecke. “One thing is for sure, clusters of Itanium64, Intel 32 or similar CPUs are the technology of the future. As it now stands, we will no longer purchase vector computers or large RISC machines. Crash calculations were one of the last fields in which vector computers were used here”.

As a tip for others, Warnecke recommends paying special attention to air conditioning: “A cluster rack produces approximately 8-10 kW of heat. When using this technology on a large scale, one can easily reach the limits of the existing air conditioning in the computing centre and must replace it with a new one”. Computer manufacturers are also working on this problem – hopefully water-cooled racks will soon be available again, but users must have the appropriate infrastructure available. When it comes to cluster computing, heat transfer is definitely an important current topic.

In the 1980s, VW was one of the first companies to introduce crash simulation in the manner described. According to Volker Tietz, the most current challenges in the field of simulation are the representation of new materials such the hybrid steels, plastics, foams and fibre composites as well as assembly technology testing. These topics already go beyond crash simulation and open the door to a whole new field of application for highest performance computers: Computational Fluid Dynamics (CFD). VW is planning on using cluster technology here in future, says Tietz.