# OpenVPX:
# Architectures for High-Performance Embedded Computing

Robert Cooper (rcooper@mc.com)
Mercury Computer Systems, Inc.,
201 Riverneck Road, Chelmsford, MA 01824

Mark Littlefield (mark.littlefield@Curtisswright.com)
Curtiss-Wright Controls Embedded Computing
741-G Miller Drive, Leesburg, VA 20175

## Introduction

The high-performance embedded systems that are the most successful in the long term combine the latest, highest performance technology within open systems architectures. While any one system may take the lead in performance at a given time, the community benefits most when users can combine the best technologies from multiple sources to craft a solution to their challenging computing problem.

A tension exists between standards and technology. Users and vendors would like stable standards that allow multiple interoperable products to be built. But standards can stand in the way of innovation, both by remaining fixed on prior technologies, and by constraining system architectures in such a way that the latest technologies cannot be deployed in the most appropriate way.

## VPX Origins

The VME form factor, as defined by a number of VITA standard specifications, has served the embedded computing industry very well over many years. It has provided a common environment in which components from different vendors could be quickly assembled into systems for a large range of benign and rugged deployed environments. However, by the early 2000s the base technologies used in VME were beginning to show their limits. VME could not support the ever-increasing >1 Gbaud signaling that the industry was moving towards. VXS (VITA 41) extended VME by adding some high-speed signaling, but there was a desire within the industry to break with backward compatibility with the VME form factor, in order to dramatically increase the number of high-speed pins in the backplane. In addition, such a break allowed the standardization of air- and conduction-cooled variants, as well as the introduction of a 3U form factor, increased board pitches, and 2-level maintenance features sought after by lead system integrators. The desire, therefore, was to make a "VME infrastructure for the future," and thus VPX (VITA 46) was born.

In 2007, the VITA Standards Organization (VSO), and later ANSI, approved the base specification for VITA 46 [1]. In doing so, they created a wide-open playing field for much higher-throughput designs.

The result provides generous per-slot bandwidth, and unprecedented bandwidth versus area and volume. Table 1 compares the number of high-speed lanes and module sizes of 6U and 3U VPX with ATCA [2] and BladeCenter [3, 4].

| | 6U VPX | 3U VPX | ATCA | BladeCenter |
|---|---|---|---|---|
| **Lanes** | 96 | 32 | 100 | 20* |
| **Area ($cm^2$)** | 373 | 160 | 902 | 1093 |
| **Volume ($cm^2$)** | 759 | 325 | 2750 | 3169 |

*Designed for 10 Gbaud signals

**Table 1: Form factor comparison**

Subsequently, a number of "dot" specifications have defined implementations on this base specification such as VMEbus, Serial RapidIO, Ethernet, and PCI Express communications fabrics, and system management functions. However, as the first generation of VPX products were fielded, it became apparent that there were integration issues between products from different vendors. The flexibility introduced by the standard (for good, technical reasons) meant that different suppliers were implementing features in different ways, which caused problems for system integrators. The result was that while VPX was highly successful in enabling users to build complex, high-speed system designs, and in fostering a marketplace of multiple vendors, interoperability between vendors was lagging.

Thus, there is a desire to define a "system specification" that leverages the baseline specifications and applies the best practices of industry to define a common set of implementations, or profiles. This standards effort is being driven by an organization called the OpenVPX Industry Working Group.

## OpenVPX

The outcome of the OpenVPX activities will form a standards specification (VITA 65) that, when formally approved, will form a guide for both vendors and system integrators. The benefits include:

- Quick lab integration through the use of standardized, off-the-shelf chassis and backplanes
- Reasonable expectation of interoperability between boards
- The ability to replace hardware with upgraded modules, perhaps from competitors, with reasonably little effort

This will result in a reduction in the time needed to define a system concept, build the system, and verify it—resulting in lower costs and shorter time to deployment. In addition, true interoperability will provide users with a greater choice of off-the-shelf components and the benefits of second sourcing at the board level.

OpenVPX activities are divided into two phases. Phase one involves the definition of a system taxonomy and glossary to create a consistent language to express system-level concepts, the review of various system topologies adopted by the industry to-date, and a definition of a common "utility plane" comprising power and utility functionality.

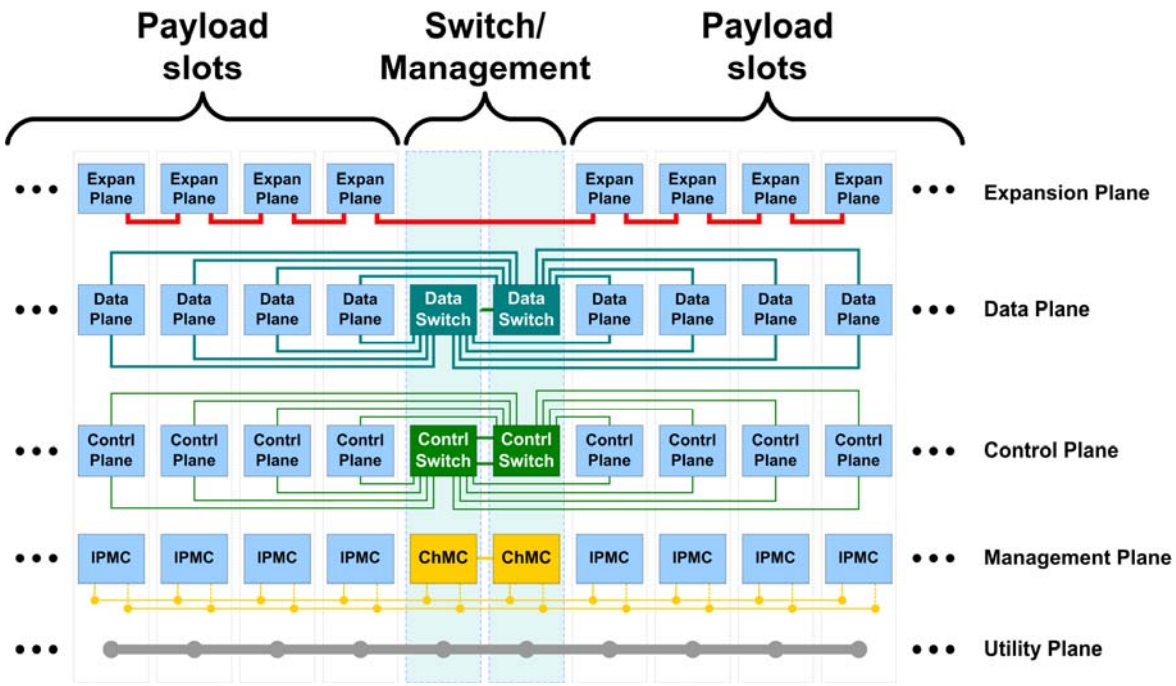


**Figure 1: Example OpenVPX topology made up of a number of common "planes"**

Building on this, phase two will distill a common set of system slot and connectivity profiles. Activities include defining common slot and backplane profiles, and a set of standard development chassis.

## Planes

A major organizing principle of the OpenVPX work has been the notion of logical communication or functional "planes" (Figure 1). This concept originates in the of management, control and data planes in the telecom world. OpenVPX has extended this notion into five planes:

- The utility plane provides basic services such as power, common reference clocks, and low-level signals such as system-wide reset.
- The management plane is used for monitoring and managing hardware resources based on the IPMI standards. The management plane can be powered up independently from main board power and thus can be used for enabling main power on each board.
- The control plane is intended for low latency, moderate bandwidth, packet-based communication. By separating control and data traffic, each can be guaranteed its own bandwidth, and congestion in one plane will not affect others. The control plane typically uses a switched Gigabit Ethernet arranged in a single- or dual-star topology.
- The data plane carries high-bandwidth, low-latency data traffic between boards and externally to the chassis. The data plane may utilize a choice of topologies such as dual- or single-star, single- or multi-hop mesh, etc. The data plane is where a real diversity of physical, logical, and network layer protocols may be used, including serial RapidIO or 10 Gigabit Ethernet.
- The expansion plane links a controlling host board with several tightly coupled adjunct processing or I/O boards. Often these adjunct boards cannot operate independently of the host. The expansion plane topology is typically a set of disjoint trees rooted at each host node. PCI Express or other very-high-speed protocols may be used. Like the data plane, the expansion plane requires high bandwidth and low latency.

## Profiles

The functional portion of the standard will be a set of backplane, module, and slot profiles. These profiles will describe a common set of pin definitions, communication protocols, mechanical form factors, and slot interconnection topologies. Suppliers of boards and backplanes will use these profiles to guide their design decisions. System integrators, on the other hand, can use these profiles in their system designs and requirements. Over time, an infrastructure will be built within the industry aligned with these profiles, giving system integrators a range of solutions.

The standards drafting activity in OpenVPX is still underway and is scheduled for completion in October 2009. The resulting standards language will be transferred to the VITA 65 Working Group within the VSO for final review and approval before passing it on to ANSI for ratification.

## Acknowledgements

OpenVPX/VITA 65 is the combined effort of representatives from more than 30 companies and organizations. For a complete list of members, as well as additional background on the organization, please visit http://www. vita.com/openvpx

## References


[1] VITA, “American National Standard for VPX Baseline Standard,” *ANSI/VITA 4.6.0-2007*. VMEbus International Trade Association, Oct. 2007.

[2] PICMG, “PICMG 3.0 Revision 2.0 AdvancedTCA Base Specification,” PCI Industrial Manufacturers Group, Mar. 2005.

[3] IBM, “IBM BladeCenter H I/O,” http://www-03.ibm.com/systems/de/resources/systems_bladecenter_pdf_BladeCenter_IO_vs_HP_01-15-07.pdf, Jan. 2007.

[4] IBM, “BladeCenter Blade Servers,” http://www-03.ibm.com/systems/bladecenter/hardware/servers/.