



How to sell a Supercomputer

Worried about potential stagnation in the server market? It might be time to branch out into the brave new world of commodity supercomputing, says Adam Oxford.

Although it seems astonishing, to occupy the number one slot in the list of the world's most powerful computers (www.top500.org) you need to adhere to the general principles of Moore's law. Even for the world's fastest computers, performance speeds are roughly doubling every year.

The standard measurement of performance is calculated in floating point operations per second, or FLOPS. Today's fastest machines are capable of a staggering 600 trillion FLOPS (teraFLOPS or TFLOPS). By this time next year, though, the petaFLOP (quadrillion FLOPS) barrier will be a dim and distant memory. Even more incredible is the fact that the headline grabbing NEC Earth Simulator, installed just six years ago at the Earth Simulator Center in Yokohama, Japan, to model the effects of global warming no longer features in the top 10 most powerful machines in the world. It barely matches 1/20 of the performance of the fastest machines today.

This may seem academic in an industry more traditionally geared to selling datacentre servers for mainstream business apps, but the knock on effect is worth noting. While virtualisation and consolidation is flattening off the growth charts for server racks, the commoditisation of hardware and rapid development of demanding applications is making High Performance Computing (HPC) one of the fastest growing markets for the channel today. Industry analysts at IDC have tracked growth in the sector at 20 per cent year on year for the last four years. In 2006, HPC clusters

accounted for a whopping 26 per cent of the entire server market – and they're not just for specialised scientific environments, either.

Off the shelf

The key to growth has been the widespread adoption of commodity parts of HPC applications. An array of standard x86 machines can offer the same computing power as an old-school Sun SMP server for around 15 per cent of the cost: £100,000 compared to £1.5m.

Julian Fielden is the managing director of OCF, the UK's largest integrator of HPC systems and the global launch partner for Microsoft's Compute Cluster Server 2003 (CCS) operating system. He sees the steady rise in the popularity of HPC as a personnel productivity issue.

"Ten to fifteen years ago," he explains, "we used to work very closely with ISVs for CAD applications. A Unix workstation would cost £20,000, the operator would cost about £20,000 and the software would cost about £20,000 – there was equality of interest and equal input into the total solution. Now, the operator is £40,000 a year, the box £1,000 and the software £4,000. So maintaining the productivity of technical/professional people has become more important.

"If you can offload some of the work he used to do at the workstation onto a cluster to be processed more quickly, so he isn't waiting around for information, and the better your ROI is."

Because HPC no longer requires specialist components but is widely available on standard x86

architectures, Fielden says, the market has changed rapidly. Once the preserve of universities and advanced research facilities, now the majority of HPC systems are sold into commercial use. "Four years ago," says Fielden, "70 per cent of our business was done with universities. Now, it's about 30 per cent, the ratio has flipped over."

As prices have fallen, the role of the HPC has also changed from bespoke machine designed for a specific purpose, to a versatile array typically running 10 or 12 applications at any one time.

"Six or seven years ago," Fielden continues, "Universities were full of people trying to get clusters to work. They were trying to design Linux systems as appliances, machines that would run one piece of code, usually that the university had written. What we're finding over the last two or three years is that the real world is accepting Linux as an operating platform and a cluster offers fantastic price/performance benefits. More and more companies are utilising HPC for their own business benefits."

Inside the cluster

There are two main approaches to HPC computing, the cluster and the larger scale grid. Cluster computing involves splitting a large compute problem down into chunks which can be distributed across many CPUs to return the answer more quickly, and relies on intelligent scheduling software to ensure the workload is balanced.

The industry standard for managing clusters is Message Passing Interface (MPI) which provides a common model for viewing network topology and assigning workloads to nodes and clusters. Thanks to the maturity of MPI as a standard and suite of programming APIs



Julian Fielden

cluster management today is much simpler than in the past, and many companies don't even hire specialist administrators.

Most new systems sold are clusters: batches of identical server node totalling anything from 12 to a few thousand CPUs. According to Dr Michael Newbury, Microsoft's product manager for HPC in the UK, the price/performance sweet spot is a 64 server rack of the kind now sold as part numbers by Dell and HP with Windows and applications pre-installed.

Given the growth figures it's no surprise that Microsoft has begun to try and challenge the ubiquity of Linux as an operating system for clusters. It will be bringing the next version of its cluster operating system, Windows HPC Server

2008, to market in a couple of months. Crucially, CCS supports plug-ins for common Microsoft apps like

Excel, allowing large computations to move between desktop and cluster seamlessly. Julian Fielding sums up the industry's attitude towards Microsoft's involvement in HPC succinctly.

"Windows CCS is still evolving," he says, "And the next version will be a lot more user friendly. I do fall for the line that Microsoft's presence in the market



Songnian Zhou

will help with wider adoption. Linux is a frightening word to anyone who doesn't understand it, whereas everyone understands Windows. The interface for Windows CCFC is very familiar."

Familiar is an understatement: in HPC Server 2008 node administration and workload assignment is carried out through a series of simple screens in a workspace that looks for all the world like Outlook 2007 running on Windows Vista. Clustering for Dummies it may be, but no-one should underestimate the potential it has for widening the market further.

Microsoft matters

"HPC used to be rather exotic, rarefied stuff for those with PhDs and ponytails," reckons Microsoft's Dr

Newbury. "Yesterday I met with three partners who wanted to get into the market before lunch. Our entire value proposition

is, in a nutshell, if you know how to use Windows Server you can use Windows HPC, so the barrier to entry has been lowered and HPC can go mainstream."

The two key areas in which HPC evolved out of the universities in the early 1990s were engineering and medicine. In engineering, modelling fluid dynamics in product design on a

computer screen is now much more popular than using wind tunnels, and in research medicine computer simulations of the effects of new drugs are used right up the point of clinical trials. In healthcare, commodity clusters are used to speed up the calculations involved in turning the data returned from MRI scans into visual maps of the brain. HPC is also critical in the development of designer drugs tailored to an individual's genome.

High yield financial products like futures, hedge funds and derivatives simply wouldn't be possible without HPC arrays capable of analysing every option on the stock market and evaluating risk at every tick. Competitive advantage can come down to a mere matter of milliseconds – time has never more literally been money.

Other well known applications for HPC within the potentially huge small to mid-enterprise markets include insurance risk modelling and film animation, but Newbury also reveals he's been working with partners in less obvious areas too, like archaeology – although he was reluctant to reveal too many details. "There's some amazing stuff coming through," he says, tantalisingly. "A few years ago we held a summer school in Southampton where we got schoolkids building a supercomputer and using it to design and build a radically unusual aircraft. They get the computer working and launched a test flight by the end of the day."

Back to the grid

As with most things beyond the desktop, though, Microsoft's ambitions in the area must be curtailed by the wider penetration of Linux in the market – and the growth of the grid network, a kind of meta-cluster, in the private sector.

Songnian Zhou is the CEO of Platform Computing, a leading designer of platform independent software for grid and cloud networks. These brings together machines or clusters from a variety of vendors running different operating systems and often located in different geographic centres, to offer an

“Customers prefer not to have an HP only version of software, or an IBM only version. A strong, committed independent software vendor is important to maximise the resources available at any one time