

## PMC, PrPMC, and High-Speed XMCs

### Abstract


The XMC standard is the next evolutionary step in mezzanine technology, providing all of the flexibility that made PMC successful, maintaining backward compatibility with legacy PMC modules today. This application note traces the evolution and design goals of the XMC, including an overview of the VITA 42 standard.

### Evolution of the PMC

PCI Mezzanine Card (PMC) standards were created about 10 years ago as a compact, high-performance option for I/O expansion. Rugged mechanical mounting, a low profile footprint and a standard PCI interface have made PMCs a natural choice for expanding the I/O functionality of intelligent controllers and single board computers. Their small size also makes them ideal for understanding and introducing cutting-edge I/O prior to the increased investments required of larger, more integrated products.

About 5 years ago, as the size of processing cores continued to decrease with the integration of CPU, memory controllers and bridges into a single chip, these same attributes made the PMC form-factor a ready-made starting point for the development of the popular *VITA-32: Processor PMC* (PrPMC) standard developed by the VMEbus International Trade Association (VITA) Standards Organization. PrPMCs have been welcomed as a small, fast and flexible standards-based computing solution for custom and scalable telecom, industrial automation, medical imaging, and defense and aerospace applications.

Figure 1. Top View of XMC



PMCs and PrPMCs have exploded in popularity—evidence of this includes the shipment to date of about a half million embedded PMC sites on VME, CompactPCI, motherboards and custom designs, ensuring the success of the mezzanine market. These mezzanines have evolved along with the PCI bus through improvements in data width and clock speed. Of course, as application complexity increases, designers desire an extended standards-based evolutionary growth path for these well-accepted embedded building blocks. Many embedded architects believe Switched Fabric Mezzanine Cards (XMCs) define this next building block.

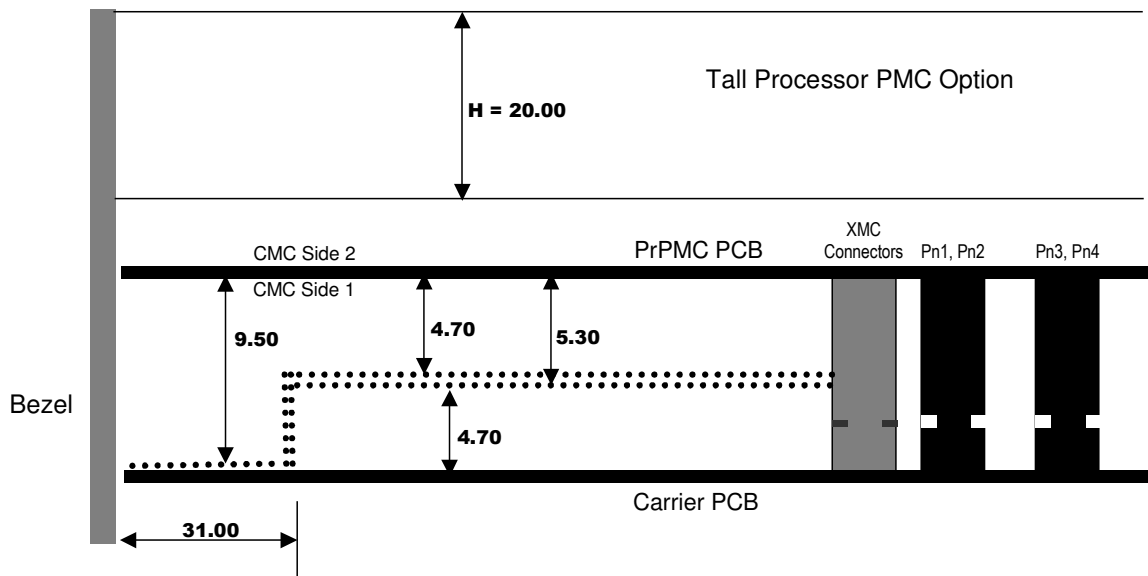
## Why XMC?

Embedded computing has always been characterized by a need for leading edge computing and I/O technology, often with severe constraints on environment, volume and power. As available processor throughputs and I/O bandwidths continue to march inexorably upward, interconnect standards are deploying to make use of improved interface technology. The drive to increased throughput, higher integration (system-on-chip) and higher functional density is currently being solved for the near term with the addition of PCI-X to PMCs and PrPMCs; refer to *VITA-39: PCI-X Auxiliary Standard for PMCs and PrPMCs*.

However, as throughput requirements continue to grow, traditional PCI-based products are being supplemented by switched fabric architectures due to restricted bus loading, routing and other issues associated and PCI associated with fast, wide buses. RapidIO, RACEway, StarFabric and Infiniband will increasingly be used for chip-to-chip local interconnects and/or slot-to-slot chassis-based interconnects in high-performance embedded systems.

XMCs build on the existing PMC standard by adding switched fabric interconnects to the existing PCI bus interface. Traditional PMC modules have two or three connectors (Pn1 through Pn3) that implement the 32 or 64-bit PCI bus interface and an optional fourth connector (Pn4) for user I/O. An XMC module adds one, or optionally two new connectors (Pn5 and Pn6) that are specifically chosen to support high-speed differential signals for fabrics or user I/O. These new connectors implement the XMC switched fabric interconnect independently of the existing PCI bus interface; they even include power, ground and basic control to allow an XMC card to function properly with just the fabric interface, and without the PCI connectors. Figures 1, 2 and 3 provide a better understanding of these updated mezzanines.

**Figure 2. Side View of XMC**




All units in millimeters. Not to scale. Refer to IEEE 1386 for details.

The key requirements in quickly creating a new standard include a dedicated, focused committee with a well-defined set of goals. The VITA 42 committee is indeed focused and moving quickly; the goals are summarized in Table 1. While each of the goals is important, the committee focuses primarily on adding two new high-speed connectors while simultaneously maintaining backward compatibility with standard PMCs and PrPMCs, including their conduction-cooled alternatives.

This group also understands the need to reduce the number of options in order to minimize confusion and maximize sales volume, thus driving down XMC costs. Finally, to speed acceptance and release of this standard, it has been important to leverage recent LVDS-compatible connector signal testing. A significant head start resulted from the collaboration between VITA and the RapidIO Trade Association (RTA). The extremely open discussions between these two groups and the RTA's detailed report summarizing connector evaluations resulted in rapid development of the initial VITA 42 draft standard, ultimately benefiting both organizations.

**Figure 3. 3D View of XMC**



### **PMC and XMC Compatibility**

Both XMC modules themselves and XMC carrier cards must be considered when discussing compatibility. Initial XMC carriers will likely include both PCI and high-speed fabric connections, allowing carriers to seamlessly support both traditional PMCs and the newer XMCs; refer to Table 2 for compatibility options. As the availability of high-speed XMCs increases, carriers will likely eliminate the PCI connectors and simply support only Pn5 and/or Pn6 mezzanines.

The PMC market evolution taught us that carrier sites are a key inducement driving mezzanine development. Since one site will support both PMCs and XMCs in the same physical real estate, carrier vendors and OEMs can economically enable XMCs. As performance and availability dictate, legacy PMCs and PrPMCs can mix and match with high-speed XMC I/O and processor modules, preserving compatibility while providing a migration path to switched fabric architectures.

As with PMC, the XMC standard supports (a) I/O and processor modules; (b) non-intelligent and intelligent carriers; (c) single and double-wide modules; (d) front panel or rear I/O connections; and (e) air cooling or conduction cooling as defined in *VITA 20: PMC Conduction Cooling* standard.

XMCs are defined by the *VITA 42:XMC Switched Mezzanine Card Auxiliary* series of standards. VITA 42.0 is the base document that defines the mechanical interfaces and connector layout, along with supplemental protocol standards that map specific switched fabric interconnects into the VITA 42 framework. The first two protocol standards are VITA 42.1 for Parallel RapidIO and VITA42.2 for Serial RapidIO. Future protocol standards will define new switched fabric interconnects as market requirements dictate.

It is expected that the very first XMC modules and carriers will use the Parallel RapidIO embedded switched fabric offering throughput rates from 1 to 4 Gbits/s initially and growing up to 5 Gbits/s with multi-channel serial RapidIO. RapidIO provides an open standard interconnect architected from the ground up to meet the latency, scalability and throughput requirements of high-performance embedded computing. As market demand grows, it is likely that PCI-Express mezzanine products are just around the corner.

One final question that is often asked concerns timeframes. The XMC standards committee has reviewed several drafts of the Base VITA 42.0 and it is moving quickly toward a task group vote and approval in the July timeframe. To build a product, of course, both the base standard and a protocol standard are required. Both of the RapidIO protocol standards will validate the structure details and interdependencies of this standard. The 42.1 standard is in review now with voting slated to complete in early August; the 42.2 serial standard is about a month behind that schedule since they share the same format.

**Table 1. VITA 42 XMC Goals**

<b>Goal</b>	<b>Benefit</b>
Add high-speed interface to PMC/PrPMC mezzanines	<ul style="list-style-type: none"> <li>• Add 2 new high-speed connectors</li> <li>• Support multiple high-speed serial fabrics</li> </ul>
PMC/PrPMC follow-on for 6U & custom applications	<ul style="list-style-type: none"> <li>• Re-use standard 75 mm x 150 mm form-factor, drawings, front panels, ....</li> </ul>
Maintain backward compatibility with traditional PMC sites	<ul style="list-style-type: none"> <li>• Facilitate adoption by minimizing cost</li> <li>• Carriers may support both PMC and high-speed XMC</li> <li>• Mezzanines are likely PMC or high-speed only; possible to support both</li> </ul>
VME/CompactPCI compatibility	<ul style="list-style-type: none"> <li>• Standard PMC height: Single card slot support</li> <li>• Optional Tall modules: custom &amp; multi-slot versions for increased thermal requirements</li> </ul>
Minimize proliferation of options	<ul style="list-style-type: none"> <li>• Facilitates adoption with focused products</li> <li>• Increase volumes to drive down costs</li> </ul>
VITA & RTA collaboration	<ul style="list-style-type: none"> <li>• Rapid head start with standard and connector</li> <li>• Whenever possible leverage compatibility with RTA development</li> </ul>
Conduction-cooled compatibility	<ul style="list-style-type: none"> <li>• Continue compatibility with VITA 20</li> </ul>

**Table 2. XMC Compatibility Summary**

<b>Configuration</b>	<b>Required</b>	<b>Optional</b>
PCI Only: PMC/PrPMC	Pn1, Pn2	Pn3, Pn4
High-speed Only: New XMC	Pn5	Pn6
High-speed and PCI: XMC including PMC/PrPMC	Pn1, Pn2, Pn5	Pn3, Pn4, Pn6