

counting the cost

Greener IT is becoming more of a necessity, says **Jean-Pierre Panziera**, chief technology director in the extreme computing division at Bull



Green IT, if you look at it in a general sense, is about being conscious of your environment and the impact that IT systems have. The way that translates into technology is a little more precise, and currently under development. When we talk about the IT industry, and more specifically high-performance computing (HPC), we are focusing on being energy conscious and making sure resources aren't being wasted. Beyond that, it's about the products and components, and the methods used for cooling those parts. If you look at the total budget over the five-year lifespan of an HPC system, energy may account for 30 per cent of that expense. This cost is not only accounted for by the electricity consumption required to run the system; it is also the additional energy that's used to cool the system itself.

All the energy put into the system is transformed into heat, which then has to be removed, and that has a cost - a figure that is gradually being driven down by new technologies. How to remove heat, and especially the considerable amount of heat that can be produced given that an HPC centre can consume several megawatt (MW) of power, is a difficult question. Within the industry there are predictions of 20MW centres in the coming years, and looking back over the course of the past five years, for every Watt used to power a system, typically an additional Watt was used to remove the resulting heat.

This one-to-one ratio has decreased significantly, to the point where most computer rooms use only half the energy produced to cool the system, and direct liquid cooling may bring that figure down further to about 10 per cent. Diminishing the overall energy consumption of an HPC system is firstly about making sure less energy is used to cool it. The second step is to use components that consume less. When considering where the energy put into a system is being used, you see that around two thirds is going into the processors, with the final third going into memory. Ensuring that these elements

The main driving force behind this shift isn't an ecological awareness, it's the fact that electricity prices are continuing to rise

can be powered more economically can go a long way to bring down costs.

The friendly approach

Most companies aim to be greener - it's a growing trend at the moment - but in reality, the main driving force behind this shift isn't an ecological awareness, it's the fact that electricity prices are continuing to rise. Take a look at Europe; on average the electricity costs in the region of 11 cents for every kiloWatt-hour (kWh). There are countries where that could be 15 cents per kWh, while in others it may only be seven cents per kWh - it all depends on the energy source being used. With computing centres facing energy costs of several million euros per year, being green is becoming more of a necessity, rather than a conscious choice.

The challenge right now is the approach used to cool the systems in HPC centres. It is quite simplistic: cold air is blown on the components to cool them down. I tend to think of it in terms of a small child who burns his hand - the first reaction is to blow on the injury, but that is not an efficient way of dealing with the burn. A more effective approach is to place the hand in cold water. In computer technology terms, it would mean immersing the system in a liquid. That liquid cannot be water, as water and electricity are not compatible. Using other fluids is possible, but that presents its own problems, especially for system maintenance, which becomes more complex.

Other solutions have been implemented, like using chlorofluorocarbons (CFCs), which are quite efficient from a technical point of view. When sprayed onto hot components, CFCs vaporise and that vapour carries away most of the heat. However, CFCs are so damaging from an environmental perspective that they have been banned in the European Union. Finally, there are other alternatives, such as circulating cold liquid in a cold plate that is in contact with the hot electronic parts, but these solutions have their own engineering issues that are still being resolved.

But, as mentioned earlier, the main aspect of saving energy within a computing centre is to use components that are less electricity hungry. As an industry we know how to manufacture parts that do consume less, but this raises new problems to be faced by the users and the application developers. The simplest and most effective way to reduce energy consumption is



to use slower components. For example, you get more compute power with many slow processors than with a few high speed ones. So, if considering global throughput, a system with many slower components is the most effective, even though each application will run slower. However, when speed is an imperative target, the application's parallel efficiency becomes crucial. An application with limited parallelism will run at a fraction of the overall system performance and the expected energy savings will not materialise.

At the core

Looking back at the historical developments within computing, the way performance has evolved has meant that if an application required a higher level of performance, it would wait for the processor to speed up. That hasn't been a problem as, following Moore's law, that processor speed has typically increased by a factor of two, every two years. This has, however, come to an end and the only way to gain that improvement now is by the addition of more computing cores. Again, this becomes a challenge for the applications as users have to distribute the job across many cores.

There are orders of magnitude of parallelism required by the applications, and that is definitely something developers are still learning when it comes to reaching those goals. It's fair to say that the most parallel applications these days run in the order of 10,000 threads. To put this in perspective, we will need to have the same applications running at one million threads in the next five or 10 years. There is a big jump in terms of efficiency that is required for the applications, and it

is not something that will be easily solved. It definitely requires a new approach in the way people architect their applications. I am confident that we can overcome all of these challenges; however, the really interesting questions will be when and at what cost?

In the HPC arena today, there are systems deployed with more than 100,000 cores and, if we take an HPC application that we are currently running on these machines and attempt to scale it efficiently to the million core level, we see that the tools currently at our disposal are not adapted. I believe that the hardware challenge will be solved to the point where we will reach an exascale level of computing power by the end of the decade, so the main question is how many applications will be able to take full advantage of those systems? It can take up to 10 years to develop a new application, so, to a certain extent, we have to guess how we should be architecting those applications for the next generation of computers.

One important point that requires improvements is that we cannot specify all the data movement in the same way that we do today; that is, specify which thread is sending data to which other thread. That's a level of detail that is not sustainable for these very large systems. Those reading this may feel that this discussion has moved away from being green to performance, however, it is when we contemplate this level of performance that these environmental issues become all the more sensitive.

Beating bandwidths

To close, there are a few issues that surround memory itself. The first revolves around efficiency in terms of capacity and bandwidth. Looking at capacity, the development curve is much like that of processors, and every few years the number of bytes that can be put into a memory chip doubles. The rate at which users can feed data from the memory to the processor, and vice versa, is not growing as fast, and so bandwidth becomes a problem.

Before dealing with how to become greener, we need to consider the issue of efficiency. The biggest improvement will be to incorporate the processing element, the central processing unit (CPU) for example, and the memory within an integrated module. Many individuals have proposed that in the past, and it would certainly go a long way to improving the amount of bandwidth that's needed. There are some other interesting benefits that come as a result of this - not least that an integrated memory and processor will draw less energy!

Memory bandwidth is definitely driving HPC now, and processor developers are continuing to use standard interfaces that require high voltage. Ensuring that the processor and memory are aligned, in terms of energy saving and performance, is something that we are working on. The emphasis within HPC has traditionally always been on performance, but if industry is to reach all its goals, it needs to pay closer attention to what it truly means to be green.

Considering global throughput, a system with many slower components is the most effective, even though each application will run slower

