

# Ensemble® 3000 Series

## CCM3012 3U OpenVPX PMC/XMC Carrier Module

*Expanded I/O Capabilities for OpenVPX Platforms*

- PCIe Gen-3 interfaces to OpenVPX data plane
- Advanced system management built-in (VITA 46.11)
- Air, conduction, Air and Liquid Flow-By packages



The Ensemble® CCM3012 PMC/XMC carrier is a compact VITA 46 VPX OpenVPXTM 3U module. This carrier brings additional processing capabilities to compliment processor modules such as the Ensemble LDS3506 by providing high-bandwidth access to additional mezzanine sites. The Ensemble CCM3012 connects to processing modules via a Gen-3 PCIe interface, allowing data to flow into a system by a simple expansion of on-board PCIe trees. Available in air-cooled, Air Flow-By™, conduction-cooled and Liquid Flow-By™ options, the Ensemble CCM3012 module can be deployed in a variety of environments with confidence.

The Ensemble CCM3012 is supported by the rich set of features available from the MultiCore Plus® software infrastructure, which allows ease of portability, while offering open software development architecture.

### PCIe Gen-3 Interfaces

The Ensemble CCM3012 carrier provides a high-bandwidth PCIe interface to the OpenVPX backplane. The switch connects 16 lanes to the backplane which can be configured as a single x16, two x8s

or four x4s. The default is two x8s. In smaller systems, the Ensemble CCM3012 can provide PCIe switching via the on-board 32-port switch. These backplane PCIe interfaces allow a physically distinct processing module to expand its PCIe infrastructure to encompass the additional resources available on the Ensemble CCM3012. The PCIe switch connects an upstream x16 PCIe link to the backplane connector, a downstream x8 PCIe link to the XMC site and a x4 link to the PCI/PCI-X bridge for the PMC site.

### Mezzanine Card Flexibility

The standard PMC/XMC site on the Ensemble CCM3012 module can be configured with off-the-shelf mezzanine cards using either PCI-X or PCIe protocols. PMC cards are supported with a PCI/PCI-X interface at up to 133 MHz on each site. XMCs are supported with x8, x4, and x1 PCIe interfaces, linked via the J15 connector per the VITA 42.3 standard. The XMC connectors utilized on the Ensemble CCM3012 are VITA 61 compliant connectors, optimized for high-speed signal integrity as well as rugged operation. The PCIe interfaces are capable of Gen-1, Gen-2, and Gen-3 PCIe data rates. XMC user I/O is mapped to the backplane via the J16 connector, in accordance with the standard VITA 46.9 X12D+X8D pattern for differential signals. Additional single-ended XMC user I/O signals are also available.

*Mercury Systems is a leading commercial provider of secure processing subsystems designed and made in the USA. Optimized for customer and mission success, Mercury's solutions power a wide variety of critical defense and intelligence programs.*



ACQUIRE



DIGITIZE



PROCESS



STORAGE



EXPLOIT



DISSEMINATE

The Ensemble CCM3012 provides ample power to the PMC/XMC site, allowing the site to support high-powered mezzanines capable of drawing more than 30 watts. By separating these high-powered mezzanines from a module with on-board processing capabilities, the dissipation of thermal energy can be balanced at the system level, maintaining a high MTBF figure and removing the need for costly thermal management designs. The Ensemble CCM3012 supports these cost-saving considerations while simplifying software design, the Ensemble CCM3012 is configured as a simple logical extension of the support software on the neighboring compute module.

## System Management Plane

The Ensemble CCM3012 module follows OpenVPX design principles in leveraging the robust, scalable, and well-tested system management infrastructure from the VITA 46.11 architecture. Using the standard I2C bus and IPMI protocol, the on-board, system-management block implements the Intelligent Platform Management Controller (IPMC), in accordance with the VITA 46.11 standard. This allows for the Ensemble CCM3012 module to:

- Report sensor values
- Report and set sensor thresholds, allowing an application to react to voltage or current variations that exceed those thresholds
- Reset the entire module
- Power up/down the entire module
- Retrieve module FRU (field replaceable unit) information
- Be managed remotely by a Chassis Management Controller at the system level, such as implemented on the OpenVPX SFM3010 Module

## VPX-REDI

The VPX (VITA 46) standard defines 6U and 3U board formats with a modern high-performance connector set capable of supporting today's high-speed fabric interfaces, such as PCIe. VPX is most attractive when paired with the Ruggedized Enhanced Design Implementation standard – REDI (VITA 48). The Ensemble CCM3012 module is implemented as a 3U conduction-cooled implementation of VPX-REDI, with air-cooled variants in the same VPX form factor available for less rugged environments.

Targeted primarily for harsh-environment embedded applications, VPX-REDI offers extended mechanical configurations supporting higher functional density, such as two-level maintenance (2LM). 2LM allows maintenance personnel to replace a failed module and restore the system to an operational state in a limited time period, minimizing the potential of damage to the board.

## Open Software Environment

Mercury leverages over 30 years of multi-computing software expertise, including multicore processor expertise, across its many platforms. This strategy is fully applied to the Ensemble CCM3012 module and includes support for off-the-shelf open software products such as OFED and OpenMPI.

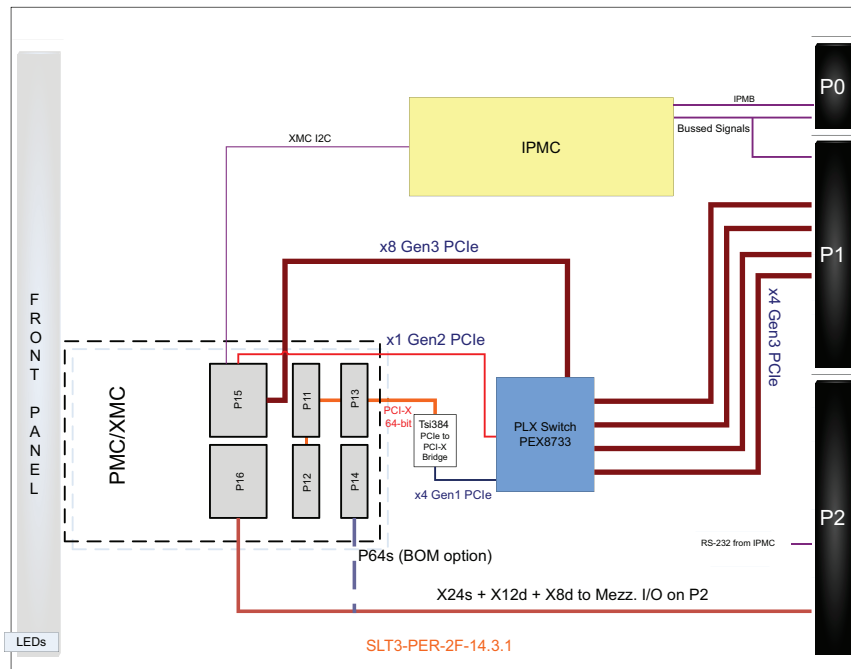


Figure 1 - CCM3012 functional block diagram

## Mercury Sensor Processing Ecosystem

Sensor chain awareness is having the technical expertise and resources to design and build capable, compatible solutions along the whole sensor processor chain. From RF/microwave, digital/analog signal manipulation to dense, size, weight, and power (SWaP) optimized processing resources to actionable intelligence dissemination; Mercury's rugged compute subassemblies leverage the best commercial-item technology, enabling prime contractors to win more business. Modern sensor compute subassemblies are customized assemblies of interoperable building blocks built to open standards. Mercury's hardware and software portfolio of building blocks are physically and electrically interoperable as defined by international industrial standards, including VITA's OpenVPX standards.

### Rugged air cooling, Air Flow-By

Air- and conduction-cooled subsystems rely on filtration to remove contaminants from their cooling air streams. Mercury's Air Flow-By™ technology eliminates filtration with the most elegant cooling solution available within a sealed and rugged package. Fully compliant to VITA standards (including VITA 48.7), Air Flow-By maintains OpenVPX's 1-inch pitch requirement, is highly resilient to liquid and particle contamination, boosts SWaP, reduces operating temperature, extends MTBF by an order of magnitude and enables embedded deployment of the most powerful and reliable processing solutions.

### System bandwidth – the effect of interconnect performance

Bandwidth is critical and especially applicable to switch fabric resources which dictate the responsiveness of the entire subsystem. All modular solutions, including OpenVPX, are to some degree interconnect-bound; they are restricted by the limitations of the best industry interconnect technology. To address this, Mercury fabricates system interconnections with innovative technology which mitigates insertion-loss and cross-talk while maintaining full OpenVPX compliance. The result and significant system performance boost is especially discernible as subsystems scale larger and data rates increase.

## Specifications

### Module Specifications

- One PMC/XMC site
- On-board 32-lane Gen-3 PCIe switch
- Backplane PCIe interface to data plane
- System management capabilities in accordance with VITA 46.11
- Air-cooled, conduction-cooled, Air Flow-By and Liquid Flow-By packages

### PMC-X/XMC Sites

- PMC-X sites 1
- PCI-X-to-PCIe bridge - Connects PMC sites to on-board PCIe switch
- PCI support 33 and 66 MHz
- PCI-X support 66, 100, and 133 MHz
- PCIe XMC sites per VITA 42.3
- Supports x8, x4, or x1 PCIe interfaces from XMC to on-board PCI Express switch
- XMC user-defined I/O from J16 to backplane
- Differential signals per VITA 46.9 X12D+X8D pattern
- X24s signals VITA-46.9 mapping from XMC J16 site to VPX P2

### Management plane

- I2C interface to backplane
- On-board IPMC
- Voltage, current, and temperature sensors

### I/O

- RS-232 serial interface from IPMC to backplane interface
- Support for PCIe resets per OpenVPX standard

### Specification Compliance

- OpenVPX System Specification (VITA 65)
- VITA 46.0, 46.3, 46.4, 46.6, 46.11, and VITA 48.1, 48.2 (REDI)

For information on cooling options, see our [MOTS+ extreme environmental protections](#)

**Reviewed and approved for public release by Mercury's Export Compliance Office. Dated 05/09/2016, reference number 16-003**

The Mercury Systems logo is a registered trademark of Mercury Systems, Inc. Other marks used herein may be trademarks or registered trademarks of their respective holders. Mercury products identified in this document conform with the specifications and standards described herein. Conformance to any such standards is based solely on Mercury's internal processes and methods. The information contained in this document is subject to change at any time without notice.

Copyright © 2016 Mercury Systems, Inc.

3167.02E-1218-ds-CCM3012



INNOVATION THAT MATTERS™

#### CORPORATE HEADQUARTERS

50 Minuteman Road • Andover, MA 01810 USA  
(978) 967-1401 • (866) 627-6951 • Fax (978) 256-3599

#### EUROPE MERCURY SYSTEMS, LTD

Unit 1 - Easter Park, Benyon Road, Silchester, Reading  
RG7 2PQ United Kingdom  
+ 44 0 1189 702050 • Fax + 44 0 1189 702321

