CompactPCI® + Ethernet

The Winning Alternative to Aging VMEbus Systems

Solutions White Paper

January 2008

www.pt.com

History of VME Systems

There was a time, not all that long ago, when all computing based solutions were built from the ground up using raw components. No matter which industry or segment one looked at, each had their tiering of manufacturers who believed the only way to sustain or increase their market share was to design and build everything themselves. Times have changed, and while the purchase and integration of higher level sub components such as boards etc. is not new, the adoption and use of Commercial Off-The-Shelf (COTS) subsystems is growing rapidly. The complexity and capabilities of base technologies are increasing while staffing levels of most application system builders are decreasing. The use of pre-integrated and embedded computing products is now distinctly advantageous and necessary to keep pace with both user demands and ever increasingly aggressive competitive landscape. This trend was seen extensively in communications industries, and other sectors such as aerospace and defense are rapidly following suit.

The introduction and commoditization of standards based computing systems has rapidly increased the choice and capability of such open COTS products. Networking and Internet developments have accelerated and enhanced the adoption of such strategies, not least of which because of the single most pervasive standard for “gluing” together disjoint computing and network elements – Ethernet.

Open hardware standards have taken full advantage of Ethernet and this paper will discuss how the marriage of one leading standard with Ethernet has created a powerful architecture for modern computing platform development.

CompactPCI® – The Beginning

During the early 1990s, the embedded and industrial computing community began to look for the next step forward for rugged, modular computing. There was a distinct need for greater performance linked with flexibility, as well as a route to lower cost, more commodity like components and chip sets.

In 1994 the PCI Industrial Computer Manufacturers Group (PICMG®) was formed. Comprised of a consortium of industry leading embedded and industrial computer vendors, PICMG now counts over 450 members. Based on earlier work done by Performance Technologies and others, CompactPCI® 2.X was born in 1995. Developed from the popular Eurocard format familiar to all users and adopters of VME, the addition of the PCI bus significantly extended the reach of the physical standard.

The PCI bus had previously prevailed in the merchant personal computer market and was rapidly being adopted by the server and custom markets because of its lower cost and availability of high performance silicon and software.

Recognizing the developing needs of communications applications, a high I/O pin count was also incorporated, as well as future support for hot swapping of blades, something that is still a challenge to legacy VME architectures.

The PCI Industrial Computer Manufacturers Group introduced CompactPCI or PICMG 2.X in 1995. At the time, it represented a sea-change for the embedded computing industry. The form factor offered a valuable, new choice to the industry which had previously been dominated by the VMEbus standard. CompactPCI quickly became widely adopted, particularly within the telecommunications market. The technology anticipated and accommodated leading embedded computing technology innovations, including high-end processors, high-speed bus backplanes, and state-of-the-art power delivery and cooling.

CompactPCI grows up

Over the last 11 years, CompactPCI has undergone a continuing advance of “point specs” which have incrementally improved the architecture, adding features such as hot-swap, redundancy, telephony bus, and support for telephony mezzanine modules. Most dramatically, in 2001, PICMG 2.16 was developed to extend the life of existing CompactPCI systems by combining the inherent robustness and reliability of CompactPCI with packet-switched Ethernet fabrics.

This eliminated bottlenecks in the traditional bus approach taken by VME and the original CompactPCI technologies. Based on a network of independently switched nodes, it not only removed the limitations on the number system elements, but also on the overall throughput capacity of the system.

The PICMG 2.16 spec describes a redundant, switched 10/100/1000 Ethernet network within a CompactPCI chassis, providing IP connectivity using a “star” topography routed across the backplane.

The spec allowed for flexibility in selecting the connection speed and higher-level protocols, based on specific application needs. This freed OEMs to mix and match best-of-breed switches, single board computers (SBC), chassis, and sub-components to exactly meet their particular needs. The modular makeup of 2.16 system design simplifies deployment, increases serviceability and greatly improves reliability. Also, with the switch and wiring embedded in the chassis, external cabling is minimized, lowering overall costs.

Decline of VME

According to some analysts, VMEbus technologies still account for over 25 percent of all embedded market applications. However, with the recent rise and growing acceptance of 2.16 CompactPCI and MicroTCA™ in the aerospace and defense segments, many are left wondering – has VME been taken as far as it can go through incremental improvement? Should it be phased out in favor of other form factors that are designed from the beginning to accommodate the technology infrastructure needs of the future?

CompactPCI has now gained broad market acceptance and in 2005 moved past VME in shipments. In a presentation given in January 2007 at the Bus and Board conference, Paul Zorfass, director of Venture Development Corp’s (VDC) Embedded Hardware and Systems practice showed excerpts from their research which placed CompactPCI at the head of the pack with a 32% market share. Growth curves indicated that CompactPCI would continue to widen this lead.

From Communications to Wider Uses

The communications market has been the strongest adopter and supporter of CompactPCI technology, which is logical given its roots. The aerospace and defense segment is the next largest and where the benefits of 2.16 CompactPCI are gaining in popularity. As evidenced by the 450 plus member companies of PICMG, added to numerous other non-members who supply compliant or related products, there exists a thriving ecosystem for CompactPCI developers to take advantage of in their design activities. Whether one needs chassis, CPUs, DSPs, analog/digital I/O, or communications blades, there are plenty of choices amongst the many thousands of products available.

In the early days of CompactPCI, the delineation with respect to VME was clear. The PICMG committees took an approach that was intentionally focused on building a technology with communications in mind. Thus the specs mandated a greater number of pins reserved specifically for I/O and accommodated the ability to utilize higher performance and more powerful and power hungry processors, e.g. Intel® Pentium® brand. The overall system bandwidth was greater, as well as the ability to add communication specific sub buses and private backplanes. Hot swap capabilities, and the attendant increase of system uptime, were also key to the communications industry.

As we have all observed, the computing industry is rife with examples of technologies that were designed for one purpose and wound up used and modified for other, equally lucrative applications far from the original intention.

CompactPCI began with a telecom base, both in the central office and remote sites, but has grown to encompass military, industrial and other governmental applications. In particular, with the rise in military budgets since the early 2000s in European nations and in the United States, 2.16 CompactPCI has come to represent a major threat to established bus architectures such as VME, especially within new aerospace and defense projects.

Aerospace and defense as well as industrial equipment integrators have the same fundamental needs as telecom equipment manufacturers and have been moving over to CompactPCI systems for fixed applications including factories and remote sites, as well as mobile needs such as airborne, ship board and military vehicles. In these scenarios, VME-based systems struggle under the weight of throughput limitations due to their bus-based approach.

CompactPCI Sub-Specs and Backplane

As the base specs took hold and gained momentum, PICMG turned its attention to the development of a series of sub-specs which have led to greater differentiation and applicability for multiple markets. Details of all CompactPCI related specs can be seen in Figure 1.

CompactPCI Spec	Description
PICMG 2.0 R3.0	CompactPCI Core Specification, with PICMG ECN 002 on Self-Describing Slot Geography
PICMG 2.1 R2.0	CompactPCI Hot Swap
PICMG 2.2 R1.0	VME64x on CompactPCI
PICMG 2.3 R1.0	PMC on CompactPCI
PICMG 2.4 R1.0	IP on CompactPCI
PICMG 2.5 R1.0	CompactPCI Computer Telephony
PICMG 2.7 R1.0	6U CompactPCI Dual System Slot Specification
PICMG 2.9 R1.0	CompactPCI System Management
PICMG 2.10 R1.0	Keying of CompactPCI Boards and Backplanes
PICMG 2.11 R1.0	CompactPCI Power Interface
PICMG 2.12 R2.0	Hot Swap Infrastructure Interface
PICMG 2.14 R1.0	CompactPCI Multicomputing Specification
PICMG 2.15 R1.0	PCI Telecom Mezzanine/Carrier Card Specification
PICMG 2.16 R1.0	CompactPCI Packet Switching Backplane Specification
PICMG 2.17 R1.0	CompactPCI StarFabric Specification
PICMG 2.18 R1.0	CompactPCI Serial RapidIO Specification
PICMG 2.20 R1.0	CompactPCI Serial Mesh Backplane Specification

Figure 1: CompactPCI-PICMG 2.x Specifications

From the perspective of architectural flexibility, especially for communications and higher bandwidth applications, one sub spec has added significant extra functionality to the original base spec.

The most broadly adopted extension to the PICMG 2.x family of specifications overlays an Ethernet packet based switching architecture on top of CompactPCI to create an Embedded System Area Network (ESAN). The CompactPCI Packet Switching Backplane provides a redundant, switched 10/100/1000 Ethernet network within a Compact PCI chassis providing connectivity between all slots using a star or dual star topology. It was specifically designed to coexist with 64 bit CompactPCI and telephony H.100/110 buses.

Applications for CompactPCI

CompactPCI 2.16 has now been adopted by many industry segments and across a broad range of applications. The greatest popularity has been for solutions where communications and flexible, high bandwidth capabilities are required. All major protocols and connectivity interfaces are supported amongst the wide selection of available products. With the industry push towards Next Generation (NGEN) architectures, CompactPCI systems will be utilized for multi-media routers, Gateways, voice servers as well as numerous 3&4G wireless network nodes such as HLRs or RNCs. With a similar focus on network based communications architectures now becoming core to aerospace and defense applications, CompactPCI, especially those linked to the 2.16 Ethernet extensions, is rapidly gaining popularity and share of new system designs. These applications currently include airborne and vehicle based communications, multi-spectral imaging/IP encapsulation/dissemination, ship-board defense systems, air-borne data recording, remote radar control/imaging, and aircraft retrieval systems, to name a few. The benefits are clear and we will examine these in the upcoming section on 2.16.

Time is a key driver for today's developments and with this in mind it is clear that one can not disassociate the hardware platform choice from the software. In the past with many VME based solutions, a significant part of any software solution was "home grown" and rarely was there any "off the shelf" system or application management software available to reduce the burden on developers.

Software and hardware together make the end solution and must work in unison to achieve the optimal results. CompactPCI 2.16 has advantages in the area of platform management and many solutions offer integrated software environments that are essential to creating a cohesive and timely end product solution.

CompactPCI meets Ethernet – PICMG 2.16

As introduced in an earlier section, systems based on CompactPCI that implement the 2.16 extensions offer the greatest advantages in terms of flexibility and design independence. As we shall see, the primary reason for this is the reduction or elimination of the reliance on a bus-based backplane for inter-blade and data communication. Let's begin by discussing the fundamentals of CompactPCI 2.16.

The original goals and objectives for 2.16 were clear and simple. First, bus-based limitations that may hamper system design needed to be eliminated. Second, there was a real need to elevate systems design to the next level of performance while making it easy and quick to implement. The solution was to implement a packet-switched backplane.

Key requirements were to:

- Improve density
- Increase performance
- Leverage IP/Ethernet protocols
- Reduce integration time
- Increase system reliability
- Provide architectural scalability
- Leverage existing CompactPCI mechanicals

The fundamental concept of 2.16 is the creation of an intra-system Ethernet network. Utilizing a switched 10/100/1000 Ethernet based network topology, data can be independently moved between elements using high speed, independent links, obviating the need for a slower shared bus.

A 2.16 enabled backplane contains discrete embedded Ethernet infrastructure. The J3 connector is used to gain access and join the network. Due to the use of this separate connector, 2.16 and non-2.16 boards may participate in the same system. Up to two central switches link the individual slots. When configured with two switches, there is inherent redundancy built into the system as each slot is connected to both switches and the switches in turn are linked to each other. It can be seen, therefore, that 2.16 simplifies both the hardware and software requirements for delivering high available capabilities. A diagram of the internal network can be seen in figure 2.

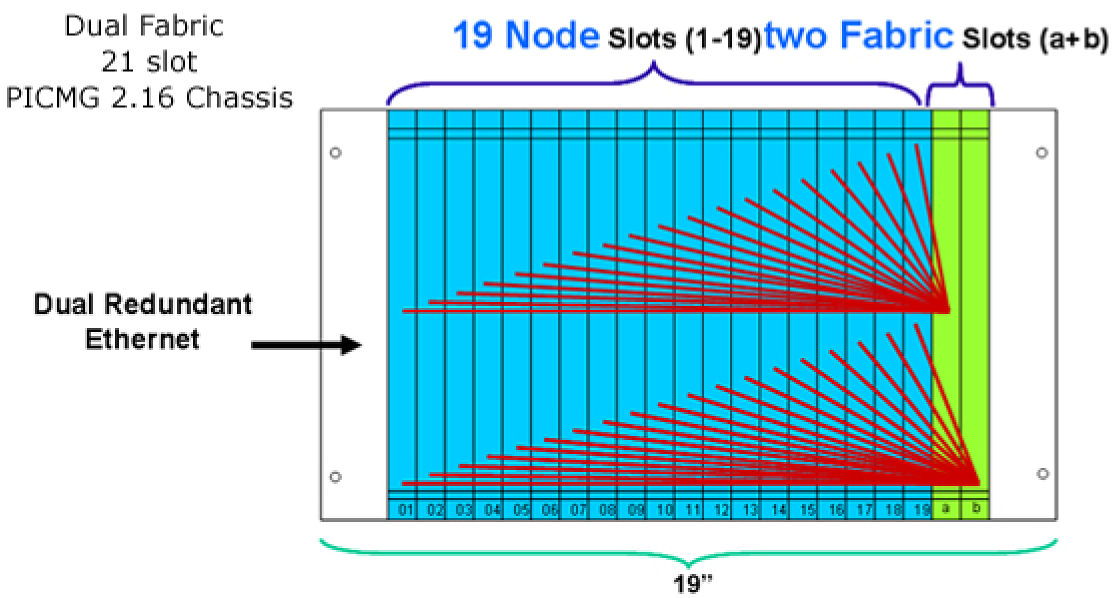

Figure 2: PICMG 2.16 Dual Star Topology

The base definitions for CompactPCI include the concepts of System and Peripheral slots. PICMG 2.16 adds to these with the Fabric and Node slots. All of these can be present in a system at the same time, if desired. The 2.16 specific Fabric and Node slots are described as follows:

Fabric slots – These slots contain the Ethernet switches. One or two Fabric slots can be supported in a 19-inch chassis and each one is capable of either routing or switching with bandwidths up to 96 Gb/s. The switched backplane is accessed through pins in the J3 and J5 connectors where 10, 100 or 1000 Mbit connections may be selected as desired. A Fabric slot can support up to 21 Node slots.

Node slots - As discussed above, a 19-inch chassis can support up to 21 Node slots. Each Node slots connects to both of the Fabric slots with a maximum available bandwidth of 4Gb/s.

Another key element of this architecture is that both Fabric and Node slots are hot-swappable, further enhancing a highly available configuration.

The ability for 2.16 and non-2.16 slots to coexist anywhere on the backplane makes it simple to begin to migrate older applications. Having overlain this high speed networking mechanism on CompactPCI, increased architectural freedom and flexibility comes into play. Since it is possible to have CompactPCI cards that can participate in both the regular PCI bus as well as the packet switched backplane, users can utilize stand alone “Blades” that may function on their own. Any 2.16 “Blade” that is inserted into the backplane or “network” can communicate to its partners through Ethernet without regard or care as to its processor, architecture software, or function. This freedom, to create, mix and match blades, capabilities and functionality within the same system is where the true power of PICMG 2.16 exists.

The capability for each blade to utilize its own independence within the system design means that changes in processor, memory configuration or even operating system from blade to blade has no regressive impact at the system level. As time and requirements change individual blades or subsystems may be upgraded for performance or functionality and simply be reinserted into the Ethernet-based hierarchy with no significant re-integration work required. Since the elements connect at the network level, not the application or operating system level, changes to the elements are far easier to manage.

While technologies such as VME have attempted to adapt by creating their own sub-bus or serial communication standards, this has only fragmented the market and limited the possible set of solutions. The simplicity and breadth of offerings for 2.16 is a clear winner. Another distinct advantage of the Ethernet-based system is the ability to extend the application outside of a single chassis. By bridging chassis using external Ethernet up to 10Gb/s, it becomes simple to create a distributed application whether for scale and extension or availability reasons. Disaster recovery options are broadened as such a distributed architecture does not care whether a backup node is in the next slot, next room or the other side of the country.

The latest set of specifications from PICMG define AdvancedTCA® (ATCA). Based on the networking technology of 2.16 CompactPCI, this is a much larger and power hungry form factor. A 10Gb option has been defined, although the overall ecosystem for this new technology is severely limited and it will be a number of years, at best, before it has a chance of filling out to meet the depth and breadth of the ecosystem for 2.16 products. If one has an application need that MUST have 10Gb links to individual processing elements within the chassis to be viable, then ATCA may hold the answer in the future. Today, no processing elements can fully utilize the bandwidth that a 10Gb pipe represents. For the vast majority of solutions, 4Gb/s or less is more than adequate. Compact, economical solutions can be created using existing, established and well proven CompactPCI 2.16 products.

PICMG 2.16 and the Triple Play Meet the Military

With Ethernet and IP based backbones now pervasive across all telecom and commercial communications networks, 2.16 technologies are ideally positioned. The continued development of solutions to take advantage of the “triple-play” that will deliver Voice, Data and Video across an all IP network is a current need for many users across all industries. Aerospace and defense providers have watched this trend closely and have clearly seen how “triple-play” solutions can benefit a multitude of applications from Command and Control all the way to the battlefield.

As far back as 2003 the Department of Defense (DoD) saw Ethernet as becoming a major component in their computing systems arsenal. They began to mandate the use and adoption of IPv6 with a goal to have all its network backbones fully capable by 2008. Notwithstanding the improvements for security that are topmost among all government agencies, the ability to create a seamless control and communications plane from the battlefield all the way to the Pentagon is a key driver for adoption of an extensible distributed Ethernet-based architecture.

As an example, let's take a look at the U.S. Air Force RQ-4A/Global Hawk program. This aircraft is an extremely high altitude, autonomous, Unmanned Aerial Vehicle (UAV). It provides near real-time, high resolution images and intelligence to field commanders.

The aircraft utilizes a CPC6600 2.16 switch from Performance Technologies as the basic networking interconnect between all communication systems and the platform sensors. These elements include VME-based sub-systems, packetized radios, and satellite links. The CPC6600 switches packets from multi-spectal imagers and other sensors, as well as tying together different systems on the plane.

Given Ethernet's pervasiveness, it was a relatively easy task to tie these network nodes together, as well as utilize well known networking protocols like multicast to enable image streams to be subscribed to by a general in a war room or by a "man on horseback" in the next valley.

Compact 2.16 technology is rapidly being adopted by programs and updates to existing platforms as diverse as the U.S. Navy's SSDS, the U.S. Marines' EFV/AAAV, FAA radar systems, the Royal Navy's SSN Submarines, U.S. Navy's EA-6B Prowler, National Weather Service weather radar, and the U.S. Air Force's E4-B.

In each case, 2.16 based systems are fulfilling critical aerospace and defense needs, including providing on-board LAN communications, data storage, sensor arrays integration, and packet processing.

Rugged Designs, Remote Management

It is easy to see why aerospace and defense are growing markets for CompactPCI. Equipment vendors need to ship equipment that is tough and reliable, built to withstand extremes in temperature, shock and vibration.

Due to its original market focus, just about all CompactPCI 2.16 products are built to satisfy the rigorous requirements of the NEB Level 3 Telecom standards. This means that just about every board can, in fact, handle the MIL-Spec requirements dictated by government programs. Even if they don't meet the requirements at the outset, most can be readily modified to do so. Conformally coated, ruggedized and conduction cooled CompactPCI blades and systems have been deployed for everyday use worldwide by prime contactors and governmental users.

NEBS Level 3 has strict specifications for fire suppression, thermal margin testing, vibration resistance (earthquakes), airflow patterns, acoustic limits, failover and partial operational requirements (such as chassis fan failures), failure severity levels, RF emissions and tolerances, and testing/certification requirements. Many times, these requirements meet or exceed the comparable MIL-Specs.

Since they deal with a large number of remote locations, military and defense system managers must be able to remotely manage that equipment for software upgrades, alarm conditions and system control. Thanks to network-based blades and system management modules common in CompactPCI 2.16 systems, they can easily do that.

Network based systems often use commercial packages based on Simple Network Management Protocol (SNMP), File Transfer Protocol (FTP), or its secure cousin SFTP, to remotely control and update individual system elements via an open or secure network.

Similarly, the chassis itself can be completely controlled via an out-out-band connection to the PICMG 2.9 Intelligent Platform Management Interface (IPMI) controller or set of redundant controllers, which act as the single management entity responsible for overseeing all of the boards, power supplies and fans within a system. The IPMI system manager can automatically react to stimuli such as an over temperature condition, a board watchdog event or power supply failure, while at the same time sending an alarm to a central site. Many 2.16-based systems are specifically architected to shunt system workload to an alternative resource using load balancing/cluster control software, either within the same chassis or to an external standby system.

Where We Are Now

A number of new technical approaches have been tried within VME to increase system performance in order to staunch its loss of market share, but as often happens introducing newer technologies, not all of the shrinking number of suppliers have adopted each of the technologies, leading to market fragmentation and loss of product ecosystem. Multiple choices have to be made between new quasi-proprietary buses and a number of networking topologies. In turn, that plethora of choices leads to a confused vendor and customer base and limits the number of suppliers that an integrator can call on for solutions using the particular topology he or she has chosen.

In contrast, systems based on Ethernet, Gigabit Ethernet and now 10Gb Ethernet have come to dominate CompactPCI, leading to a focused, mature ecosystem and broad choice of competitively priced vendors who offer an array of compatible blades. For example, there are now over 30 suppliers of single board computers for the CompactPCI 2.16 market. Integrators are able to choose from a wide selection of competing boards that cover everything from a low-cost Intel Pentium to a high performance, dual core x86 and PowerPC® solutions – along with a very impressive array of operating systems and applications for those blades.

The EcoSystem

Among the many vendors of CompactPCI products there are more than 200 2.16 vendors who offer an exhaustive range of systems, boards and blades. All functions required are available ranging from high end CPUs to communications specific network I/O cards, storage blades and DSP engines. While a system solution can easily be built from the combination of multiple vendor products, there are a few companies who offer a more holistic approach and can provide as complete 2.16 platform ready for application developers to begin work. Such an approach provides developers with a great "head start" over their competition, reducing the need for time consuming hardware integration. Software also plays a big part of any system and utilizing a pre-integrated platform that can be delivered with OS, HA/Fault Tolerant Middleware and networking installed and working is a huge advantage.

Performance Technologies, the inventor and leading developer of PICMG 2.16 products, is one such company. Along with a comprehensive portfolio of discrete CompactPCI products we offer a product line called the "Advanced Managed Platform™" a fully integrated and customizable solution.

Keeping up with the Joneses

Backwards compatibility is another major feature required by military system designers. The military has been using embedded systems for more than forty years. Unlike the PC market, where everything changes every nine months, maintaining and refreshing military systems (often based on custom designs with frozen technology) is a major challenge. The O/S and drivers used by all boards in a bus-based system must locked together, creating another impediment to keeping designs fresh (or being able to handle EOL notices). Bus-based systems have an inherently difficult time of being able to integrate the latest technologies and techniques to adjust to user needs.

Using 2.16 in COTS systems changes that paradigm. A switched packet backplane approach grants a much finer granularity to system changes, since the now independent system nodes use the lingua franca of Ethernet to communicate at the network level, severing the master-slave co-dependency of bus based systems. In this new scenario, the user simply changes one single board computer or communication node and puts in a higher performance unit to upgrade the whole system, greatly lessening the impact of obsolescence and upgrades.

CompactPCI 2.16 represents the best, balanced solution, because regardless of individual connection speeds, the fundamental technologies, topologies and protocols all work together. It is a standard that is evolving over time, but is inherently is backwards compatible. Elements of switched Ethernet backplane systems can all talk at different speeds (10/100/1000 Mbps), but all speak the same language. That's why Ethernet has come to dominate modern networked world, and is in use in over ninety five percent of networks deployed today, driving network hardware and training costs down while adding capabilities on a continuous basis.

Real World Solutions

Performance Technologies' portfolio of 2.16 compliant products encompasses over 20 families and includes complete systems, chassis, systems management, Ethernet switching, single board computers, packet processors, as well as specialized telephony and serial I/O, communications and storage blades.

Performance Technologies' family of embedded products has the most innovative and comprehensive set of hardware, software and solutions offerings in the market. With its unified set of products - OEMs, integrators, and service providers receive highly available and fully-managed systems with time-to-market, performance and cost advantages.

As mentioned earlier, it takes more than just best-of-class hardware to build an end solution and the integration of hardware and software is critical to meeting time to market goals. Enabling such timely and critical development is Performance Technologies' NexusWare® suite - a highly intergrated, comprehensive suite of Linux® OS, development tools and management environment, that is POSIX compliant and CGL registered.

Central to the NexusWare suite of software is NexusWare Core, which enables system engineers using Performance Technologies' embedded products to build packet-based systems. It provides the foundation environment and tools required for leading edge developments. This greatly accelerates the inclusion of both standard Linux packages and Performance Technologies' existing and future value-add NexusWare software packages, including NexusWare C7 (MTP-2), NexusWare WAN (Radar, TADIL-B, HDLC, X.25, Frame Relay…), NexusWare SIP, and NexusWare Voice Processing.

Highlights include:

- A full Linux operating system
- Kernel specifically tailored for embedded applications
- Complete suite of development tools & compilers, including the Eclipse platform
- Powerful APIs for all onboard hardware resources
- Integrated drivers - no need for external bus drivers
- Application development with Windows® XP or Linux
- Pre-packaged applications and protocols
- IPv6 Support
- Support for NexusWare installable software packages

NexusWare Core features an integrated suite of field tested device drivers that provide full access to onboard resources (H.110, Serial, Ethernet, PCI, T1/E1/J1, T3, CAS) as well as a powerful NexusWare API. As a result, products offered by Performance Technologies, with NexusWare Core integrated drivers communicating over TCP/IP, operate within any host system that supports TCP/IP - without the need of bus-based device drivers. The NexusWare Core API provides a common interface for application development and allows system integration using either the CompactPCI bus or the PICMG 2.16 architecture. Additionally, NexusWare Core allows application development to occur across several different operating systems, including Linux and Windows XP, translating to greater flexibility for developers.

Rounding out the capabilities of NexusWare is its ability to be installed onto either x.86 based or PPC based SBC's, providing the same capabilities as any Linux OS. Coupling these capabilities with its extensive API and development environment greatly reduces development time for applications utilizing I/O and storage modules. The power of NexusWare is its ability to provide a not only a common OS but also a Development environment spanning all I/O, Storage and Compute sub-systems offered by Performance Technologies.

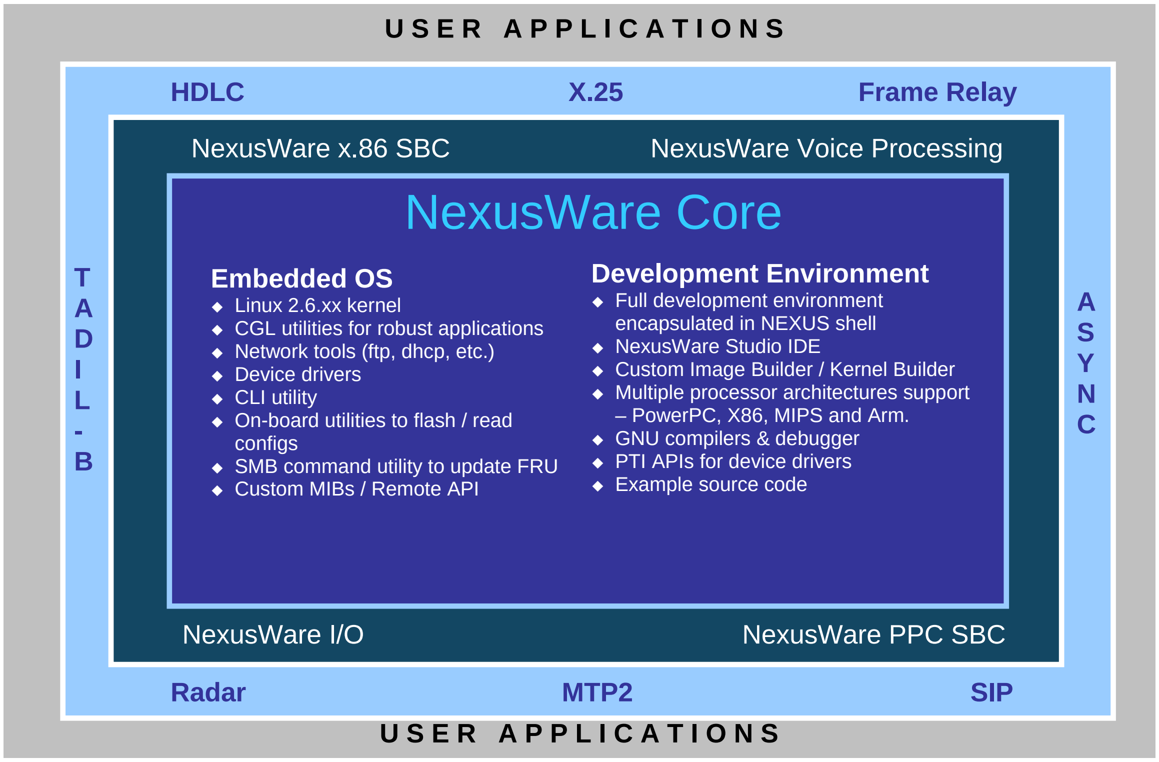

Where Do We Go From Here?

Industry analysts support the idea that use of CompactPCI 2.16 systems will expand, even as the new AdvancedTCA and MicroTCA specifications become more widely accepted. For example, IMS Research forecasts that further penetration into applications including aerospace and defense will cause the CompactPCI 2.16 market to experience strong growth over the coming years. Although VME still has its place in legacy systems, and AdvancedTCA solutions are an excellent choice for core telecom or other high-performance applications, growth in the CompactPCI market will continue for more and more applications that do not require the outright performance– and associated price, power consumption and footprint – of AdvancedTCA.

While nobody knows exactly what the future may be, the hold that Ethernet has on the future of computing and communications is set in stone. Probably more so than any other technology Ethernet is truly ubiquitous. The creation of the PICMG 2.16 extension that links Ethernet to the base bus standard has created an environment with immense power and flexibility. Global economic and business trends continue to place pressure on all industries and the move towards a network based, all IP infrastructure for aerospace and defense systems makes the advantages that can be gained by adopting CompactPCI 2.16 evident. Ethernet and IP are here to stay and those that can most easily and quickly harness the power of this combination will be the long term winners.

Performance Technologies is the originator of the CompactPCI PICMG 2.x and 2.16 standards and offers a host of compatible products and supporting software. For more information please visit www.pt.com.