CompactTCA and AdvancedTCA – Complementary solutions for high-availability applications

By Dick Somes

The development of CompactPCI has been incremental, starting in 1994 as an industrial packaging scheme for bused PCI systems. Starting with PICMG 2.5, which added support for the H.110 TDM bus, CompactPCI has evolved into an architecture with a menu of optional management and interconnect options suitable for a wide variety of applications. Many of these optional features are intended to tailor the CompactPCI platform for telecommunications use.

But this incremental process has also demonstrated the confusion that results when a large number of optional features are added to an architecture on an ad hoc basis. As telecommunication system architecture began the transition from bused to crossbar switched control planes, it became increasingly clear that a coherent vision would be essential for the next-generation of platforms. As a result the PICMG 3.0 base specification, and its subsidiary link specifications, were created in a hierarchical structure under the name AdvancedTCA which reflects this change of philosophy.

Recognizing the evolution of AdvancedTCA from CompactPCI, and the complementary relationship between the two architectures, the CompactTCA initiative was begun in 2003 by a small group of companies and handed over to PICMG (see Figure 1). The sponsors of the initiative are established suppliers of CompactPCI platforms, who have also invested heavily in the development of AdvancedTCA and are committed to the long term success of both.

It is intent of the CompactTCA initiative to formalize a set of compatible practices for managed PICMG 2.16 platforms limiting implementation options so as to enhance interoperability. These platforms are intended to support the portability and scalability of telecommunication applications, as well as other application classes, between this new platform and PICMG 3.0 compliant AdvancedTCA systems.

Realizing the inherent value a new specification with immediate product support, it is the group's intent to retain backwards compatibility with existing PICMG 2.16 board level products (see Figure 2).

Eliminating support for PCIbus signaling on the CompactTCA backplane is consistent with the goal of positioning this new PICMG 2.x specification with respect to PICMG 3.0.

Common requirements

The CompactTCA platform is intended to meet many of the same requirements as AdvancedTCA:

- A Base Interconnect for IP based control plane traffic
- A variety of data plane options

Figure 1

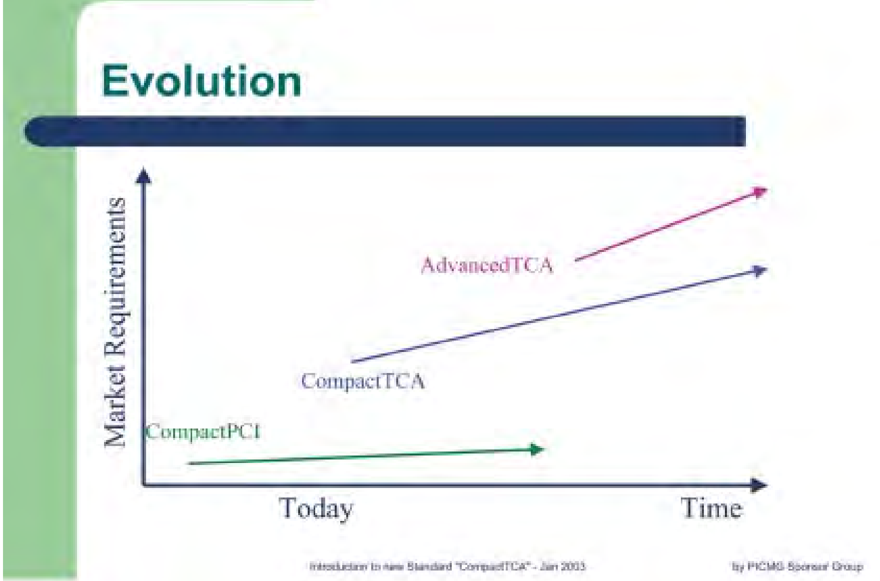

Figure 2

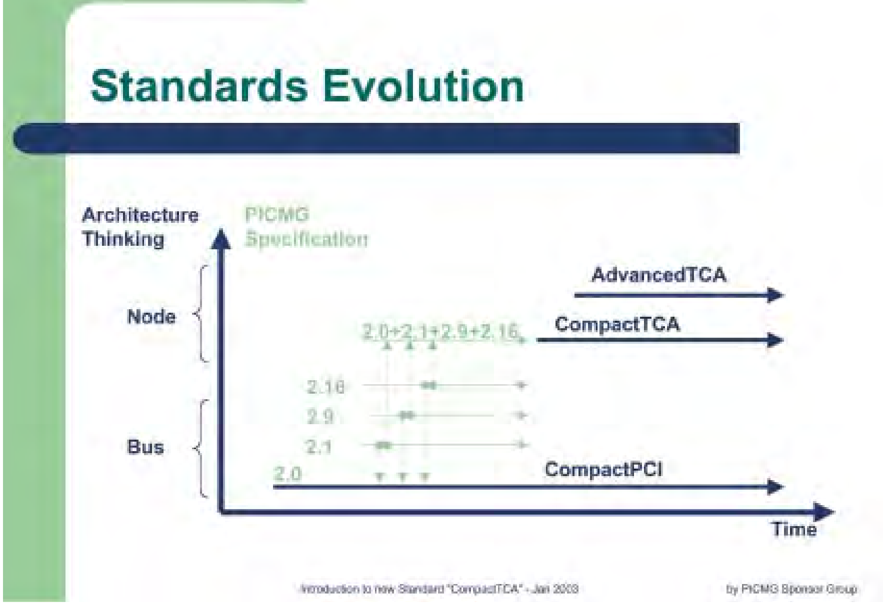

Reprinted from ***CompactPCI Systems*** */ May-June 2003*

- An shelf management infrastructure capable of remote visibility and dynamic allocation of platform resources
- A mechanical, electrical, and interconnect environment suitable for telecommunications applications

Table 1 shows how these high level requirements are satisfied.

Requirement	AdvancedTCA	CompactTCA
IP Control Plane	10/100/1000 BASE-T	PICMG 2.16
Data Planes	PICMG 3.1 Ethernet/FC PICMG 3.2 Infiniband PICMG 3.3 StarFabric PICMG 3.4 PCI Express	PICMG 2.5 H.110 PICMG 2.17 StarFabric PICMG 2.18 Serial RIO PICMG 2.20 CSMB
Shelf Management	IPMI V1.5 with extensions	PICMG 2.9 ECN 001 with selected PICMG 3.0 enhancements
Package	6U by 160mm boards 0.8 inch slot pitch 21 slots in IEC 60297 subrack (450mm aperture) 35 to 50 watts per slot	8U by 280mm boards 1.2 inch slot pitch 16 slots in ETSI subrack (500mm aperture) 150+ watts per slot

Table 1

Differences

CompactTCA and AdvancedTCA take different approaches to power distribution, with CompactTCA providing an environment equally hospitable to AC and DC power input The traditional logic level power supplies, 5 and 3.3 VDC, are produced by shelf level regulators and are distributed within the shelf.

AdvancedTCA is optimized for redundant DC power feeds, but supports frame level AC power as well. This simplifies shelf level power conditioning, but places an additional responsibility on individual boards for producing logic level voltages.

The two architectures are optimized for different levels of functional density as well.

The AdvancedTCA shelf typically ranges from 12U to 14U in height with 14 to 16 slots each. This size shelf can stacked three high in a typical 42U frame. With a dual P4 processor in each board, it's possible to deploy 84 CPUs in a frame.

The CompactTCA shelf, typically 9U in height, has 14 slots of which 12 slots are usable for compute or I/O functions. With four shelves stacked and rear to rear (eight per shelf) up to 112 CPUs can be packaged in a frame.

The option for back to back stacking results when shelves are configured with front I/O only, and package in shelves with a maximum depth of 300mm. This option is not available with AdvancedTCA shelves, which require the full 600mm depth of a standard frame.

While the size of the CompactTCA board seems to provide greater functional density as measured by the raw number of processors that can be packaged in a volume, the AdvancedTCA form factor is better suited for higher powered more complex multiprocessing nodes (see Figure 3).

Figure 3

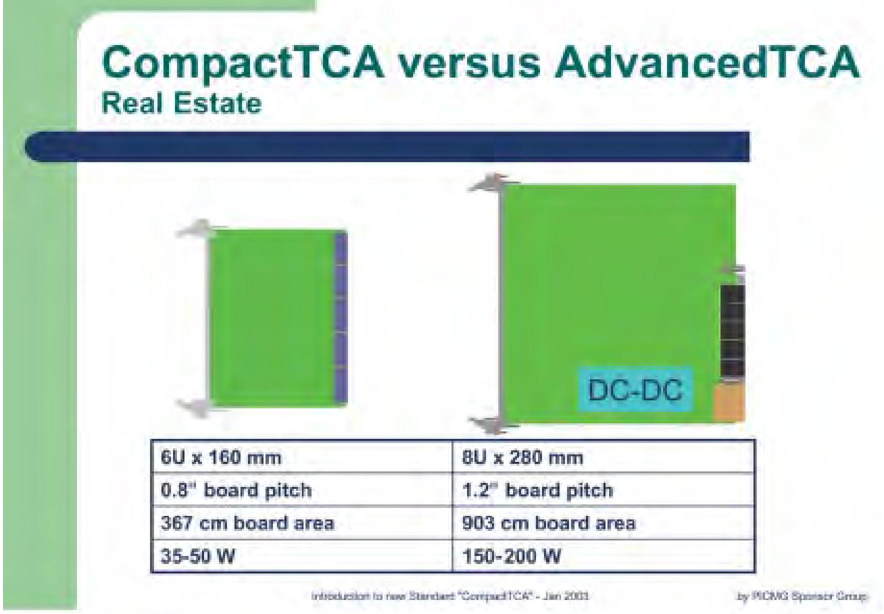

CompactTCA versus AdvancedTCA

Real Estate

6U x 160 mm	8U x 280 mm
0.8" board pitch	1.2" board pitch
367 cm board area	903 cm board area
35-50 W	150-200 W

Introduction to new Standard "CompactTCA" - Jan 2003

by PICMG Sponsor Group

Performance

Both the CompactTCA and AdvancedTCA platforms use switched Gigabit Ethernet to provide base interconnect capability, and therefore provide similar capabilities for applications which require only this level of performance (see Figure 4). There should be transparent portability between the two for such applications. The dual-dual star topology of Gigabit Ethernet links provides an aggregate of about 38 Gbits/sec assuming full-duplex usage of both switch fabrics.

Both platforms also provide for fabric interconnects that augment communications between nodes, but these capabilities clearly differentiate the two architectures.

CompactTCA provides support for bused H.110 connectivity between nodes at an aggregate of about 250 Mbits/sec, which is modest indeed. This connectivity still provides dedicated bandwidth between individual nodes, and is therefore logically a fabric.

CompactTCA also provides for the use of StarFabric, as defined in PICMG 2.17, concurrent with the Ethernet base interconnect. Using 622 Mbits/sec link technology, StarFabric is capable of roughly two times the performance of the base interconnect, or about 80 Gbits/sec aggregate, in support of transparent PCI to PCI bridging or higher throughput H.110 traffic.

The recently announced PICMG 2.18 initiative, which maps Serial RapidIO onto the CompactTCA platform, offers another factor of two increase in performance (160 Gbits/sec aggregate) over StarFabric.

All of three of the CompactTCA fabric interconnects discussed so far use the familiar 2mm HM connector in all five interconnect zones of the backplane. PICMG 2.20 introduces a version of the higher performance ZD connector to provide full mesh connectivity across the backplane at a 2.5 Gbits/sec rate. Across 19 slots this results in an aggregate throughput of 855 Gbits/sec (see Table 2).

Technology	Topology	Slots	Links/slot	BW/link	Aggregate
GigE	Dual Star	19	2	1	38
StarFabric	Dual Star	17	2	2.5	85
Serial RapidIO	Dual Star	17	2	5	170
PICMG 2.20	Mesh	19	18	2.5	855
PICMG 3.0	Mesh	16	15	10	2400

Table 2

AdvancedTCA, which provides for full mesh connectivity across 16 slots at 10 Gbits/sec, delivers 2400 Gbits/sec.

Figure 4

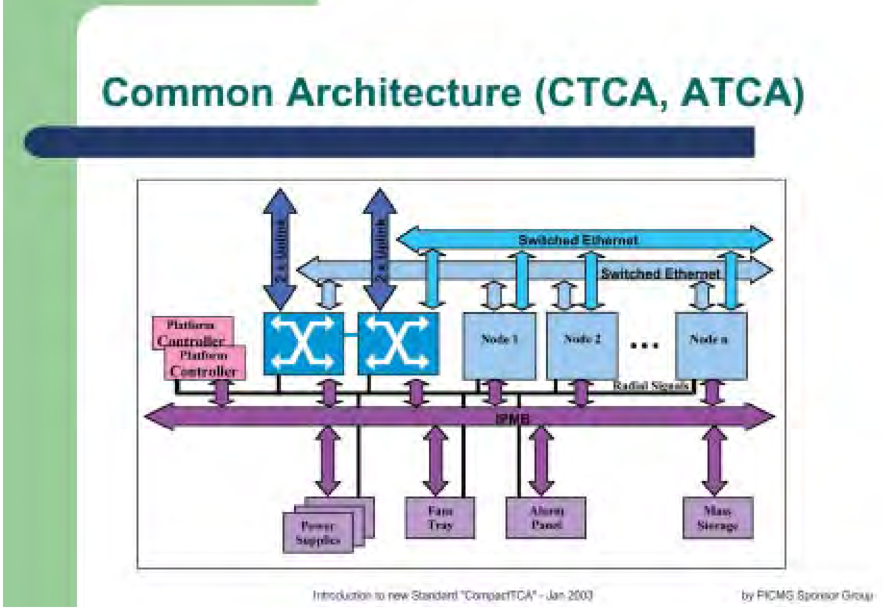

Introduction to new Standard "CompactTCA" - Jan 2003

by PICMG Sponsor Group

Complementary solutions for high-availability applications

AdvancedTCA and CompactTCA have much in common. Both platforms eliminate the dependence on high performance bused interconnects in favor of switched serial connectivity. Both have mandatory IPMI based management infrastructures providing transparent remote visibility and control of platform resources. Both provide electromechanical environments hospitable to high-availability operating systems, middleware, and applications.

But there is also scalability across this TCA family, which allows an application developer to optimize his solution while preserving a migration path for the future.

CompactTCA offers a cost optimized platform at the low end of the performance range. It uses the tried and true packaging and connector technology of CompactPCI, and takes advantage of an existing supply of PICMG 2.16 components.

CompactTCA provides a traditional power distribution with logic voltages on the backplane. It is equally convenient to condition these voltages from AC or DC input. The thermal capacity of the CompactTCA platform is in the range of 35 to 50 watts per slot, when cooled with forced air.

In terms of physical space consumed, the CompactTCA shelf is traditional. Vertical heights of 9U to 12U are typical depending on the strategy for packaging power conditioning and storage. The shelf is consistent with mounting in 300mm of depth when no rear I/O is used. It also offers a flexible array of rear I/O options.

Aggregate I/O performance of the CompactTCA platform depends on whether the traditional PICMG 2.x connector system is used in the backplane. When this is the case, backplane topology is limited to dual-dual star and there is little overlap with AdvancedTCA.

When the PICMG 2.20 enhancements are used to enable full mesh connectivity, the overlap in performance with AdvancedTCA is somewhat greater, but with some impact on backward compatibility with the PICMG 2.x legacy.

AdvancedTCA offers a platform optimized for the higher power density, complexity and aggregate I/O required by the next generation of general purpose multiprocessor nodes as well as arrays of network and digital signal processing functions. The area of the AdvancedTCA board approximates that of an AT form factor motherboard, making it easier to reuse desktop and server designs. While the minimum sized shelf is larger, 12U high and requiring 600mm of depth, it eliminates the need for active front end power conditioning within the shelf. The approach to rear I/O is even more flexible than CompactTCA, allowing the board designer to arbitrarily select connectors and providing a path for the introduction of fiber optic connectivity (see Figure 5).

The distribution of bulk power at a nominal –48 VDC recognizes the obsolescence of the traditional 5 and 3.3 volt logic supplies, and the increasing necessity of onboard DC-DC conversion.

AdvancedTCA provides the higher density thermal environment, supporting approximately 200 watts per function and 3.2 KW per shelf, required by higher performance functions. It also offers aggregate I/O performance of 2.4 Terabits/second for the current generation of transceiver technology, with the promise of 2x and perhaps even 4x enhancements as transceivers improve.

But in the final analysis, we must recognize the complementary relationship between CompactTCA and AdvancedTCA. PICMG 2.16 inspires the AdvancedTCA Base Interconnect, and PICMG 2.16 support is a mandatory feature of CompactTCA. The IPMI management infrastructure defined as a CompactPCI feature in PICMG 2.9 inspires the management architecture AdvancedTCA, and it is also a mandatory feature of CompactTCA. Both architectures support first generation StarFabric, and it is likely that Serial RapidIO and other fabric interconnects will be mapped onto both.

With this commonality of management infrastructure, base connectivity, and a growing number of common fabric interconnects, both architectures support many of the same software applications. For such applications, integrators can choose a CompactTCA platform for first deployment with confidence that more powerful compute engines with higher I/O capability will soon be available on AdvancedTCA platforms.

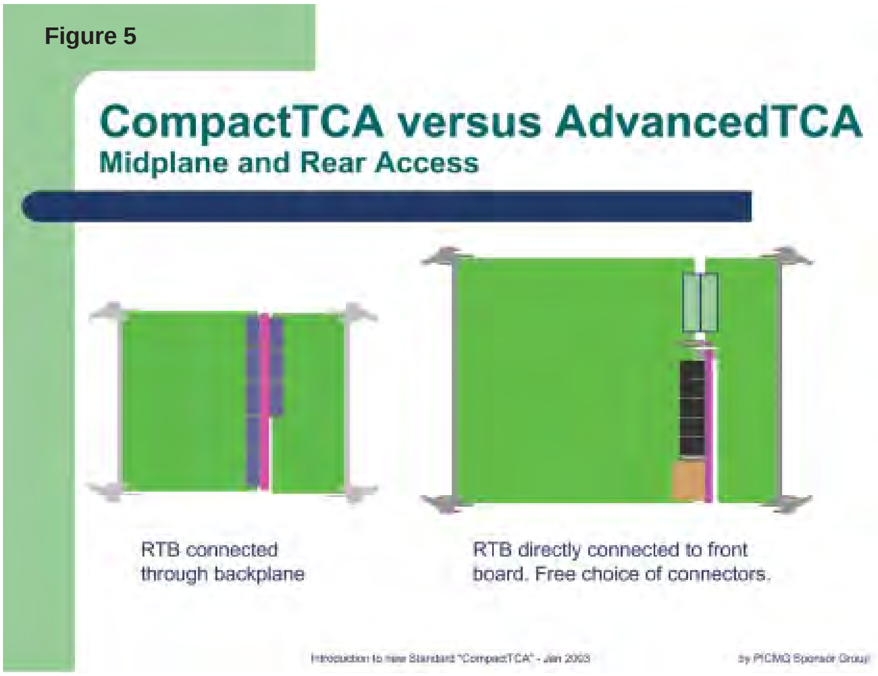

Figure 5

CompactTCA versus AdvancedTCA

Midplane and Rear Access

RTB connected through backplane

RTB directly connected to front board. Free choice of connectors.

Introduction to new Standard "CompactTCA" - Jan 2003

by PICMG Sponsor Group

Status of the CompactTCA Initiative

In January 2003 the sponsors of the CompactTCA Initiative announced their intent to develop a PICMG specification describing the architecture. The press release can be found on a number of Web sites including Force Computers Inc. at www.forcecomputers.com.

Motorola Computer Group, Performance Technologies, Pigeon Point Systems, StarGen, and Znyx are also mentioned, though there are more sponsors who chose not to be included.

The sponsors elected to disclose their intent to develop this specification around the same time as the AdvancedTCA specification was adopted by PICMG in order to emphasize their belief in the continuing value of the PICMG 2.x family of specifications in guiding the implementation of standards based telecommunications equipment.

By the time this article is published it is anticipated that the PICMG subcommittee will be formed, and that the preliminary specification will be under review. Completion of the specification is projected by the end of 2003 (see Figure 6).

Figure 6

Available CompactTCA Products — today

Note: All CompactTCA references are made to the current Sponsor Group's draft specification

Force's Centellis 25k & Centellis 21k systems with optional Intel, SPARC & PowerPC boards

Performance Technologies chassis with optional CPU boards, Line Cards, Ethernet Switches, and Intelligent Shelf Managers

Motorola Computer Group's MXP system with 2.16 PSB and optional 2.20 Distributed Serial Mesh

Introduction to new Standard "CompactTCA" - Jan 2003

by PICMG Sponsor Group