

PICMG[®] 3 -- Frequently Asked Questions

(Compiled by PICMG President Joe Pavlat: Pavlat@picmg.org)

What is PICMG 3?

PICMG 3 is a family of specifications, currently under development, that defines a new generation architecture for building high end, “carrier grade” equipment. The PICMG 3 specifications are oriented around switch fabric technology instead of a conventional parallel bus.

The specifications provide enough information to allow board, backplane, and chassis vendors to independently develop products that will be interoperable when integrated together. Details include board dimensions, equipment practice, connectors, power distribution, and a robust system management architecture that can be used independent of the switch fabric link technology.

What markets are targeted with the PICMG 3 family of specifications?

The PICMG 3 efforts are designed to provide an open, multi-vendor architecture that is aimed at Central Office telecom applications.

Unlike CompactPCI[®], which began as an industrial computer specification and later evolved to address a variety of telecom and communications markets, PICMG 3 is being written specifically to address the \$100B per year (and still growing) market for central office grade equipment. As “convergence” moves the architecture of core telephone switching equipment from circuit switched designs to packet switched designs, core telecom manufacturers such as Lucent and Nortel see a need for an open industry specification that maintains the rigorous standards they need, such as NEBS and ETSI compliance, high availability, robust system management, DC power distribution, etc., but is open and non-proprietary. By outsourcing more hardware in the future, these companies can concentrate on their core competencies, which include services and software.

While it is not a target market, there is a surprising amount of interest on the part of military equipment suppliers in PICMG 3, as the high bandwidth communications capability and extremely robust mechanical and electrical definitions are attractive for next generation military communications equipment.

While PICMG 3 provides unprecedented processor density and data bandwidth, the specifications are not targeted directly at classic “server farm” applications, which generally do not require central office-grade reliability and robustness.

What is PICMG 3.x?

PICMG 3.x is a term sometimes used to describe the family of specifications. It is largely interchangeable with PICMG 3 as a descriptor. The specifications under development now include:

- PICMG 3.0. This is the overall general specification that defines mechanics, board dimensions, power distribution, power and data connectors, and system management. This is a “fabric agnostic” specification that is intended to serve the needs of a variety of fabrics currently on the market.
- PICMG 3.1. This defines an Ethernet switch fabric over the generic backplane fabric interconnect. It provides for data rates up to 10 Gbit/sec. per link.
- PICMG 3.2. This spec defines how **InfiniBand®** systems are built within the architecture and will specify link physical layers, protocols, and protocol mappings.
- PICMG 3.3 defines a **StarFabric** implementation over the backplane providing TDM, cell, control, and packet connectivity over the same fabric.

Will the PICMG 3 family grow?

Yes. Any three PICMG Executive Members can propose a PICMG 3 subsidiary specification, and the PICMG Technical Officer is already receiving inquiries.

It seems reasonable to assume that mappings of Fibre Channel, 3GIO, and the cell based interconnect being proposed for PICMG 2.20 will also be developed for the PICMG 3.0 platform.

Will the same backplane work for different fabrics?

The PICMG 3 family of specifications uses the Tyco/ERNI ZD high speed differential connector for data transport. This connector is capable of data rates up to 5 Gbit/second per pair and is a good choice for Ethernet, InfiniBand, and StarFabric. A number of other high speed fabrics can also be mapped to PICMG 3 and the ZD connector.

A sophisticated electronic keying methodology is used as part of the PICMG system management infrastructure to determine what boards are plugged into the system and what fabric topologies and fabric technologies they support. In general, an entire shelf will be populated with boards of a single fabric technology. PICMG 3 defines 8 high speed differential pairs per data link, each pair capable of data transmission up to 5 gigabits/second.

Can proprietary technologies be deployed in the PICMG 3 environment?

Yes. Any transport technology compatible with PICMG 3 backplane topologies and signaling characteristics may be deployed. It is anticipated that the System Management infrastructure will recognize the presence of user defined links and configure according to

the same sort of interoperability rules used for those link technologies, which are explicitly supported.

What fabric topologies will be supported?

The standard PICMG 3 backplane is a “full mesh”, wherein every slot has a dedicated link to every other slot. This provides the potential, in aggregate, for terabit/second data transfers within a single shelf, or chassis.

Very high performance systems will use the full mesh for maximum throughput. The popular “dual star” and “single star” topologies are also supported in PICMG 3 on the same backplane, as a single or dual star is just a subset of a full mesh, with only a portion of the links used. Star topologies have a lower total bandwidth than a full mesh, but are less costly to implement.

What is “AdvancedTCA™”?

AdvancedTCA™ is the name PICMG has chosen to describe the new architecture. It stands for Advanced Telecom Computing Architecture, and it is pronounced “Advanced Tee-See-Aye.” This is analogous to the term “CompactPCI” being used to describe the PICMG 2.x family of specifications. PICMG is building a web site specifically devoted to the new architecture with descriptive information, white papers, contributions from member companies, plug-fest announcements, etc.

When will the specs be completed?

PICMG is currently on track to finish the core 3.0 specification in late summer, with 3.1 (Ethernet) to be completed shortly thereafter.

A “plug-fest” is scheduled for fall 2002 to begin testing first generation boards, backplanes, chassis, and software.

What size are the boards?

After lengthy deliberations, sophisticated thermal simulations, and a lot of customer feedback, PICMG 3 boards are 8U (322.25mm) high and 280mm deep. This size was carefully arrived at after considering cooling, front panel space, backplane size, and rear panel I/O requirements. Boards are spaced at a 1.2” (6HP) pitch. The wider pitch accommodates taller components like next generation CPU’s with integral heat sinks, off-the-shelf memory modules, and high power DC-DC converters. The wider pitch also improves cooling as more air volume can be circulated over a card.

This results in a board area of about 140 square inches, compared with about 54 square inches for 6U VME and CompactPCI. This additional real estate will be essential in the design of high performance computing blades with on-board fabric switches, network processors, general purpose processors, and memory.

PICMG 3 is based on the well-proven IEC60297 equipment practice, sometimes called the “inch based Eurocard standard.” This is the same equipment practice that is used in VME and CompactPCI systems. It is well understood, well tooled, relatively inexpensive, and is available from a wide variety of suppliers worldwide. As PICMG is planning on finishing the core specification this year, choosing an existing, proven equipment practice was important.

The 8U height was also chosen to optimize the use of a standard 42U (73.5”) rack because three 12U shelves (which include fans and cooling plenums above and below the card rack) fit into the 42U rack with about 6U left over for power conditioning circuitry and an alarm panel.

How many boards fit in a shelf?

Using the 1.2” board-to-board spacing, 16 boards can fit into a standard 600mm ETSI cabinet, which has a 500 mm frame aperture opening. Standard 19” racks, as defined by the IEC60297 family of specifications and used worldwide, accommodate 12 boards with a 450 mm Frame aperture opening.

Will Rear Panel Transition Modules be supported?

Yes. But the method being used is a little different – and much more flexible – than the technique used in CompactPCI, wherein front panel plug-in cards mate with rear panel I/O transition modules via double ended connectors that pass signals through the mid-plane.

The PICMG 3 approach still uses a rear panel transition module, but one that is 60mm in depth. The RTM is typically attached in a semi-permanent fashion in the rear of the chassis and mates directly with the front card. In order to do this, the vertical extent of the backplane ends just below the rear panel I/O area.

By mating the two cards directly together, a wide range of connectors can be used, depending upon the application. These include high density connectors like the 2mm family, controlled impedance or differential connectors for high speed signals, and blind-mate optical connectors for optical signals.

How much power can each board dissipate?

One of the basic limitations in CompactPCI is how much power can be dissipated per board. With the 2mm connectors and 0.8” board pitch used in CompactPCI, a practical upper limit of 50 watts per board using forced air cooling exists. Emerging high end microprocessors and support logic now dissipate as much as twice that amount.

The goal of the PICMG 3 family is that each and every board can dissipate 200 Watts. This adds up to over 3 kilowatts in a single chassis and almost 10 kilowatts in a standard

42U high rack. Extensive and detailed thermal analysis indicates that this power level is achievable, but 200W boards will likely require individual thermal design and simulation. A practical upper limit, without exotic thermal design and cooling, appears to be 150W using either high performance fans or blowers, even in the presence of a single fan failure.

How will power be distributed within a shelf and within the rack? What voltage will be used?

With power dissipations potentially being in the 3 kilowatt range for a single shelf, circulating logic voltage levels no longer makes any sense. Distributing traditional supply voltages such as 5 or 3.3V at 600 to 1000 amps is not practical. Moreover, the proliferation operating voltages within modules requires the distribution of a single higher voltage. Because PICMG 3 products are targeted at core central office applications, it has been decided to use the traditional dual, redundant -48VDC feeds available in central office locations. Dual -48VDC feeds will arrive at the rack to a power conditioning and distribution panel in the standard fashion, and then will be distributed to individual shelves. Each board will use local DC-DC converters (one or more) to provide logic level voltages for use on that board.

A sophisticated method of power management has been developed that provides capacitor pre-charge without glitching DC feeds. It also accommodates live insertion and removal (Hot Swap) and provides early System Management power after a card has been inserted.

Will the PCI bus be supported?

No, not directly. Switched fabric technologies like StarFabric that convert a high speed serial fabric to conventional parallel PCI can be used to create local PCI buses as needed, but there is no PCI on the backplane.

Is H.110 supported?

No, not directly, but some fabric technologies like StarFabric also support TDM traffic providing a port for bridging from a board level H.100/H.110 compatible bus.

The only signals on the backplane are power pins, grounds, the data fabric itself, geographic addressing pins, and a system management bus that is common to all of the fabrics.

VITA is developing next generation packaging in the VITA 34 specification. It also uses an 8U card. Why two separate efforts?

VITA 34 is intended to be a general packaging standard for a wide variety of markets, including military, aerospace, industrial control, and scientific instrumentation. PICMG 3 is very focused on central office telecom and is optimizing the spec for that market.

Also, PICMG 3.0 will be completed this year, with a plug-fest scheduled for the fall. This means that the PICMG 3 effort must use an existing equipment practice that is well understood, tooled, and widely available. The IEC 60297 “inch-based Eurocard” standard used by PICMG 3 (and by VME and CompactPCI) meets those criteria. VITA 34 is defining an entirely new equipment practice, and it will take some time before customers adopt it, and packaging companies tool up and make components available in volume.

Also the VITA 34 board is 220mm deep, instead of PICMG 3’s 280mm. This difference is important for the target markets. VITA 34 is appropriate for aerospace applications with legacy packaging requirements – for example, ATR boxes. Emerging aerospace applications need the higher power and better cooling provided by VITA 34, but may not be able to use the larger boards defined for PICMG 3.

VITA 34 also supports liquid and spray cooling for high power applications. This is not acceptable for central office telecom environments and their short MTTR requirements.

###