



By Chris A. Ciuffo, Editor

Virtualization yields hardware optimization and new embedded architectures



I've written in this space many times about multicore processors, serial switched fabrics, and virtual environments. But for the years I've been pontificating, I've always assumed that each stood on its own merits – that Freescale 8641D multicore CPUs, for instance, would add more performance to an embedded system. Or that a serial switched fabric such as RapidIO would increase the data bandwidth within a system's interconnect. Or that a virtual environment from Wind River would provide a secure partition so that multiple operating systems could be run on the same machine.

All are true, of course. But a series of very recent desktop PC and enterprise server virtualization announcements has made me wonder how the embedded space is going to react to the combination of some new COTS technologies. You might've missed these announcements; I very nearly did. (Some embedded board vendors are going to have to rethink their architectures – and even their CPU choice.) But connecting the dots between the following announcements paints a picture of radical architectural change when overlaid onto the embedded markets:

- Altera's "DO-254 Global Partner Network" and "DO-254-Certifiable Nios II Embedded Processor"
- TenAsys' "Embedded Virtual Machine eVM"
- PCI-SIG's "MR-IOV"
- NextIO's "ExpressConnect" and "nControl"

Let's get right to it, shall we? DO-254 is the hardware analog of the RTCA/FAA's DO-178B avionics software spec. DO-254 essentially says that instead of running System Function A on a dedicated but separate system from System Function B, they can be combined on the *same hardware* as long as certain hardware provisions and certification steps are achieved. This has the effect of huge savings in Size, Weight, and Power (SWaP) because

one piece of hardware – processors, memory, I/O, subsystems – can safely perform multiple functions. This is sort of like brute-force virtualization, but it's a watershed in defense systems looking to save SWaP and cost. FPGA heavyweight Altera has embraced this concept with a vengeance by making their soft-core Nios II processor certifiable, along with myriad pieces of IP for their FPGAs. And voila – one FPGA can now *safely with assurance* replace entire boards in defense systems.

Virtualization vendor TenAsys has also amped up their embedded campaign through the introduction of eVM, which now allows just about any RTOS to coincide with Windows (32-bit Windows 2000 on up), and not just with the company's own INTime RTOS but with many guest OSs. For now, the hypervisor runs on a Windows OS host with multicore Intel CPUs, takes advantage of all of those Intel VT-x hardware hooks in existence, and intentionally closely couples the two virtual operating environments. The significance here is again SWaP – using one host to act like multiple machines – but there's so much more than meets the eye. In military systems, this means that any old operating system environment can be ported to work with Windows, and that includes all those old mil-specific I/O sensors that have been out of production since the 1970s. This minimizes the effort of creating virtual device drivers, keeps performance to a maximum due to data passing between OSs, and is specifically architected for determinism in real-time, deployed systems. This will get even faster when Intel multicore CPUs and chipsets implement VT-d (VT for directed I/O) to relocate bus master DMA devices¹.

And since virtualization really shines when the IC hardware is designed for it, more chipsets and I/O devices (Ethernet, SATA, graphics, PCIe, and so on) are adding provisions. The PCI-SIG just finished their Suite of Specifications, announcing

Multi-Root I/O Virtualization (MR-IOV) as we went to press. If you remember the death of PCI Express-ASI, you'll understand how MR gets it right this time by extending the PCIe root complex with a protocol that virtualizes *I/O and processor nodes* for high-end servers in the enterprise. This effectively turns PCIe into a smart, fast fabric (albeit in the traditional PCI vines approach) that puts even 10 GbE to shame in data movement. Best of all: It's built into most of the ICs you'll use in new system designs.

The company NextIO provided their IP for MR-IOV, and also just announced a series of extremely impressive server devices that rely on their own PCIe 64-lane/16-port nonblocking crossbar switch that turns PCIe effectively into a bridge and reflective memory extender. In a server architecture, for instance, what once required four OSs, four servers, and 12 I/O devices (Ethernet, Fibre Channel, and SAS) can now be *reduced to only* four OSs running on two servers and three I/O devices. Amazing. The MR-IOV spec from the PCI-SIG is implemented in their switch to maximize the loading on the more-than-capable hardware that was previously underutilized. The SWaP benefits are staggering, not just in enterprise servers and ISPs, but in similarly configured embedded architectures.

Call me whacky if you want, and maybe I'm reading the tea leaves all wrong here. But we haven't even seen the *beginning* of what COTS virtualization is going to do for us. When it hits the embedded market, there will be a virtual upheaval among the dual and quad PowerPC boards with hard-to-understand serial fabrics and odd programming models. I virtually guarantee it.

Chris A. Ciuffo
Group Editorial Director
cciuffo@opensystems-publishing.com

¹ Intel's Virtualization Technology (VT) roadmap is complex and evolving. Search for "Intel Virtualization Technology processor virtualization extensions" and you'll find a hugely useful PDF.