By Ashley Eikenberry,
Group Marketing Manager for OEM Board Products,
Sun Microsystems.

CompactPCI in the Era of Convergence

Today solutions need to support the higher levels of availability, and provide easy and rapid serviceability, and the ability to scale. Scalability is the only way to meet ever-changing market demands. To manage this and still remain competitive, adherence to industry standards and open technology have become a strategic imperative. It is critical that systems incorporate components based on open, industry standards. CompactPCI is such a standard, as explained in this article.

With today's emphasis on services in the telecom and Internet infrastructure, the last thing a carrier or service provider wants is to be distracted by the need to maintain its own proprietary compute platform and operating system. Telecom and service providers need to concentrate on their core competencies and cannot afford to be deflected from that focus. What they want is the ability to build solutions with rich third-party content that supports faster and easier deployment of services with lower cost of ownership.

To that end, solutions need to support the higher levels of availability, and provide easy and rapid serviceability, and the ability to scale. Scalability is the only way to meet ever-changing market demands. Telecom and service providers want the freedom to quickly scale up to larger servers. At the same time, they want to be able to scale horizontally with increments of thin compute systems. With a platform that supports availability, serviceability and scalability, carriers and service providers can focus on their application services and system integration.

To manage this and still remain competitive, adherence to industry standards and open technology have become a strategic imperative. It is critical that systems incorporate components based on open, industry standards. The sheer amount of development around an open standard not only ensures greater reliability but also makes the components much more cost-effective and extends the number of choices for solutions providers.

Building upon standard hardware and software, telecom, service providers, original equipment manufacturers (OEMs) and network equipment providers (NEPs) can focus on adding more value-whether in the central office or in the field-instead of having to reinvent foundational technology for each project. Fortunately, access to abundant third-party computing capabilities for both hardware and software has created a thriving market for commercial telecom and service provider solutions. As a result, NEPs and OEMs are building systems using third-party, standards-based solutions founded on commercial, off-the-shelf (COTS) components built around the widely accepted CompactPCI standard.

WHY COMPACTPCI?

On the hardware side, the biggest decision is choosing what board-level technology to use as a standard. Careful consideration should be paid to selecting a standard that is highly reliable, yet flexible enough to allow systems to be tailored to specific needs. For rack-mount systems, the choice narrows down to PCI, VME and compactPCI (CompactPCI).

The PCI bus is the industry standard for millions of desktop systems. Unfortunately, it does not provide the higher levels of reliability or, consequently, the uptime needed in a high-availability system. There is no easy way to cool this type of board and it incorporates edge connectors that are notorious both for being somewhat unreliable and for making board replacement difficult. In its favour, however, the PCI standard leverages the tremendous advantages of the enormous PC industry, making it very cost-effective, reliable and extremely flexible given its wealth of robust device drivers and proven, inexpensive silicon.

The VME standard, on the other hand, was specifically developed for industrial applications, where high availability has long been a major concern. As a result, it offers superior reliability and is easily installed or removed. It is, however, an IEEE standard targeted at high-performance niche markets. For example, VME's 33-MHz performance is just 40 MB/sec for 32-bit operation and 80 MB/sec for 64-bit operation. Moreover, any software ported to the VME environment requires

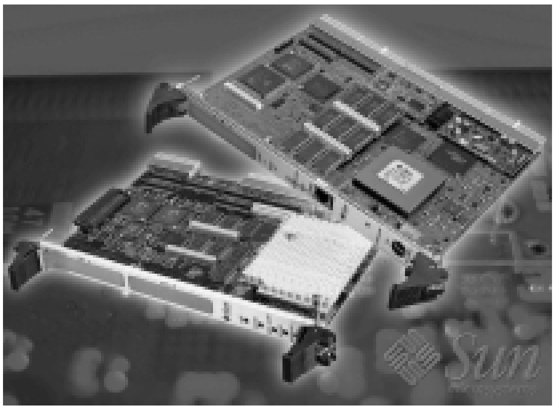

Figure 1. Sun592BP.

Copyright 2001 by Dedicated Systems Magazine - 2001 Q3 (http://www.dedicated-systems.com)

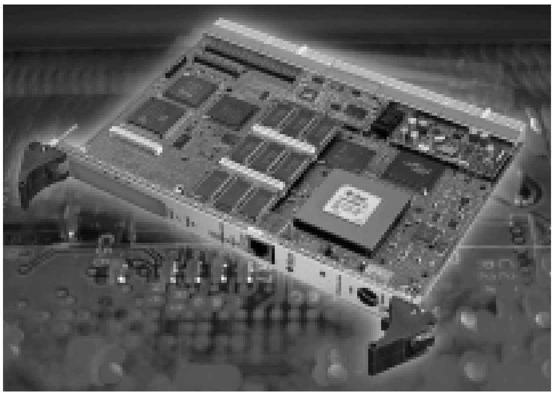

Figure 2. Sun59-40.

customisation, increasing the verification and support tasks.

To address the limitations of these two established standards, a consortium of over 400 computer suppliers and manufacturers, including Sun Microsystems, worked closely together to create the CompactPCI specification. This standard deliberately merges the performance, scalability and reliability of VME with the cost-efficiency and flexibility of the PCI standard. Network, telecom and service provider manufacturers are embracing this new approach. The benefits of CompactPCI include:

- **Standard form factors**: CompactPCI is offered in two different card sizes, making it an easy replacement for VME-based solutions.
- **Compatibility**: It is electrically equivalent to the desktop PCI bus and supports exactly the same interface chips as those used in desktop PCs and workstations.
- **High performanc**e: Operating at 33 MHz, CompactPCI's performance is 132 MB/sec in 32-bit systems and 264 MB/sec in 64-bit mode.
- **Scalability and expandability**: Through standard bridging practices, CompactPCI offers support for up to 256 PCI buses capable of concurrent operation, although physical limitations often restrict this considerably. CompactPCI also supports eight slots per system as compared to four in PCI systems. This can be easily expanded with bridge chips to support up to 32 slots. Strictly speaking, the limit is seven loads per bus on CompactPCI and four loads per bus on PCI. One can have multiple buses in both PCI and CompactPCI systems.
- **Reliability**: This standard incorporates a 220-pin, 2-mm "hard metric" connector that ensures adequate shielding and grounding for low ground bounce and reliable operation in noisy environments. CompactPCI was also designed to support vertical card orientation for superior cooling.
- **Ease of service**: Rear-panel connectors minimise the service time needed to replace a failed board. When a board fails, it can be replaced without unplugging cables. This has substantial service benefits, e.g. faster MTTR, a simpler repair process, etc. NEPs and OEMs can choose from two different fundamental rackmount architectures for third-party solutions: the industry-standard CompactPCI, and the proprietary VME architecture.

CompactPCI meets the need among NEPs and OEMs for a flexible, carrier-grade, rackmount platform. It also conforms to the rugged Euro-card form factor long favoured by telecom companies. Plus, NEPs, OEMs, service providers and telecoms want open standards and APIs, both of which CompactPCI offers. For example, CompactPCI incorporates the PC-standard I/O, leveraging the wide range of PC device drivers, I/O drivers and operating system support. Equally important, CompactPCI provides comprehensive hot-swap and high-availability features.

CompactPCI is also flexible enough to accommodate a wide range of applications. With all these benefits, CompactPCI is rapidly gaining acceptance in the industry, resulting in Sun selecting it as the hardware foundation for its high-availability board products.

NEXT-GENERATION COMPACTPCI CARRIER-GRADE BOARD SOLUTIONS

For all its strengths, CompactPCI, in and of itself, cannot support the stringent high-availability requirements-the so-called "five nines" or 99.999 percent high availability-that the telecom industry demands. To achieve this level of high availability, the entire system architecture, from the hardware and operating system on through to the applications, needs to be high-availability aware.

However, the current specification for CompactPCI focuses on only the hardware aspects of achieving high availability. This specification, developed by PICMG, the standards body that governs CompactPCI development, allows CompactPCI products to support hot-swap capabilities, but provides no API or framework for device drivers, middleware or application software to consistently take advantage of hot swap. No one vendor has stepped up to offer a complete high-availability solution. The result has been a series of point solutions in which the card vendor writes a high-level device driver or middleware component to handle failover and hot swap, creating a spider's web of paradigms and interfaces that must be maintained.

To meet the many needs of the complex, evolving telecom and service provider markets, some CompactPCI vendors are extending the basic standard to provide support for higher levels of availability and uptime, crit-

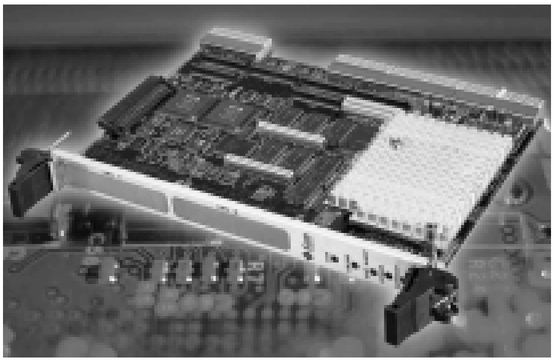

Figure 3. Sun59-60.

ical in today's data-centric networks. For example, Sun's NetraTM CP2000 high-availability platform offers a suite of board solutions that make high-end, data-center-class RAS features available on cost-effective, industry-standard CompactPCI boards, and at an extremely attractive price point. Designed to deliver the scalability, high availability and performance demanded by today's and tomorrow's telecom and service provider infrastructures, this platform is built around standards-based hardware and software components, with Sun's 64-bit UltraSPARCTM Ilex microprocessor technology at its core.

Specifically, the platform includes NEBS Level 3-certified cards along with an innovative backplane architecture that improves system resilience after single failures. The solutions include a set of universal cards that function as a system controller, an alternate system controller or as a satellite card. Sun has designed another CompactPCI board to be a satellite board that provides a low-cost solution easily customised for satellite processing.

A key feature of these cards is that they can be installed either in a system controller or a satellite slot, making it possible to construct distributed or loosely coupled multi-computers. For most systems, the system controller is running a UNIX operating environment, with the satellite cards hosting either the same UNIX environment or some real-time operating system (RTOS)/firmware operating environment. The system has a mechanism to allow booting across the backplane; mainly, satellite cards boot from the system controller. It also has a comprehensive set of management utilities.

The CP2000 architecture implements a variety of techniques to improve the resilience of the overall system. For example, it provides comprehensive hot-swap support with user-level interfaces that allow applications to register their interest in various hot-swap events. A weakness in the current hot-swap CompactPCI specification is that cards are either on the PCI bus or they are completely disconnected. This makes it difficult to perform diagnostic work without jeopardising the system integrity. The architecture also includes an IPMI bus for an out-of-band communication link between the system controller and the satellite card. This means it can be used as a comprehensive debug port without affecting the main PCI bus.

The system controller itself, a key potential point of failure, is backed up by a second system controller on the PCI bus. A second IPMI bus allows the two controllers to agree on which is mastering the bus. Protection from the failure of the PCI bus is achieved by disabling (tri-stating) any problem cards. Moreover, the platform incorporates a variety of hardware- and application-specific techniques used in many telecom applications to improve reliability and availability.

In many existing applications, for example, the framework for failover and restart has already been implemented, though it remains complex and non-transferable. Therefore, a key part of the CP2000 high-availability architecture is to implement a checkpoint, failover and restart mechanism for applications running within a shelf. This mechanism is designed to be fast, lightweight and applicable to both the SolarisTM Operating Environment and ChorusOSTM operating-system.

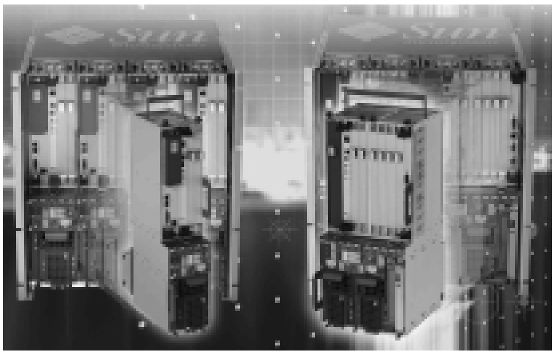

Figure 4. Sun60P.

EXAMPLE DEPLOYMENTS

NEPs and OEMs can use the CP2000 building blocks to create different types of system architectures for telecom and service provider applications. These can range from multi-computing to multi-racks.

A NEP or OEM could, for example, construct a signalling gateway. The CP2040 would be used as the system controller and up to seven I/O boards(e.g. CP2080's with 3rd party PMCs) could be added for the SS7 or IP signalling functions. Both the system controller and I/O cards would run a UNIX operating system. The I/O cards could be hot swapped or set up in redundant node pairs, depending on the needs of the telecom or service provider.

The single system controller in the above example may become a bottleneck for a signalling gateway that is running advanced applications. In that case, satellite CPUs can be added. These cards can run an RTOS or UNIX operating system, depending on the application need. Communication among these boards could be carried out across the backplane or via an Ethernet switch.

A third example is a media gateway, where I/O cards are replaced with CP2060 boards configured as DSP cards. Due to the high cost of trunk-capacity DSP cards, it may be preferable to employ an N+1 arrangement. Another recommendation is to use two host controllers in a system controller/alternate system controller approach. In this scenario, one would function as the main system controller, while another serves as an alternate controller. In the case of a failing main host, this redundant controller would step up and take control of the database engine with no perceptible interruption to the end service.

HIGH-AVAILABILITY READY

By starting with the CompactPCI standard and carefully enhancing where appropriate, Sun's CP2000 program delivers a complete high-availability architecture for the telecom industry. It does this by extending the ability of the base CompactPCI standard and incorporating a secondary IPMI bus combined with optimal levels of redundancy. Together, these enhanced

Copyright 2001 by Dedicated Systems Magazine - 2001 Q3 (http://www.dedicated-systems.com)

AD TEXAS INSTRUMENTS

CompactPCI boards form the foundation for building highly available systems that can be integrated with an integrated stack that includes operating system support, middleware and higher-level uptime capabilities. These features enable providers like Sun to offer low-cost platform solutions with the critical uptimes for today's data-centric networks

---

Ashley Eikenberry is Group Manager for compactPCI and AX board products at Sun Microsystems. She joined Sun in 1997 as Product Manager for compactPCI board products, and from that time has been involved with the development of all Sun's compactPCI products, from Sun's first compactPCI processor board, the CP1500, through to the current Netra CP2000 range. In her current role, Ashley also manages Sun's full compliment of PCI-based Netra AX motherboard products.
Prior to Sun, Ashley worked in a marketing function at Versant Object Technology, an object database company focused on the telecom market. In this role she worked on product alliances with companies such as Sun and Oracle to position Object Technology's value proposition. Ashley has a degree from the University of Rochester and an MBA from the University of San Francisco (USC) with an emphasis on telecommunications.

Copyright 2001 by Dedicated Systems Magazine - 2001 Q3 (http://www.dedicated-systems.com)