

Dual-node SBCs: A hardware-based approach to red/black architectures

By Michael Slonosky

An industry perspective from Curtiss-Wright Defense Solutions



As the need for data security increases, the need to support both encrypted data and less sensitive data within the same system is also on the rise. One approach for providing so-called red/black separation of sensitive and encrypted data in embedded computer systems is the use of partitioning operating systems such as MILS (Multiple Independent Levels of Security/Safety). For some users, though, a software-based partition solution is unappealing, because of concerns about robustness, design risks, or the associated costs of commercial operating systems, which have to undergo a rigorous National Security Agency (NSA) certification process.

A better solution for many of these customers is to handle the architecture separation in hardware. Partitioning done in hardware provides the most secure and easy to validate alternatives. One way of accomplishing this is to provide a fully separated, dual-node single-board computer (SBC) on a single 6U VPX card. Essentially, you're making a two-headed SBC.

It's possible to run different applications on a multiprocessor, treating individual cores or combinations of cores as an independent processor. However, the process of defining the partitions so that application performance can be guaranteed and determining how peripherals – such as PCIe and serial ports – will be shared can be challenging and time-consuming. The multiprocessor approach requires a hypervisor on the board to properly handle the separation and the sharing of resources, which can add complexity and reduce overall performance.

While there have been single-card dual-node SBC hardware designs in the past, they failed to provide true independence in that their onboard SBC nodes still depended to some extent on shared resources such as mezzanine

sites, memory, or data networks. A truly independent dual-SBC card must provide each node with its own memory, FPGA, XMC site, backplane I/O, and power supply. While there is a compromise in I/O and data plane information that results from the reduced amount of onboard real estate available to a single node, the advantages in terms of robustness and size, weight, and power (SWaP) optimization for many applications can make it an attractive tradeoff. Each of the card's nodes effectively functions as though it were in its own slot.

Even better, using a dual-node SBC to reduce board count in technology refresh applications also significantly improves mean time between failure (MTBF). The robustness of the system in terms of MTBF is only affected by half of the VPX board. If one half of the card fails the other half will continue to run, maintaining its functionality. That means that the MTBF of the card isn't the MTBF of two cards, but is instead the MTBF of a single card.

Another advantage of a dual-node SBC is the ease with which legacy software can be reused when moving an application to a new platform. Since each node is independent, and associated with its own processor, the system designer can take the applications that used to run on two separate SBC cards and port them each to a node on the dual-node SBC. This approach frees the system designer from having to worry about rewriting the applications to function together on a single processor. For example, many legacy applications have a significant investment in Altivec code; when migrating such an application to a new platform, the task of rewriting that code can be daunting.

A dual-node SBC designed with the latest Power Architecture processor, such as the T2080, retains support for Altivec while delivering nearly three



Figure 1 | The VPX6-195 single-board computer (SBC) offers two fully independent processor nodes on one 6U VPX board. The board's SBC nodes are isolated from one another so that failure in one does not conversely affect the other.

times the performance of earlier devices. That means that just one dual-node card can deliver the GFLOPS performance of as many as three previous SBCs. In addition to the performance improvement, the SWaP advantage is significant: Compared to the two SBCs it replaces, the slot count and weight can be halved while the power burden is reduced by nearly two-thirds. What's more, the cost for the single dual-node SBC will be competitive with the cost of just one of the two earlier SBCs that it replaces.

An example of a dual-node Power Architecture OpenVPX SBC is Curtiss-Wright's VPX6-195 (Figure 1). It has two fully independent processor nodes on a single 6U VPX board. Each node features a Freescale 1.5 GHz quad-core T2080 processor and is also provided with its own power, I/O, FPGA, and XMC expansion site. This processing engine is designed so that both of its SBC nodes are isolated and incapable of impacting the other. To ensure the highest level of security, the board also supports Curtiss-Wright's Trusted COTS and Freescale's Trusted BOOT technology.

Michael Slonosky
Product Marketing Manager for
Power Architecture Single Board
Computers, C4 Solutions Group
Curtiss-Wright Defense Solutions
www.cwcdefense.com