


# VME versus CompactPCI

## Who Cares?

By Dick Somes



*Much has been written on the relative merits of VME versus CompactPCI over the last eight years and much of that has been focused on the raw performance of the two bus architectures. It's probably not a bad idea to set the record straight on the issue of bus performance before moving on to the real issue, which is the obsolescence of both busses and indeed bussed interconnects in general.*

VME (introduced in 1981) has its origins as a bussed version of the MC68000 memory and I/O bus, a sort of second-generation Unibus for the processor that the PDP-11 should have grown up to be. It's worthwhile to note the similarity of the MC68000 architecture, with its unified memory and I/O address spaces, its orthogonal instruction set, and its sixteen 32-bit registers, to the PDP11 and VAX architectures.

VME was first deployed at a time when aggregate bandwidths exceeding 1 Mbyte/sec were newsworthy. A theoretical bandwidth of 57.6 Mbytes/sec has been claimed for basic 32-bit VME, and 40 Mbytes/sec is widely accepted as a practical limit.

The claim of 40 Mbytes/sec throughput for the basic VMEbus, as described above, assumes that a Data Strobe/Data Transfer Acknowledge handshake can occur in 100 nanoseconds and that the data phase of a Block Transfer (BLT) takes place at 10 Mtransfers/sec. Thus 32-bit transfers proceed at a burst rate of 40 Mbytes/sec, and a somewhat lower sustained rate due to bus arbitration and addressing overhead.

The first significant enhancement to VME was the introduction of a multiplexed 64-bit block data transfer protocol (MBLT) which used 32 address and 32 data lines together as a 64-bit address bus in the addressing phase and as a 64 bit data path during the data phases of a bus transac-

tion. The result was a doubling of the burst data transfer rate to 80 Mbytes/sec.

Further extensions to VME signaling have been standardized, 2eVME and 2eSST, but not widely adopted.

PCI originated in the early 1990s as a chip-level interconnect, conjured up by Intel and the other members of the original PCI-SIG as an architecturally neutral bus spanning multiple generations of CPU and core logic technology. The success of the PCI movement is unquestioned, with chip level peripherals being offered by a growing number of manufacturers and core logic chip sets for multiple CPU architectures over two and three generations of architecture. With the adoption of PCI the computing industry has finally achieved commodity economies of scale for extremely sophisticated graphics, storage, and communication options.

PCI may be operated 32-bits or 64-bits wide with full interoperability between devices supporting either width on bus segments of either width. The theoretical burst bandwidth is 133.33 Mbytes/sec for 32-bit transactions and 266.67 Mbytes/sec when operating 64 bits wide. These are theoretical maximum limits, which compare to the 40 and 80 Mbytes/sec figures often quoted respectively for the VME 32-bit BLT and 64-bit MBLT signaling protocols.

The most recent enhancements to PCI are the PCI-X, PCI-X DDR, and PCI-X QDR protocols extensions which result in data rates exceeding 1 Gbyte/sec, but over short interconnects which are normally useful only for dedicated chip-to-chip connections.

In a nutshell, the PCI bus architecture is newer and faster than VME, but it is also targeted as the solution to a different problem. The CompactPCI specification adapted the PCI electrical specification to

6U x 160mm Eurocard packaging standard in the same way that VME adapted the MC68000 bus. The higher performance of PCI in the CompactPCI form factor means shorter backplane segments as compared with VME. VME also has a superior interrupt handling architecture for many embedded applications.

CompactPCI, by contrast, has grappled more effectively with the problem of module hot swap and platform management, and has therefore reached a higher level of maturity in supporting high availability applications.

So what is the function of the backplane bus today?

The reality is this: No one designing a new platform cares very much about bussed interconnects. And why? Because switched serial interconnects are now capable of delivering aggregate performance exceeding that of the bus.

A switched Gigabit Ethernet infrastructure deployed on a PICMG 2.16 or VITA 31.1 platform delivers the rough equivalent of a 32-bit 33 MHz PCI bus segment to each slot in the backplane and with the slot limitation imposed by CompactPCI. Across 19 node slots (the number supported in a 19-inch wide enclosure) the aggregate is on the order of 2 Gbytes/sec, a 20-to-1 ratio when compared to a basic 32-bit 33 MHz PCI Bus.

In VME systems the gain is even more dramatic, compared with the 40 Megabyte aggregate bandwidth of a basic 32-bit VME bus.

Now, of course, the advantage in terms of aggregate performance is reduced with respect to either bus when compared to 64-bit operation or operation at faster clock rates, but there are other architectural advantages to switched serial interconnects.

- Switched serial interconnects are inherently easier to make redundant. In a bussed system the bus is a single point of failure no matter what else is done to enhance availability. Both PICMG 2.16 and VITA 31.1 support redundant Ethernet switches.
- Switched serial interconnects are inherently easier to hot swap. In a bussed system care must be taken to prevent data corrupting transients on the shared interconnect. In a switched serial platform each slot has a dedicated connection and hot swap transients are contained there.
- Switched serial interconnects are a better match for the functional

densities of today's board-level functions. When VME was introduced, the area of a 6U x 160mm board was barely enough to house the processor and main memory was often connected to the bus. When CompactPCI was introduced, basic processor functions could be contained on a single board, but additional slots were required for high performance peripherals. Today the 6U x 160mm board is capable of containing a fully functional compute node. Even I/O functions normally contain supporting intelligence. Peer-to-peer communications between intelligent devices is better supported in a switched serial environment than over a bus.

- The 20 years invested in the development of the network and the TCP/IP stack is better leveraged in a switched serial environment. Communication between nodes within an enclosure is indistinguishable at the application level from communication across a continent.

For these reasons, and more, Gigabit Ethernet has already effectively replaced VME and CompactPCI as the workhorse interconnect within Eurocard form factor systems. But this transition is only the first wave of change.

Additional higher performance interconnects have already been written into the CompactPCI family of specifications. PICMG 2.17, for instance, adds StarFabric support to Gigabit Ethernet providing a platform with two redundant switched serial interconnects. This mapping also applies directly to the VITA 31 family of standards.

The first generation of StarFabric uses four 622 Mbaud LVDS links to provide 2.5 Gbaud of throughput to each port of redundant switches. StarFabric provides a switched serial environment for transparent PCI-to-PCI bridging at rates capable of supporting 64-bit 66 MHz segments at each switch port. This intermediate level of performance is a factor of 2 to 1 faster than Gigabit Ethernet, yet it is still capable of operation over the legacy CompactPCI connector (the same family used for the VME64x J0 connector).

Looking beyond StarFabric, the next wave of interconnects is already upon us. Ten Gigabit Ethernet, InfiniBand, and PCI Express share a common electrical layer supporting 2.5 Gbaud signaling on a single differential pair of conductors. (Two pairs are used per link to support full duplex operation.) At these data rates, specialized connectors are needed to cleanly couple signals to what are effectively

backplane transmission lines. The connectors used by VME and CompactPCI are just not up to the job.

Both PICMG and VITA are specifying electromechanical environments for this next generation of switched serial interconnects. In the long term PICMG's AdvancedTCA family of specifications, as well as VITA 34, provide a next generation packaging solution optimized for 2.5 Gbaud switched serial technology. But there is also an evolutionary path available for CompactPCI and VME users.

About a year ago VITA began work on a new family of standards under the name VXS, a variant of traditional VME packaging with backward compatibility for legacy boards. VXS systems substitute a high performance differential connector for the 2mm Hard Metric connector currently defined as J0 and provide special slots for redundant switches. This environment is intended to support switched serial interconnects in the 2.5 Gbaud range.

Likewise a variant of the CompactPCI specification has been defined under PICMG 2.20 which substitutes a differential connector for the 2mm Hard Metric connector used in the J4 zone. The differential connector used in PICMG 2.20 systems is a member of the same ZD family selected for AdvancedTCA and VITA 34. PICMG 2.20 is backward compatible with PICMG 2.16 (Ethernet) and 2.17 (StarFabric).

So what is the role of VME and CompactPCI in these emerging switched serial architectures? Both provide a well-known power and packaging environment for switched serial platforms. Subrack and enclosure hardware is readily available. So are power supplies and fans. Even the elements of a management architecture are in place, based on Intel's Intelligent Platform Management Interface (IPMI), defined for CompactPCI platforms by PICMG 2.9 and adapted for VME by VITA 38.

Both CompactPCI and VME provide IEEE 1101.x compliant packaging familiar to the industry for more than 20 years and supported by many manufacturers. Even IEEE 1101.2 conduction cooling is now defined for CompactPCI boards in the recently approved ANSI/VITA 30.1 standard. The only real difference between one platform and the other is the backplane.

Both CompactPCI and VME distribute +5V and +3.3V as the primary logic supplies, with the capability of delivering around 40 watts from either or both supplies. The thermal characteristics of the

enclosure are unspecified for both architectures and are dependent entirely on the cooling components installed.

In short, there is very little difference between CompactPCI and VME when providing support for switched serial interconnects. The architecture for Gigabit Ethernet support is essentially identical, with the same connector and pinouts used on node cards. VME 31.1 systems use the same switch slot layout and connectors as PICMG 2.16. The architecture for 2.5 Gbaud link technologies differs between the CompactPCI and VME families primarily in the choice of connectors.

What are the criteria for choosing between CompactPCI and VME? If support for legacy boards is a requirement, the choice is obvious. But if legacy support is not an issue, the choice becomes subtler. CompactPCI has an advantage in terms of user I/O pin count over VME, even when the five-row DIN is used. CompactPCI also has an advantage in terms of defining radial control lines for hot swappable modules.

But in the final analysis, the choice of bus architecture on 6U x 160 mm boards has become irrelevant. Configuration discovery and hot swap control is no longer a function of a backplane bus, but a platform management bus. The base interconnect is no longer the bus, but Ethernet. The high-performance I/O required among intelligent peripherals exceeds the capabilities of a bussed interconnect and must be handled by the emerging 2.5 Gbaud technologies.

There is very little reason left to care about the backplane bus technology.



**Dick Somes** is Director of Standardization for Force Computers, representing Force at PICMG, VITA, and numerous industry consortia and standards

development organizations. He has been the Vice President for Technology for PICMG since 1997. Dick holds a BSEE from Tufts University and an MSEE from the University of Pennsylvania.

For more information, Dick can be contacted at:

Dick Somes  
Director of Standardization  
**Force Computers**  
115 Flanders Road  
Westborough, MA 01581