



virtensys

# Rewriting the rules of server interconnects

Reproduced from an article first published in PC Pro magazine, September 2010




## Virtensys: rewriting the rules of server interconnects

Last week, I had a meeting booked with a software company at its PR firm's offices in London. I can't remember what exactly it was that piqued my interest in this product, other than a usually reliable PR rep telling me that it was really quite something. Its website gave nothing away, so when I sat down in that office in London I expected this to be yet another dynamic, highly leveraged, cloud-enhanced, web-enabled strategic Tribble. How wrong I was. What Virtensys has come up with made me reconsider the design of servers, desktops and indeed the complete hardware design and implementation. It's radical but extremely simple in concept. Better still, it's working now, you can buy it today and it's British, run by people with an impeccable track record.

The basic idea is quite straightforward: your PC, whether it's acting as desktop workstation or server, contains a high-speed internal bus on which almost all the components sit, which nowadays is called PCI Express and works very well. It's an extremely fast connection fabric, and some of Intel's processors will soon come with it built into the main CPU itself. Branching off from this PCI Express bus is a range of other, slower interconnects: there may well be a disk controller that provides SATA or SAS buses to the disk drives; there's probably a network card that provides TCP/IP over an Ethernet connection, or over multiple connections and hence wires in most cases; if you have some shared disk storage then there may be fibre channel connections from an FC card going to a disk storage array. The back of most servers, especially at the high-end of the market, is therefore a rat's nest of wires and fibre.

Some of this tangle can be alleviated by going for a blade server, in which multiple blades, each containing a mini-computer with its own processor and RAM, can be plugged into a single intelligent frame. The frame provides services such as networking that are shared among all the blades, which works very well but has the downside of locking you into one vendor. If you're lucky you might get the option of fitting, say, a Cisco switch that was custom-designed for your particular HP blade server, but that's as far as it goes. So what Virtensys has done is to extend the PCI Express bus outside of the computer itself.

This is done at present by a simple, almost dumb interface card that drops into a spare PCI Express slot – although in future there may be a PCI Express connector on the back of the computer itself – which connects to Virtensys' own box, which is in effect a PCI Express protocol switch. Inside this box is a custom motherboard that accepts PCI Express cards connecting to other host computers. Then there's another set of PCI Express slots into which you plug expensive, high-performance PCI Express cards, typical examples being 10Gb Ethernet cards



on fibre, 4Gb fibre channel disk controller cards, high-speed RAID controllers and so forth. And between the "computer" slots and the "card" slots there's a pair of very high-performance chips that act as a virtualisation switch.

What does this mean in practice? Let's say you want to put all your servers on to 10Gb Ethernet, but you're baulking over the per-server cost of the 10Gb network cards, which can cost thousands of pounds each, so upgrading all your servers can run up a very large bill. After all, 16 servers need 16 cards. So why not virtualise a single physical network card and make it available to each of the host servers as required? It's important to

↑ Virtual machines can also boot from the Virtensys box.

realise that each server is made to believe it has its own 10Gb network card. You use the standard, vendor-supplied 10Gb network card driver for each server, and no OS changes are required. All that's happening is that the PCI Express bus has been physically extended and then put through a very clever switch, so that the one card is visible to multiple computers.

Of course, if you have 16 servers all sharing one 10Gb network card there may come a time when the card is overloaded by the demands of all these servers, but the solution is simple: drop in a second card and make it available to all the servers too. Similarly, if you want to add fibre channel to your servers to allow multiple virtualisation hosts to access a single fibre channel store, you could pay for fibre channel cards in each server, or you can now put one fibre channel card into the Virtensys box and virtualise it to each server.

In fact, you can run all your I/O needs up and down this bus, because if you think about it all your I/O is already going up and down the PCI Express bus within each server itself. All Virtensys has done is extended this bus out of the case and turned it into a switched environment. You can, for example, put a RAID controller into the Virtensys box, connected to a bunch of disks so that each server can boot from the same controller, believing it has its own boot disk. Virtual machines can boot from it too, because they also think they have their own set of storage. Although the Virtensys box isn't a cheap option, at around £30,000 to support 16 servers, the savings quickly mount up once you realise how much sharing can be done by these virtualised cards.

There's one final kicker: to get best performance, it's common practice to use direct disk access from a VMware virtual machine to guarantee high-speed disk access. Unfortunately, this creates limitations on moving VMs around between machines, because they've been hardwired into the disk infrastructure. With Virtensys you can move any VM from one host to another, and it will still "see" the same virtualised disk controller, so hot-moves will still work.

This technology is interesting because it rewrites the rulebook for interconnect, for the cost of upgrading to high-speed networks and disks, and it provides a fully flexible solution that is vendor-agnostic. The firm says it's working with every major vendor (except for Cisco, surprise, surprise). Maybe, just maybe, this piece of British innovation will provide the key to a radical change in the way we build and use servers, and thus percolate down to other devices such as desktops and even laptops. ■

↓ The Virtensys box's motherboard accepts PCI Express cards that connect to other computers.



MARKETING BRIEF