RapidIO in VME Systems

www.vmetro.com

CURTISS WRIGHT Controls Embedded Computing

VMETRO a Curtiss-Wright Company

Curtiss-Wright VMETRO Overview

Curtiss – Wright: The Foundation of Aviation

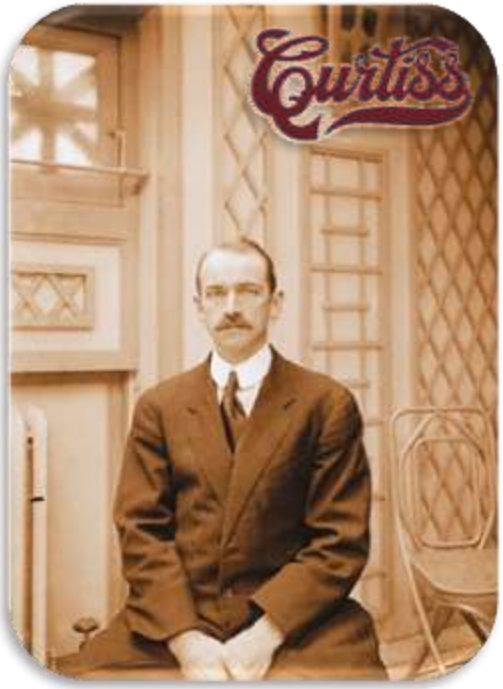

Aviator Glenn Curtiss
(1878 – 1930)

CW

- 100+ years of History
- $1.7+ B Company
- 7000+ Employees
- 800+ in Embedded Computing

VMETRO

- 21+ years of History
- $65+ M Company
- 200 Employees

Wilber Wright 1867 – 1912
Orville Wright 1871 – 1948

The Curtiss Jenny – By far one of the most successful mass produced planes in aviation history

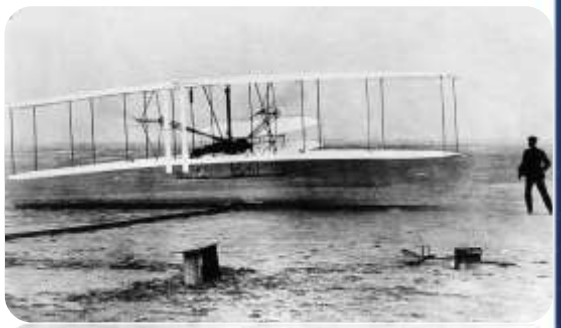

Kitty Hawk – 1903
The worlds first powered flight

A Rich History in Engineering Excellence and Innovation

VMETRO
a Curtiss-Wright Company

Curtiss-Wright VMETRO Overview

- FPGA processing
- Multi-processing
- Analog & digital I/O
- Buffer memory
- Network storage

- High Speed Data Recorders
- Ruggedized storage
- Bus analyzer – VME, PMC, CPCI, PCIE, AMC
- Form Factors – VPX, VXS, VME, PMC, XMC, FMC, CPCI, etc.

<u>Acquisition</u>, <u>processing</u>, and <u>recording</u> of *high volumes of data* in *real-time*

- Real-time Digital Signal Processing (DSP) applications
- Electronic Warfare (EW) and Intelligence, Surveillance, and Reconnaissance (ISR)

Signal & Data Processing Solutions

High-speed Data Recorders & Rugged Storage

VMETRO
a Curtiss-Wright Company

Why VMETRO-CW is supporting RapidIO

Embedded Digital Signal Processing Applications

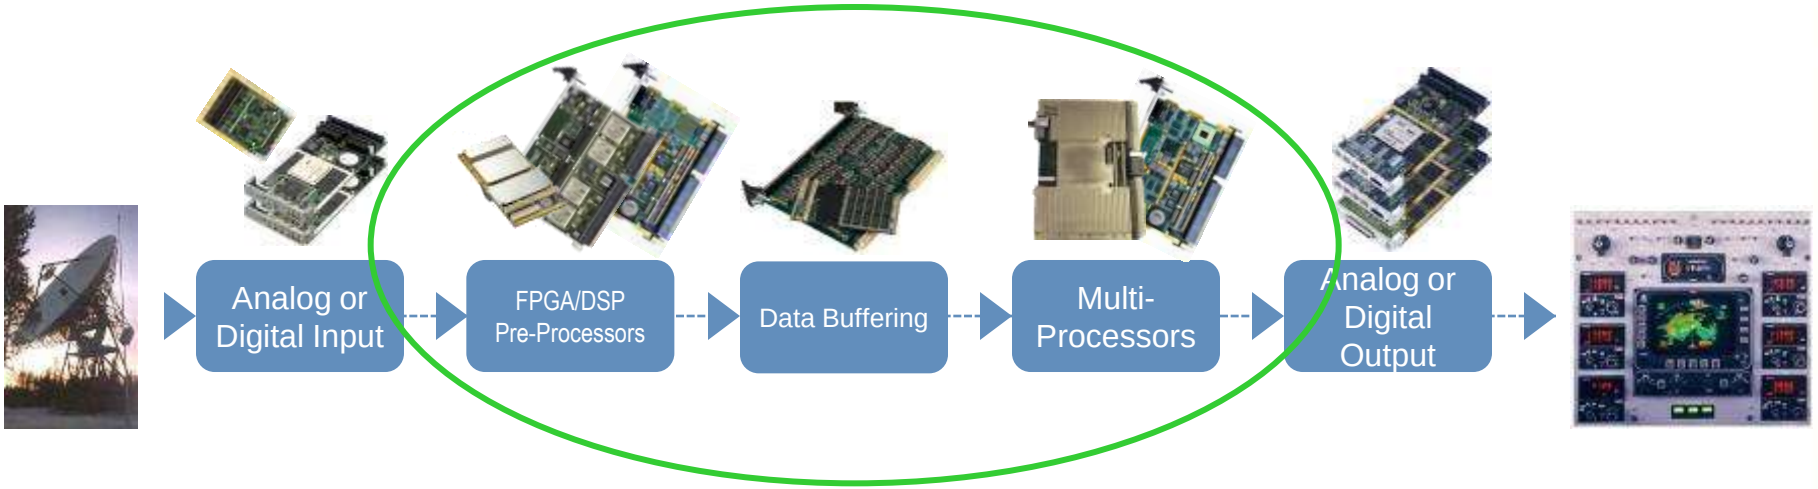

Distributed Multi-Processing (DMP)

VMETRO
a Curtiss-Wright Company

Distributed Multi-Processor (DMP) Applications

- High-end DSP applications such as military Intelligence, Surveillance and Reconnaissance (ISR) have very high performance requirements
  - very complex processing on massive amounts of data in real-time
    - No single processor exists that is capable to provide needed compute power
    - Need to **d**istribute application among **m**ultiple **p**rocessors
  - A high performance communication interconnect is necessary
- Ideal Distributed Multi-Processing interconnect characteristics
  - High throughput
  - Low latency and deterministic
  - Low overhead and low CPU utilization
  - Scalable
  - Reliable transport

VMETRO
a Curtiss-Wright Company

Interconnect Topologies and DMP Applications

Bus

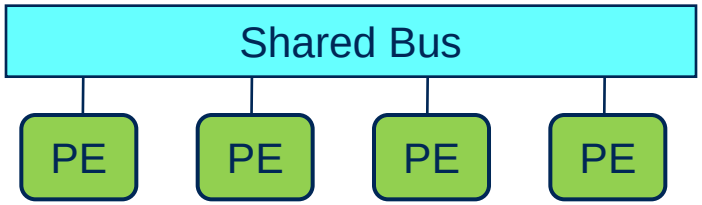

Switched - Tree

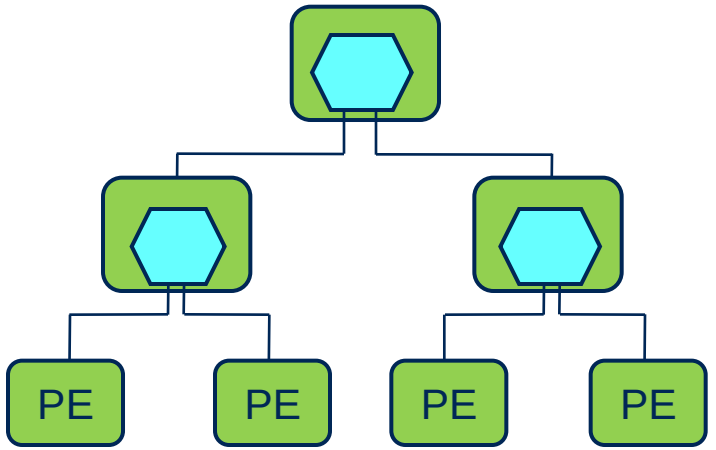

Switched - Star

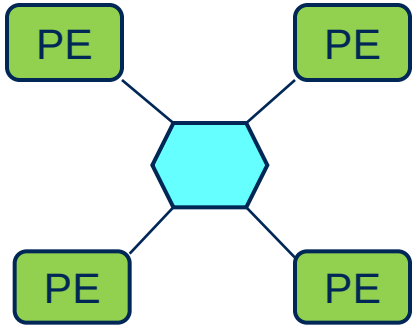

Switched - Network

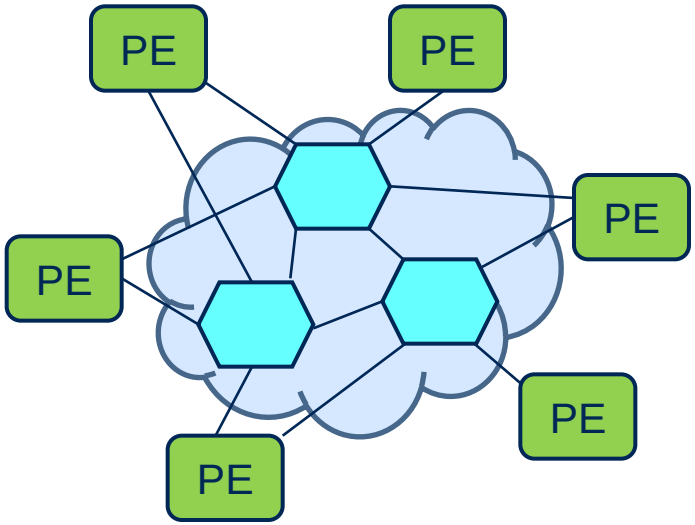

Switched – Full Mesh

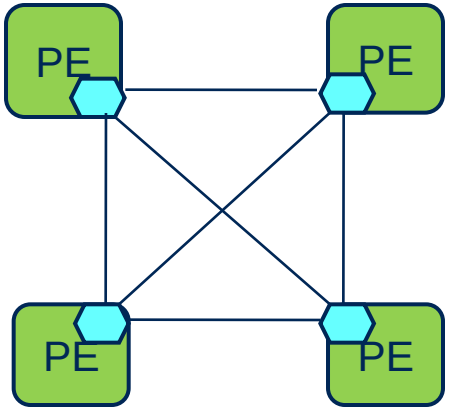

VMETRO
a Curtiss-Wright Company

VME continues to live

- VME – VME64x – 2eSST
- Parallel Bus – up to 64b, 320MB/s
- VITA 1.5 with various other extensions
- New VME extensions
  - VITA 42 (VXS)
  - VITA 46 (VPX)

VMETRO
a Curtiss-Wright Company

The evolving VME standards…

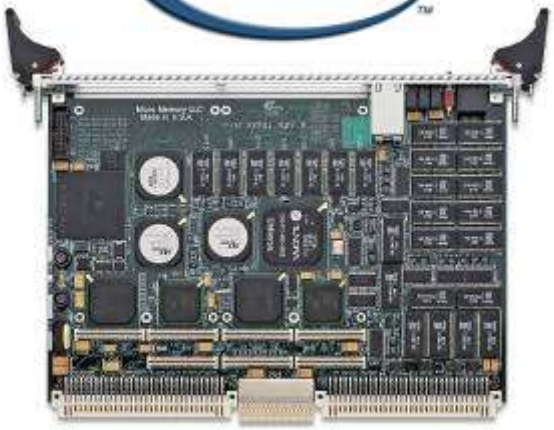

1982: VME16/32
1994: VME64
1998: VME64x
Parallel

<0.1GB/s

2004: VXS
VME with *some* High Speed Serial

2GB/s full duplex

2008: VPX
***All* High Speed Serial**

4-20GB/s full duplex

Speed improvement: 20-40x! **Speed improvement: 2-10x!**

VMETRO
a Curtiss-Wright Company

VPX – where is VME?

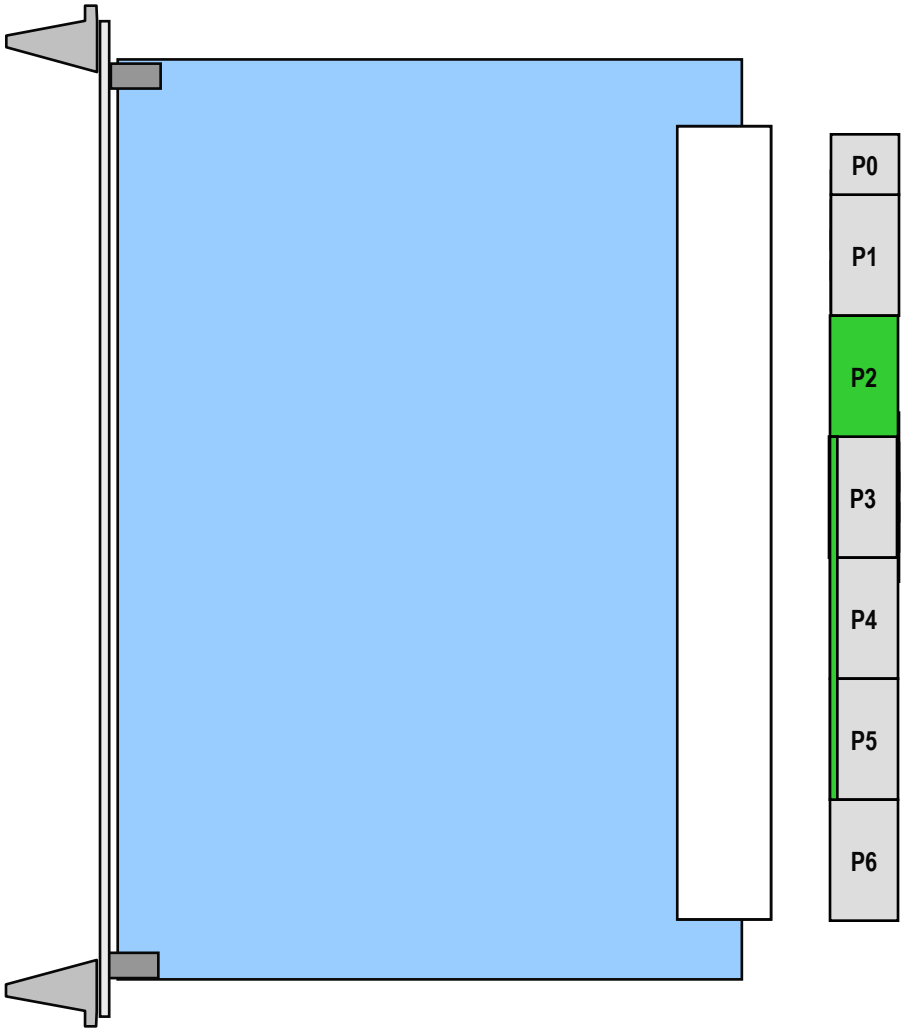

If there is VME (V46.1), it goes here – optional

A24/D16 in P2

A32/D32 in P3, P4 & P5

VMETRO
a Curtiss-Wright Company

Ideal Interconnect for DMP Applications

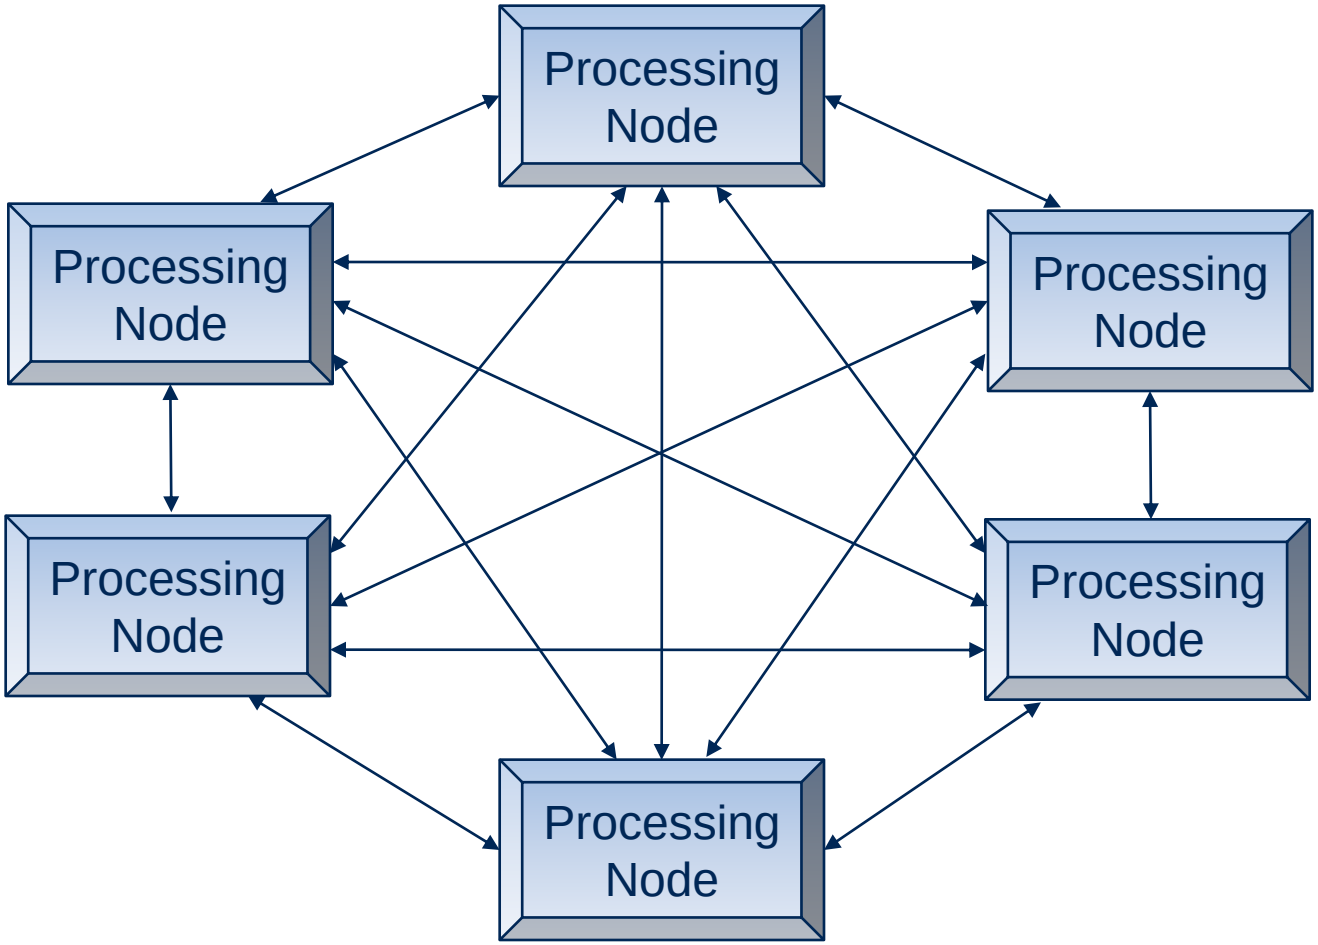

- Full mesh topology – every processing element has a direct connection to every other processing element as peers (peer-to-peer)
- Requires a fabric protocol that supports distributed switching architecture
- Inter-Process Communication (IPC) software makes use of fabric

VMETRO a Curtiss-Wright Company

Serial Protocols for DMP Interconnects

Possible Candidates:

Gb Ethernet

Let’s take a closer look at these ...

VMETRO
a Curtiss-Wright Company

Aurora

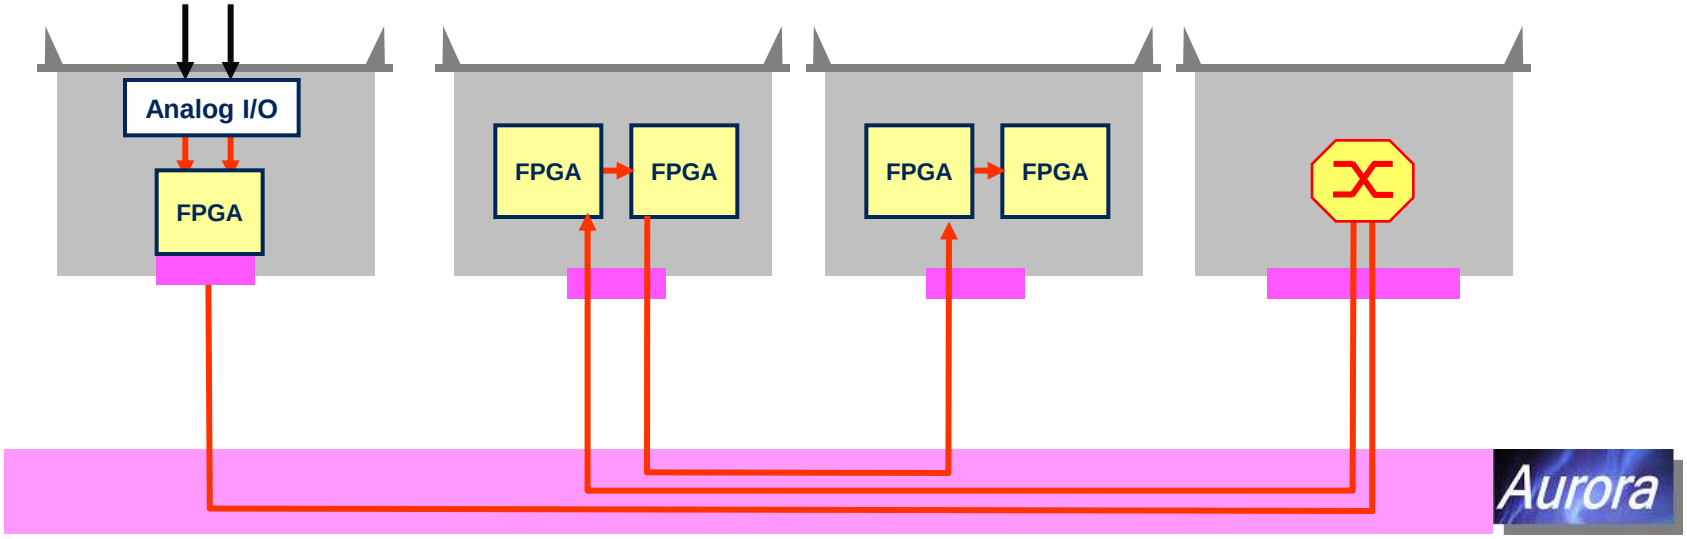

Point-to-point domain

- Point-to-point, high-speed, low latency, low overhead protocol
- Supported in Xilinx FPGAs
- Well suited for interfacing to I/O and for FPGA to FPGA communication
- Does not support the features that DMP software needs

VMETRO
a Curtiss-Wright Company

PCI Express

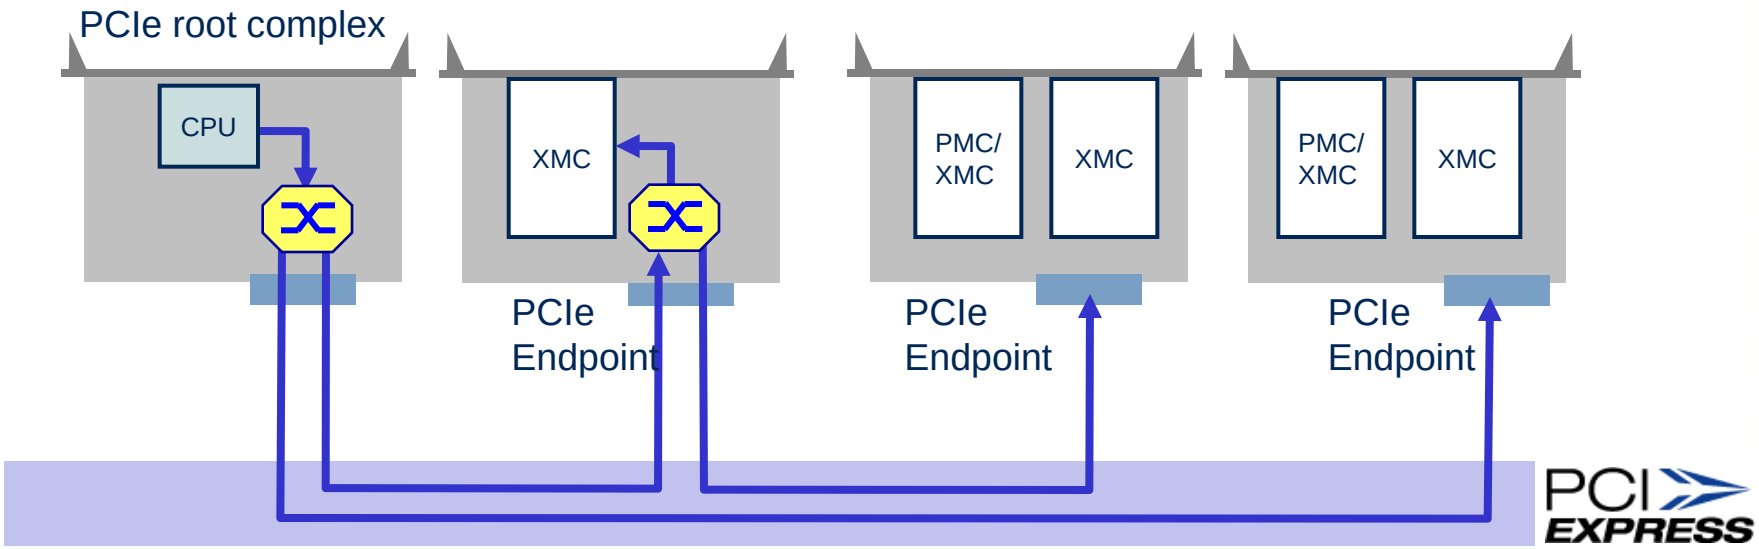

- Natural for I/O because of PCI/PCIe support in I/O devices
  - I/O expansion, e.g., PMC/XMC carriers
- Tree topology – PCIe root complex, just like PCI
  - Does not support peer-to-peer communication
- Simple SW model – memory mapped as PCI

VMETRO
a Curtiss-Wright Company

PCI Express for DMP

- Requires PCIe root complex and use of non-transparent bridges
  - Distributed switch topology not possible
- Requires careful address window mapping
  - Manageable for small systems

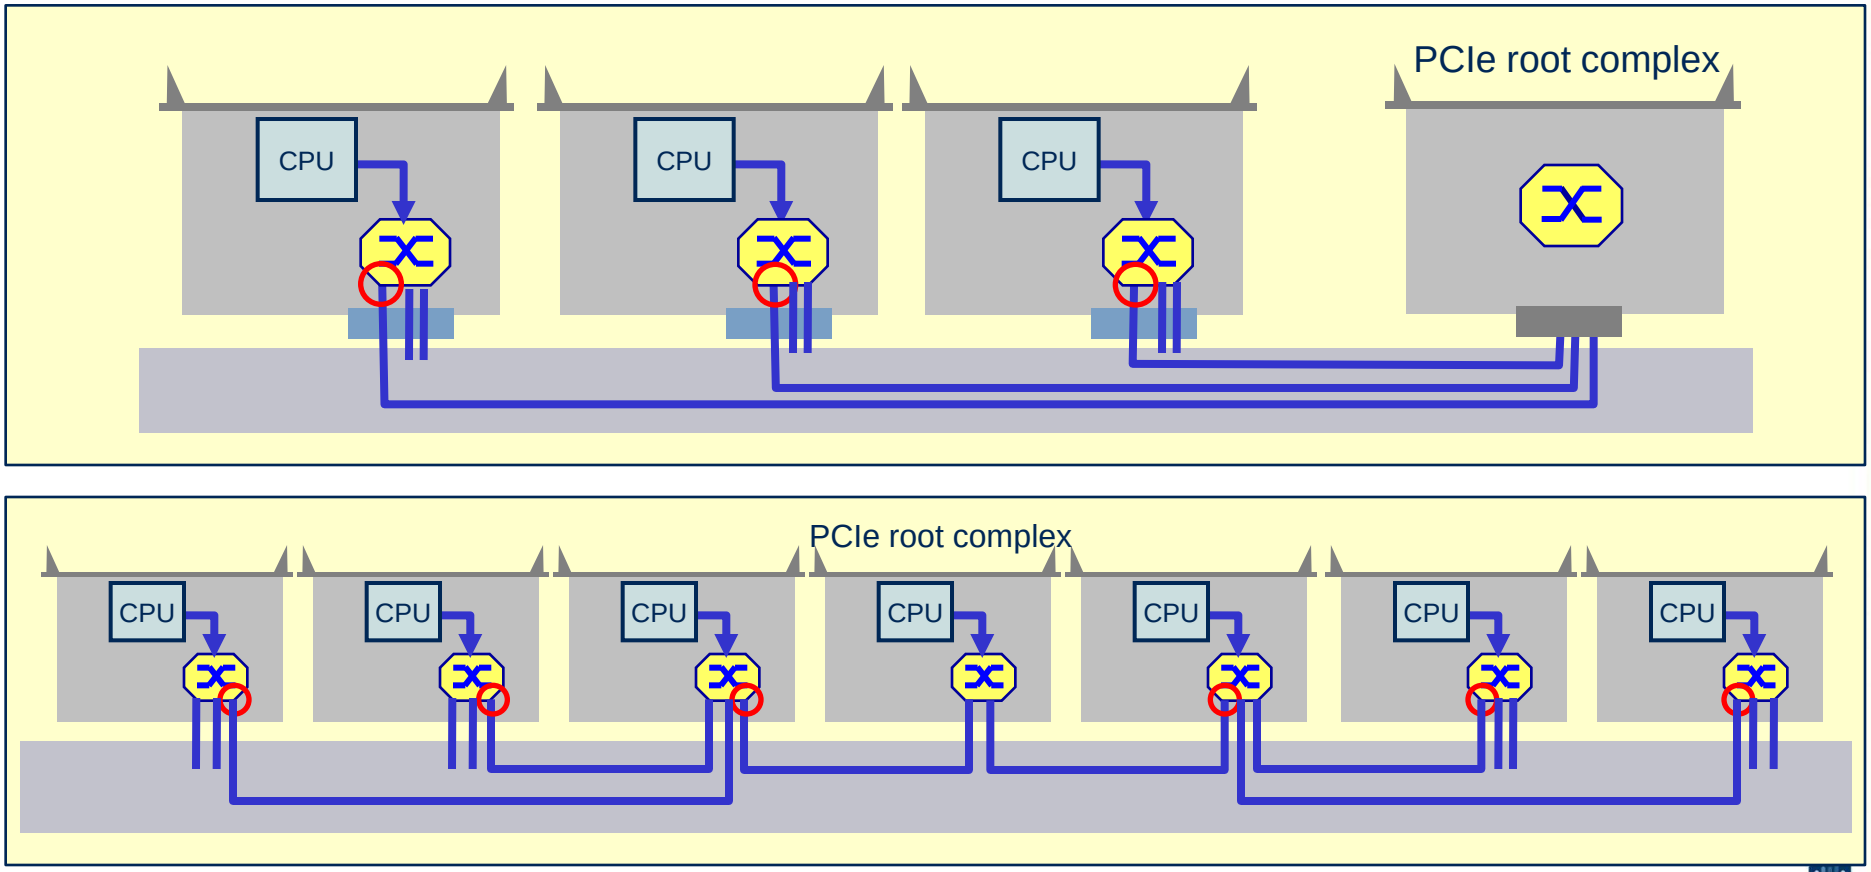

VMETRO
a Curtiss-Wright Company

Gb Ethernet

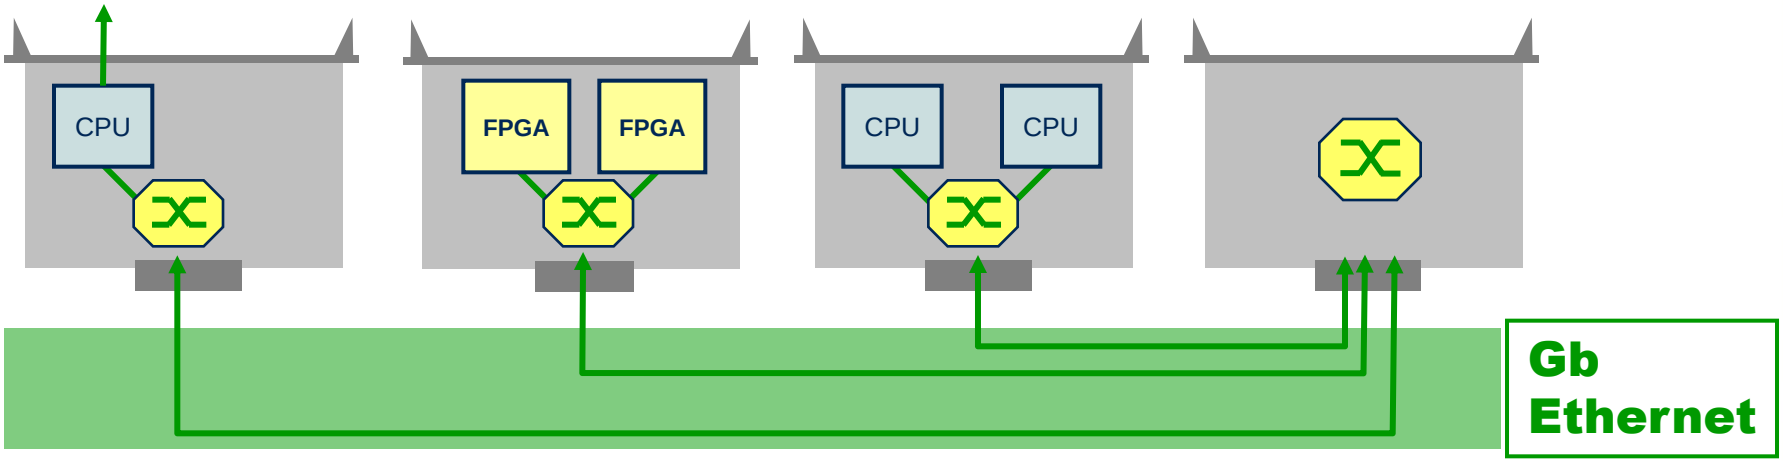

- Default standard for networking
  - Connects the embedded system to the outside world
- TCP/IP software is well suited to system control needs
- Ethernet does not provide a reliable transport without TCP/IP
  - TCP/IP requires a tremendous amount of CPU processing bandwidth
- Multi-processing can be done with Ethernet
  - Performance is the issue – determinism, latency, and CPU utilization
    - Low-end or non-real-time applications: Ethernet is sufficient
    - High-performance, real-time DMP applications: Ethernet does not meet the performance requirements

VMETRO
a Curtiss-Wright Company

Serial RapidIO

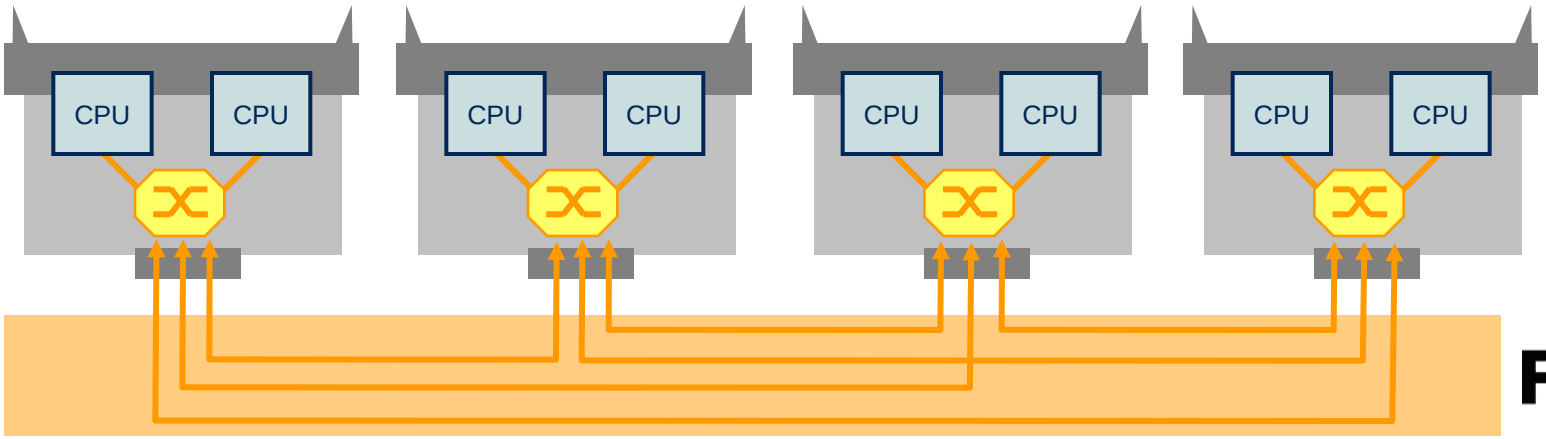

RapidIO

- True peer-to-peer, switched fabric protocol
- Topology flexibility
  - Point-to-point
  - Switched
  - Distributed switched, aka "full mesh"

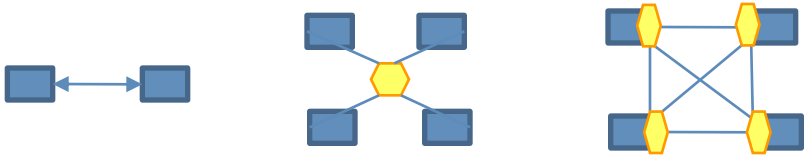

- Good connectivity without need for dedicated switch card
- Reliable transport provided in hardware rather than software
- Ideal for Distributed Multi-Processor (DMP), facilitated with Inter-Process Communication (IPC) S/W

VMETRO
a Curtiss-Wright Company

DMP Full-Mesh sRIO Architecture

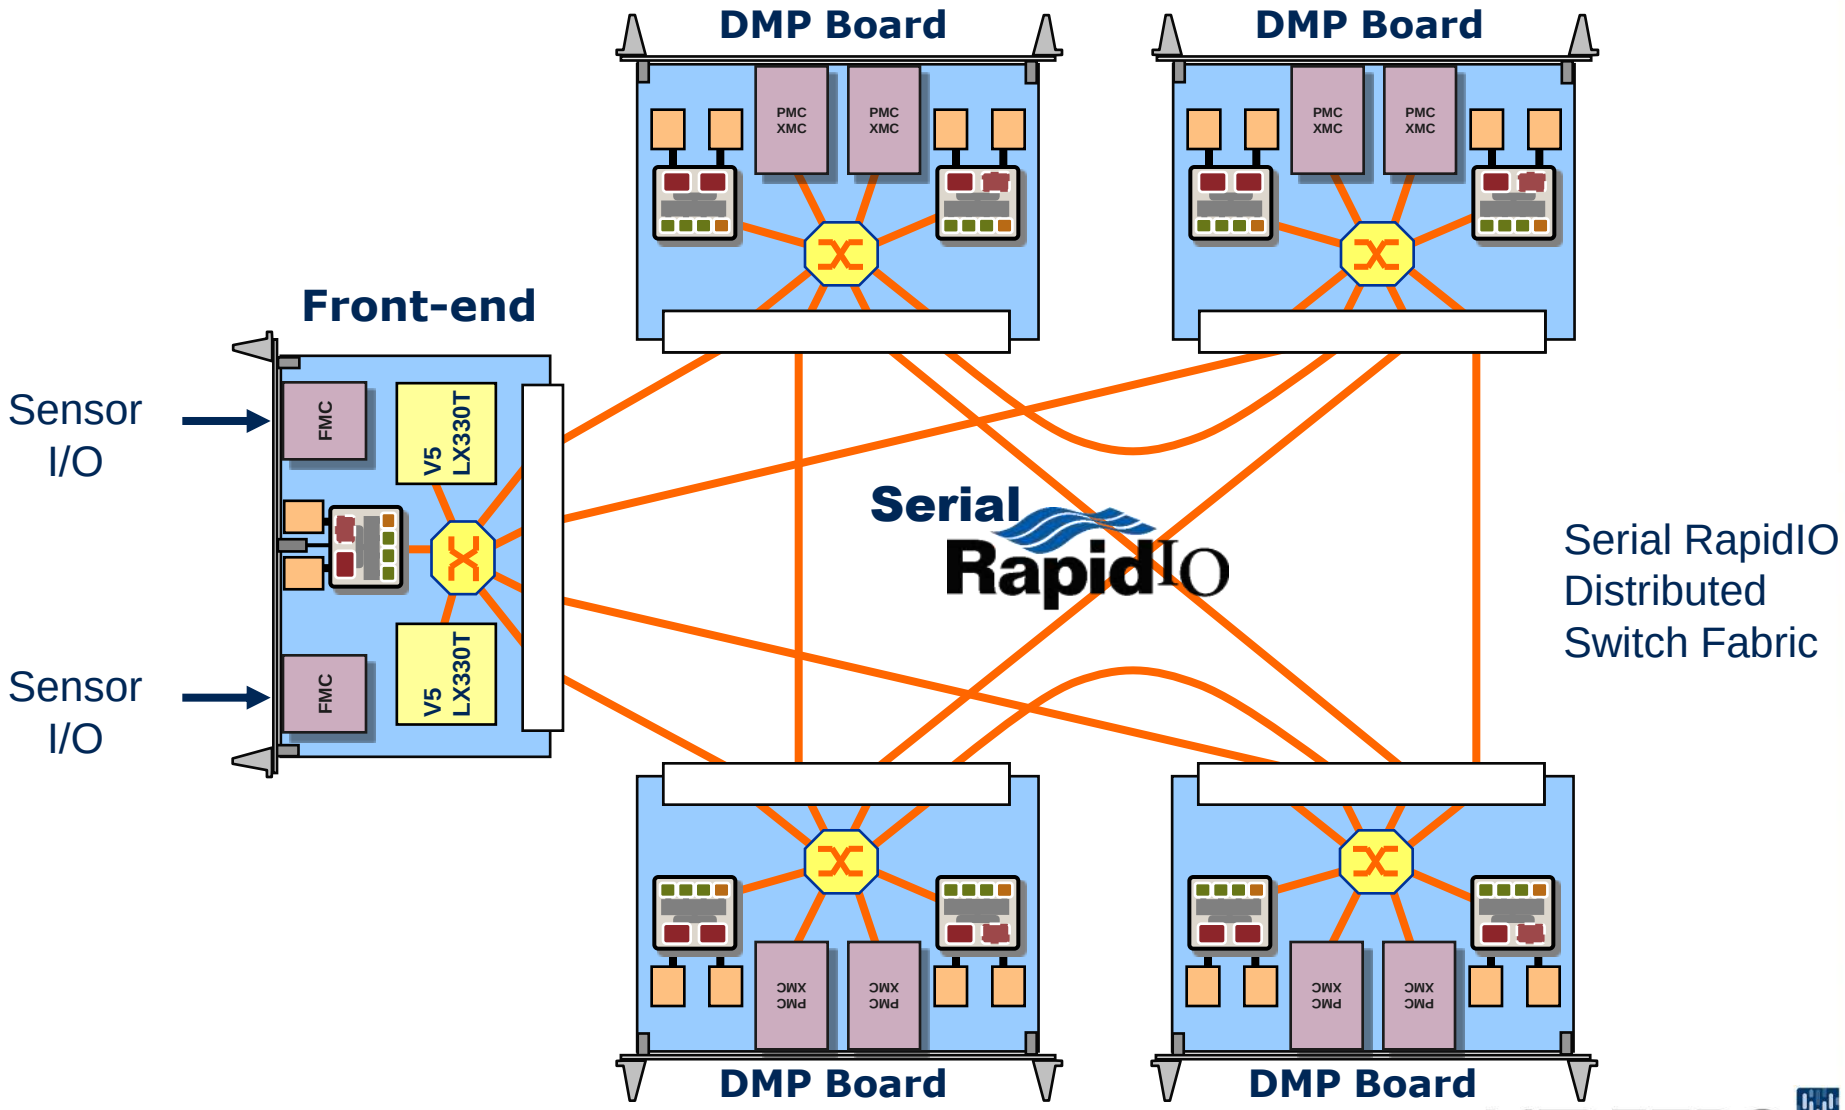

The Protocols Working Together in DMP Systems

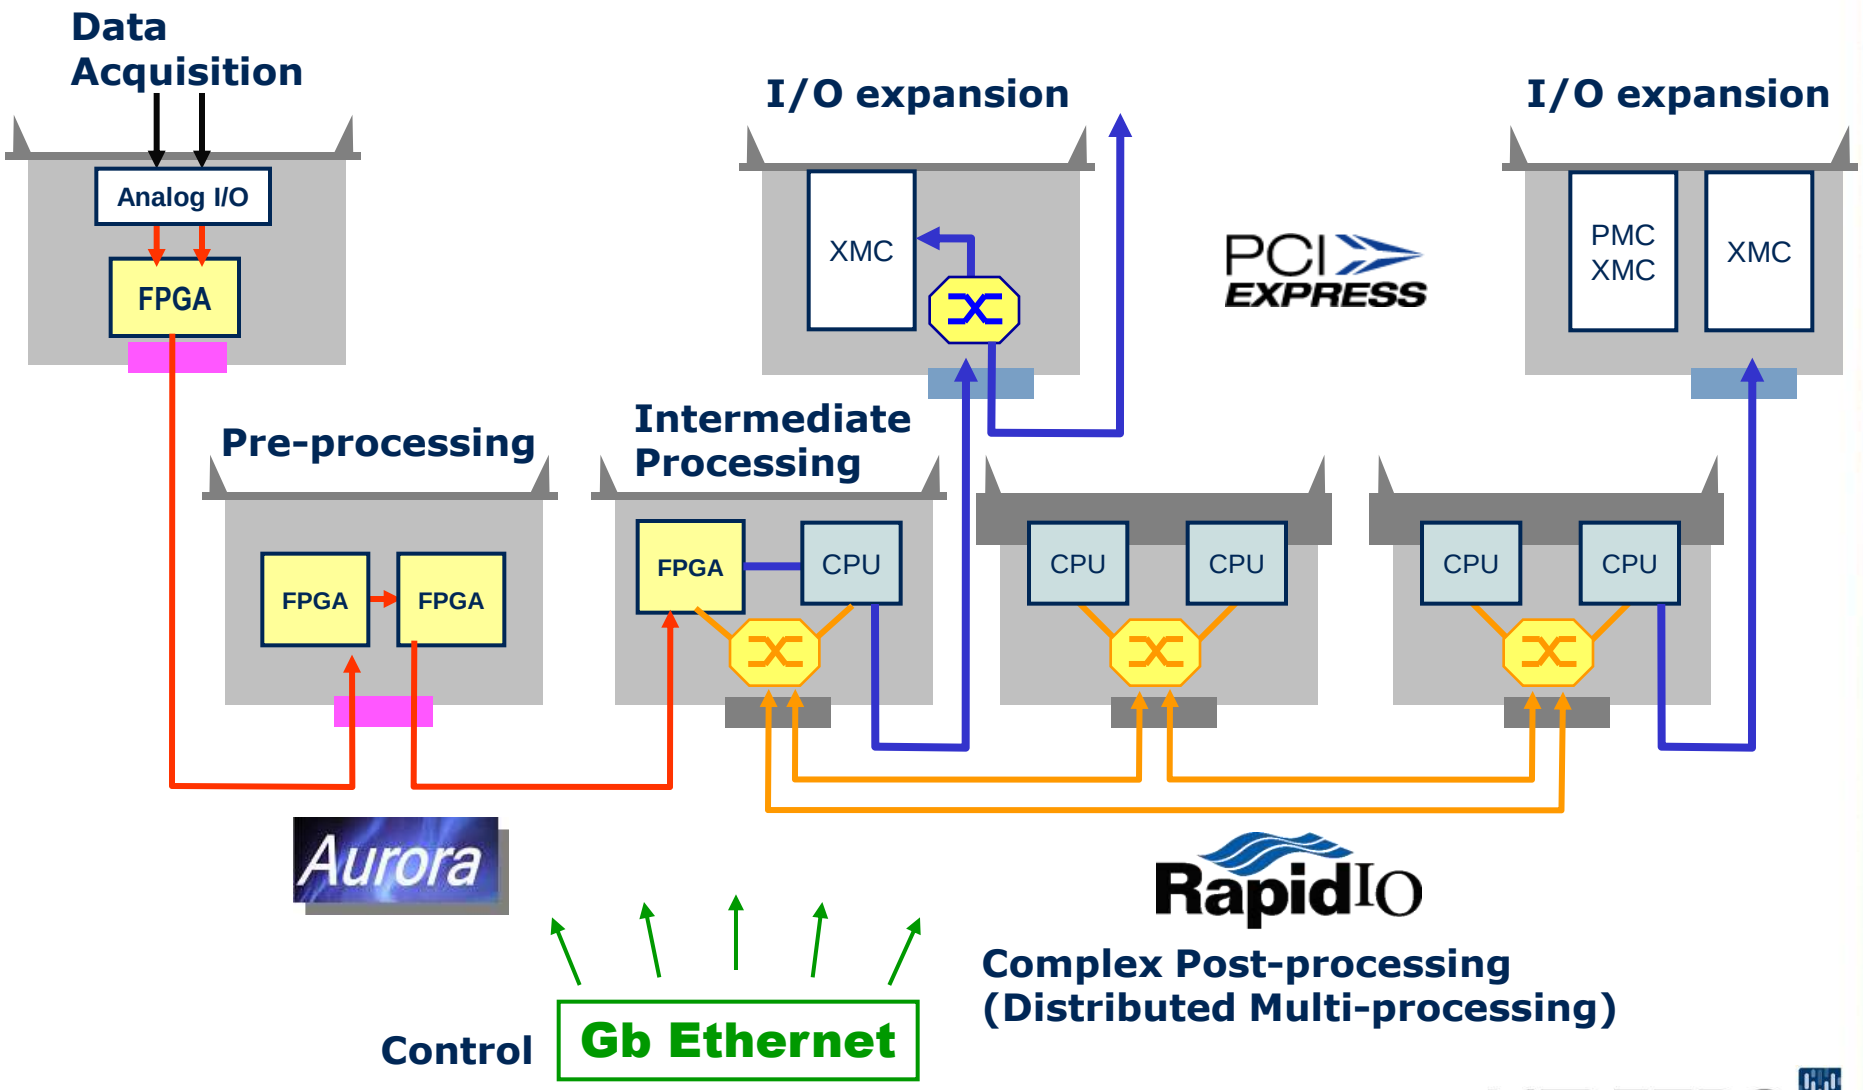

VMETRO
a Curtiss-Wright Company

RapidIO Features and DMP

- Link speeds of 1.25, 2.5, or 3.125 Gb, going to 5.0 and 6.25 Gb in 2.0
- Links can be bonded, x1 or x4 lanes, support for x2, x8, & x16 in 2.0
- Minimum transaction overhead
  - Due to its low protocol overhead and hardware-based implementation, latency is extremely low, in the range of 200 nanoseconds, and real data rates are much closer to theoretical maximums than other interconnects
- Hardware-based error handling
  - Protocol includes error recovery mechanisms for packet retry, stomp, link request/response, and CRC
- Quality of Service with multiple hierarchal flow control mechanisms
- Devices may have up to 256 outstanding transactions between two end points
- Hardware-based Message Passing and Protocol Processing
  - Message passing is hardware-based through the use of mailboxes and doorbells (similar to interrupts) which are easily utilized by software

RapidIO Ecosystem

Processors

Crossbar Switches

RapidIO IP Blocks

VMETRO
a Curtiss-Wright Company

RapidIO: Fabric of Choice for DMP

Why choose Serial RapidIO for embedded DMP applications

	Interconnects				
	sRIO	AS*	PCIe	GbE	IB
Low Latency	√	√	√	O	√
Efficient Power Consumption	√	√	√	O	O
Scalable Memory Mapping	√	√	O	√	√
True Switch Fabric Capabilities	√	√	O	√	√
Native in currently available GPPs, Switches & ASSP's	√	O	√	√	O
Simplicity	√	O	√	√	O

*** AS is not commercially available; included for comparison only**

VMETRO a Curtiss-Wright Company

IPC Software Architectures for DMP

- Message Passing
  - Communication is made by the sending of messages to recipients; application has no knowledge of hardware, examples are MPI and CORBA
  - Known as "shared nothing" approach because the message passing abstraction hides the underlying implementations for sending/receiving messages
  - Can work in a heterogeneous system (i.e., boards from multiple vendors)
  - Same application code can be used “in the box” or across networks
  - Does not exploit inherent hardware capabilities
- Shared Memory Buffer (SMB)
  - Buffers on local nodes are “shared” between all nodes – they are not under the control of the local node. Data is moved via DMA transfers. Used for bulk data transfers
  - Efficient – data does not have to be copied from user space to system space
  - Makes use of inherent hardware capabilities
  - More tied to the hardware, not easy to use in heterogeneous system, for use “in the box”
- RapidIO supports both IPC software architectures
  - SMB approach utilizes “messages” for synchronization and control
- SMB approach provides highest performance for DMP applications

VMETRO
a Curtiss-Wright Company

IPC Example using Shared Memory Buffers

Scenario

- Two processing Elements: PE-A & PE-B
- Processes running on the PEs that need to share data
- Each PE has a RapidIO endpoint and switch

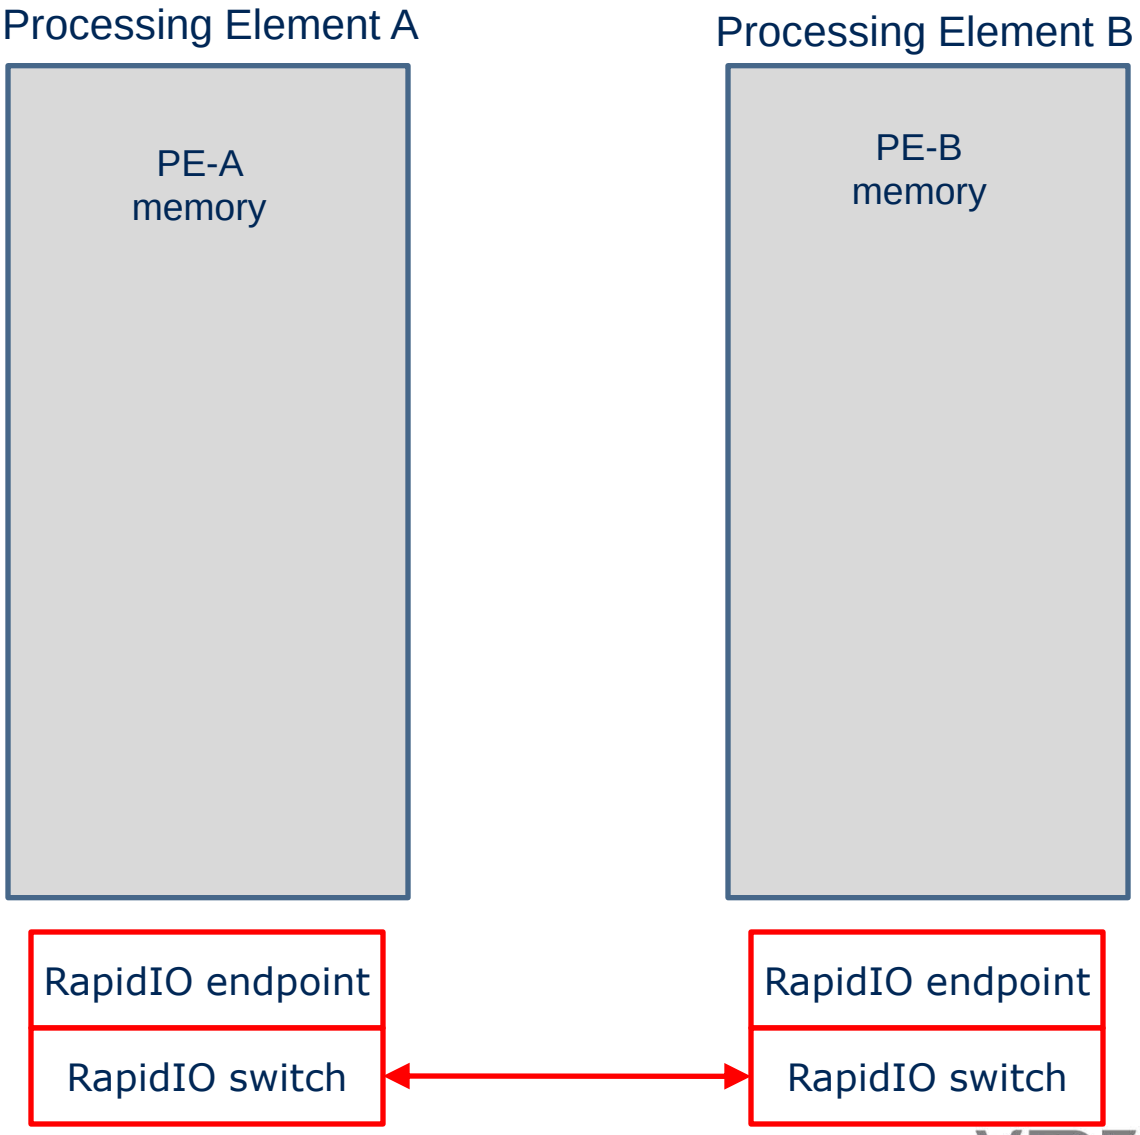

VMETRO a Curtiss-Wright Company

IPC Example using Shared Memory Buffers

System Initialization

- At system start-up, each PE's memory is split into local memory and SMB memory

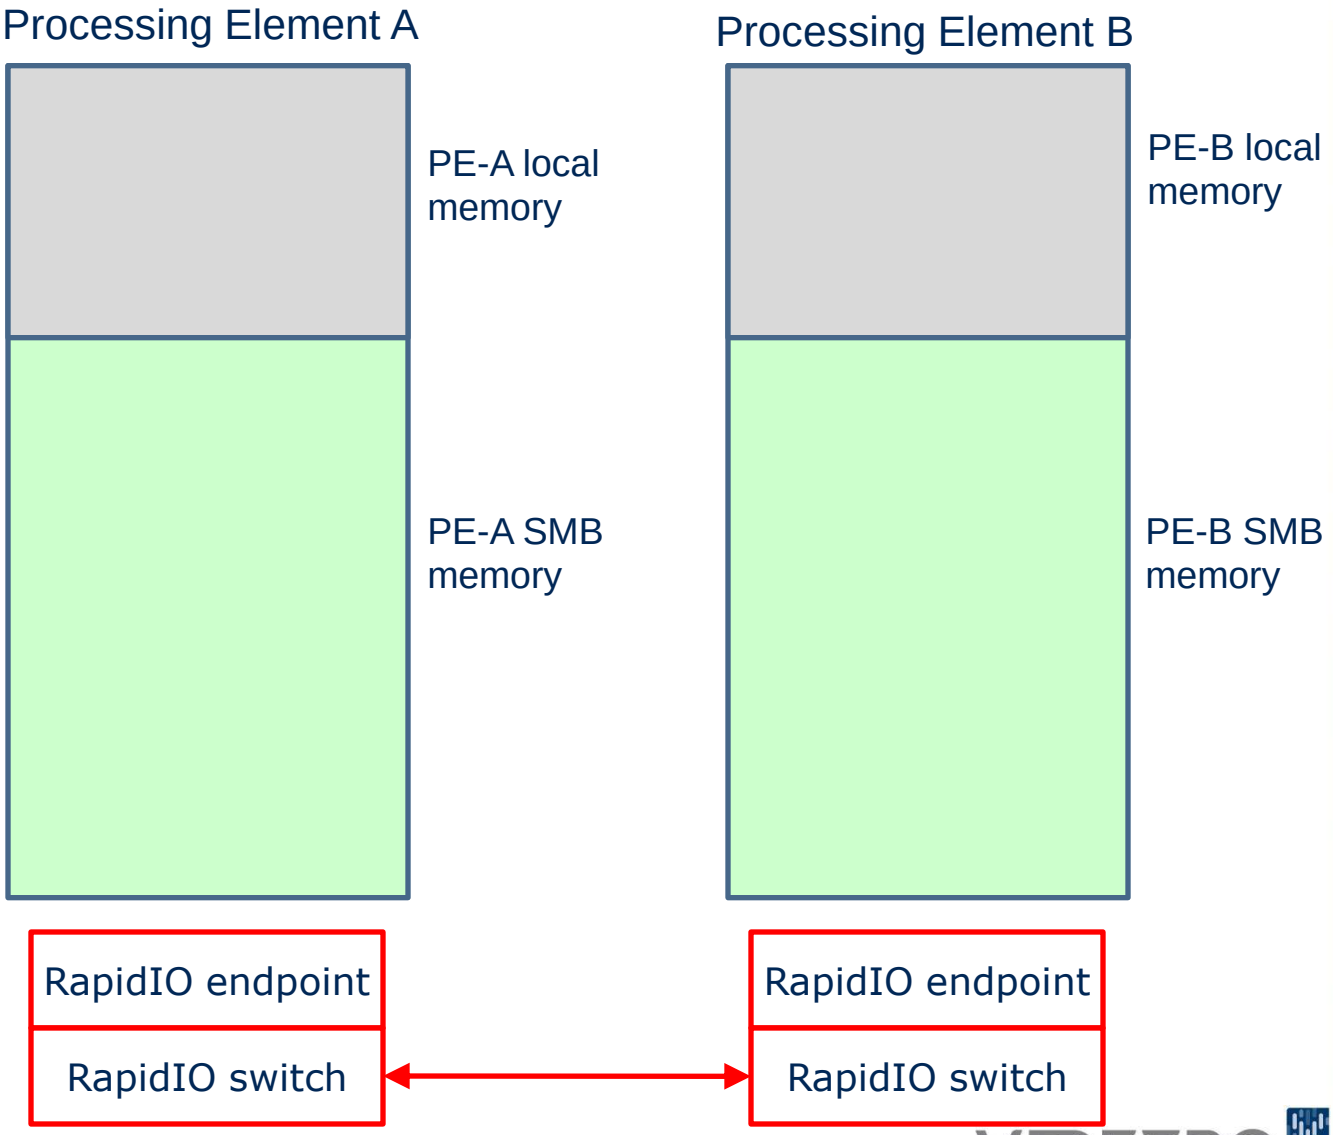

IPC Example using Shared Memory Buffers

System Initialization

- SMB are allocated by O/S
- SMB memory is a "system resource"; from a user process, SMBs can only be accessed via IPC APIs
- SMBs are "shared" between all processing elements; all PEs in the fabric can access them

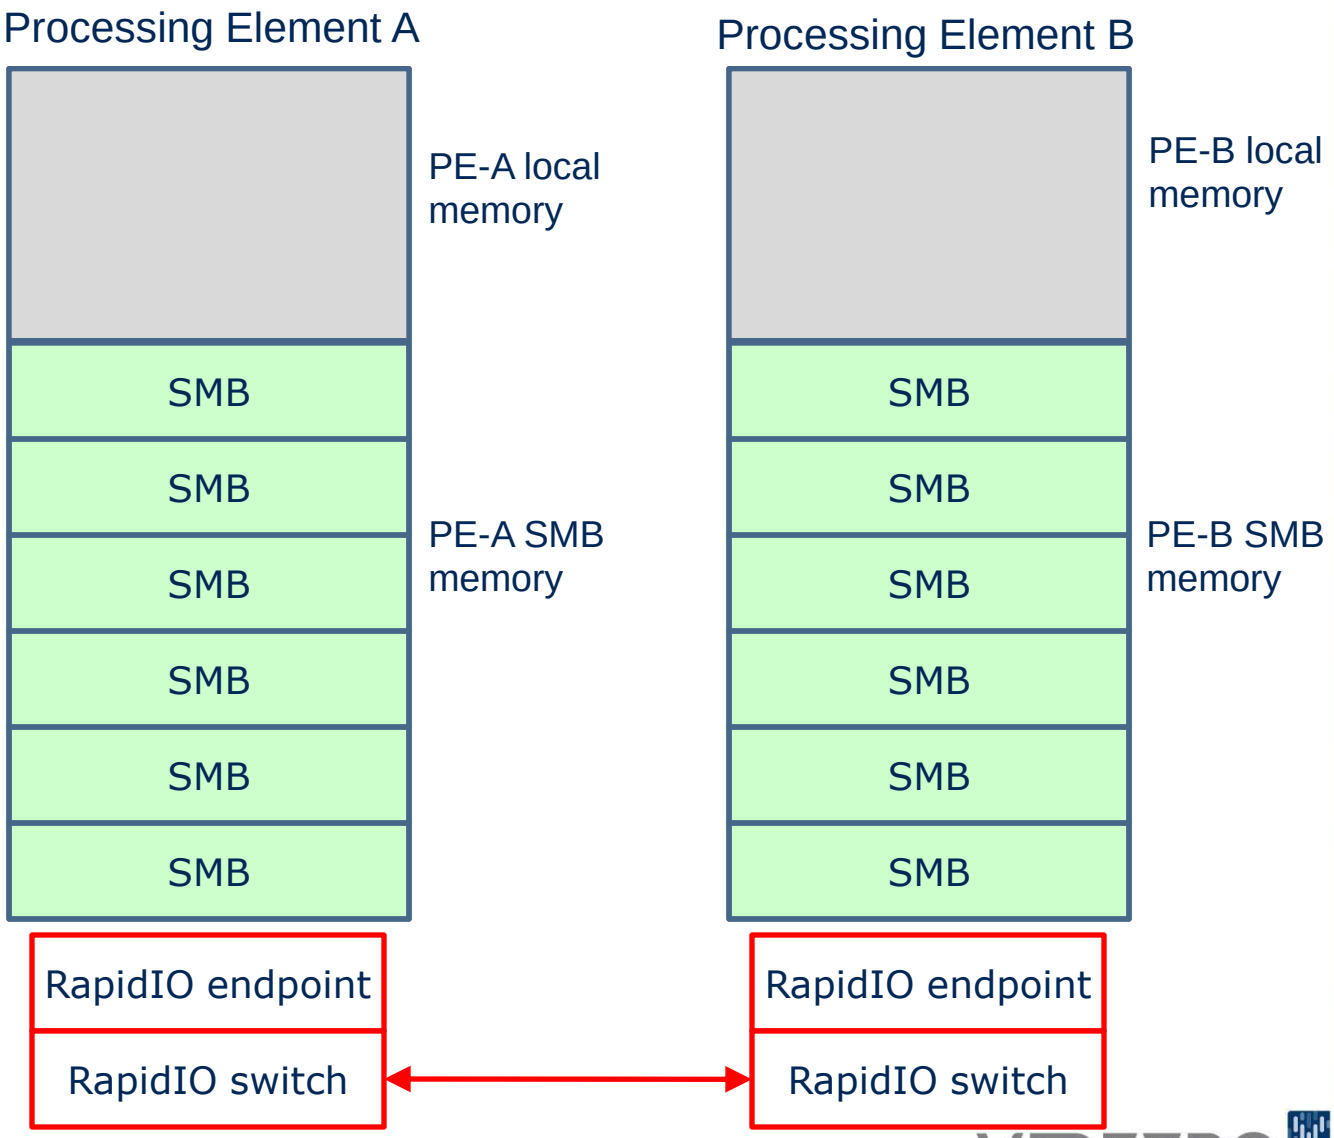

VMETRO a Curtiss-Wright Company

IPC Example using Shared Memory Buffers

Application Start-up

- The processes running on two distributed processors are started

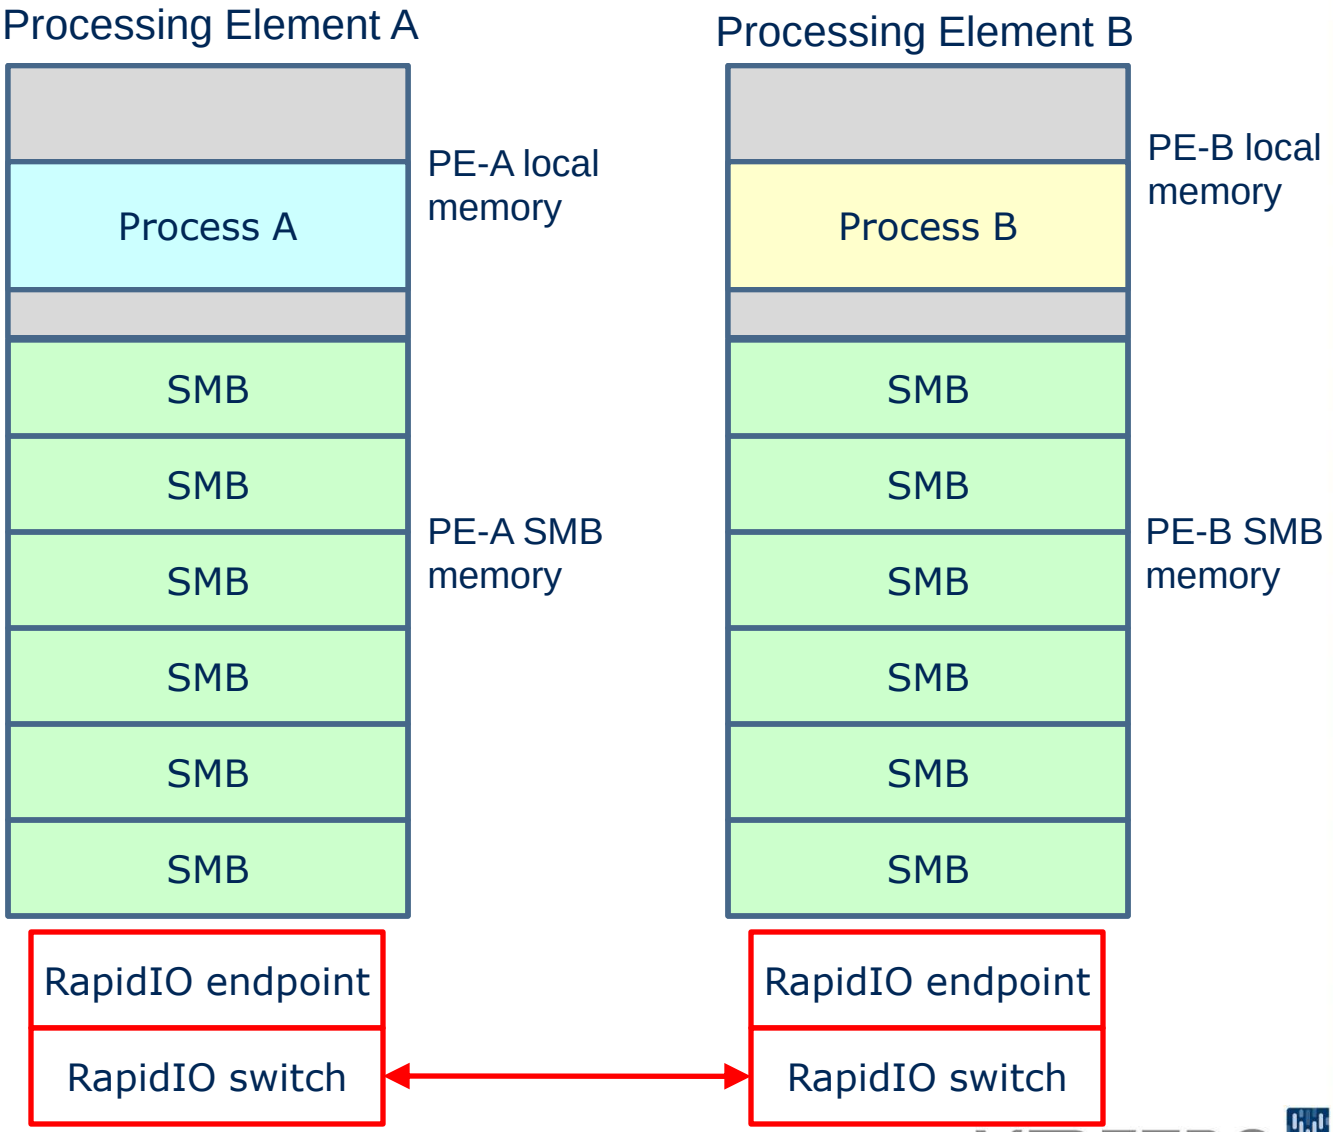

IPC Example using Shared Memory Buffers

Application Initialization

- Processes map SMBs into their address space via an IPC API
- Processes can now directly read from and write into SMBs without having to copy data between user space and system space

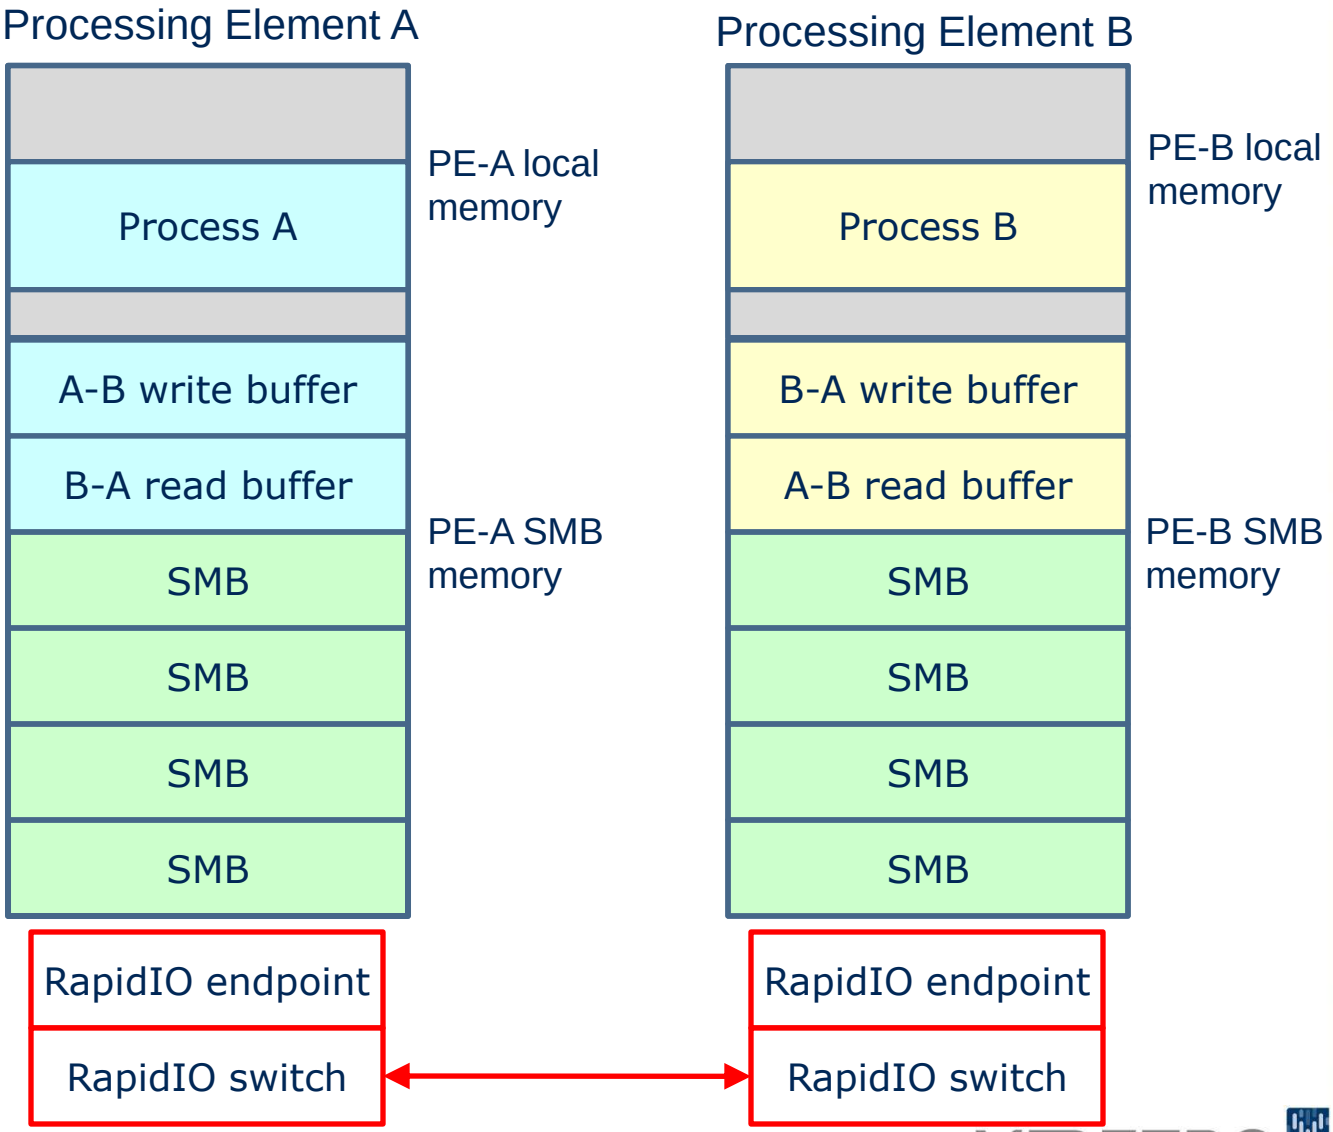

IPC Example using Shared Memory Buffers

Processing Loop

- Process A writes into the ***A-B write buffer***
- Process A initiates a RapidIO DMA transfer to ***A-B read buffer*** via an IPC API
- Using RapidIO message passing and doorbells, Process B will be notified when the data has been transferred into ***A-B read buffer***
- Process B reads from ***A-B read buffer***

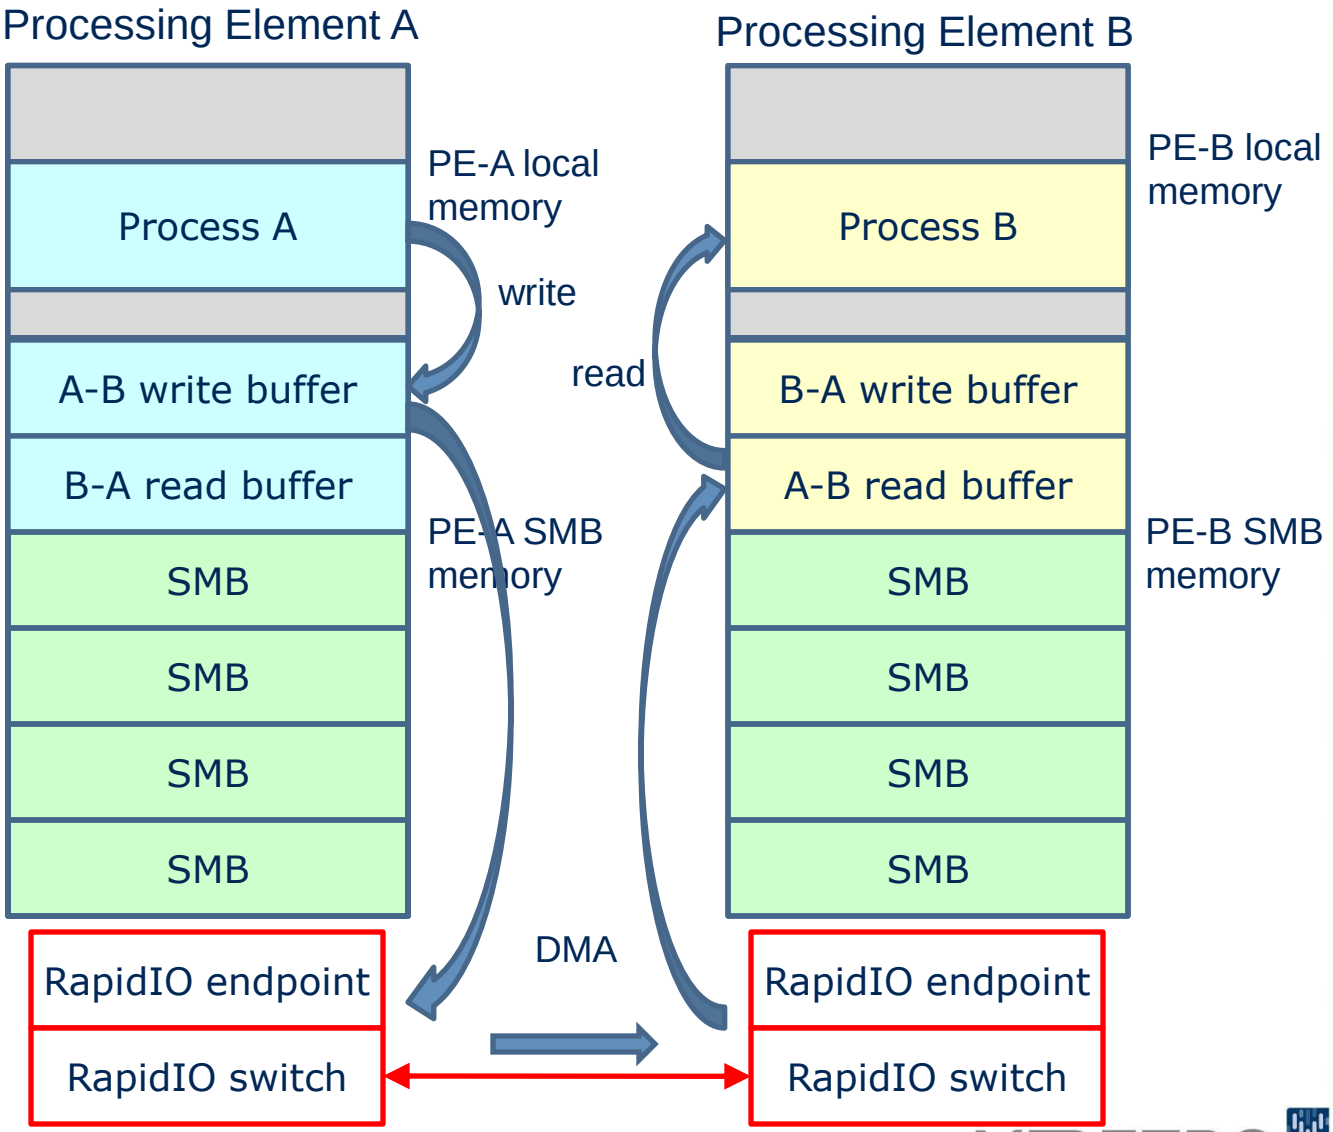

VMETRO
a Curtiss-Wright Company

Summary: *RapidIO is the Right Choice for DMP*

- Distributed Multi-Processing requires a “distributed” fabric
  - RapidIO is a switched, peer-to-peer protocol
- Hardware and protocol features to easily support IPC software model necessary for Distributed Multi-Processing
  - High throughput, low latency, deterministic, low overhead and low CPU utilization, scalable
  - Flexible topologies – supports mesh topology via distributed switching
  - RDMA support
  - Reliable transport - guaranteed delivery
  - Supports message passing and SMB IPC software architectures
    - RapidIO has built-in hardware message queues for message passing
    - Bulk data movement is handled via DMA transfers
- Robust ecosystem of processors, switches, & IP necessary for embedded market

VMETRO
a Curtiss-Wright Company

VMETRO
&
Serial RapidIO

VMETRO
a Curtiss-Wright Company

VMETRO's RapidIO IP

- VMETRO's RapidIO product development started in 2003 with a Parallel RapidIO XMC (RMC) that was successfully delivered to OEM customer in 2004
- VMETRO RapidIO products are based on Xilinx FPGAs and Freescale 8641D PowerPC processors
- VMETRO Serial RapidIO FPGA IP is based around Xilinx LogiCore RapidIO modules
  - VMETRO added Control Status Module and sophisticated Serial RapidIO corner turning, multi-ported DMA engine

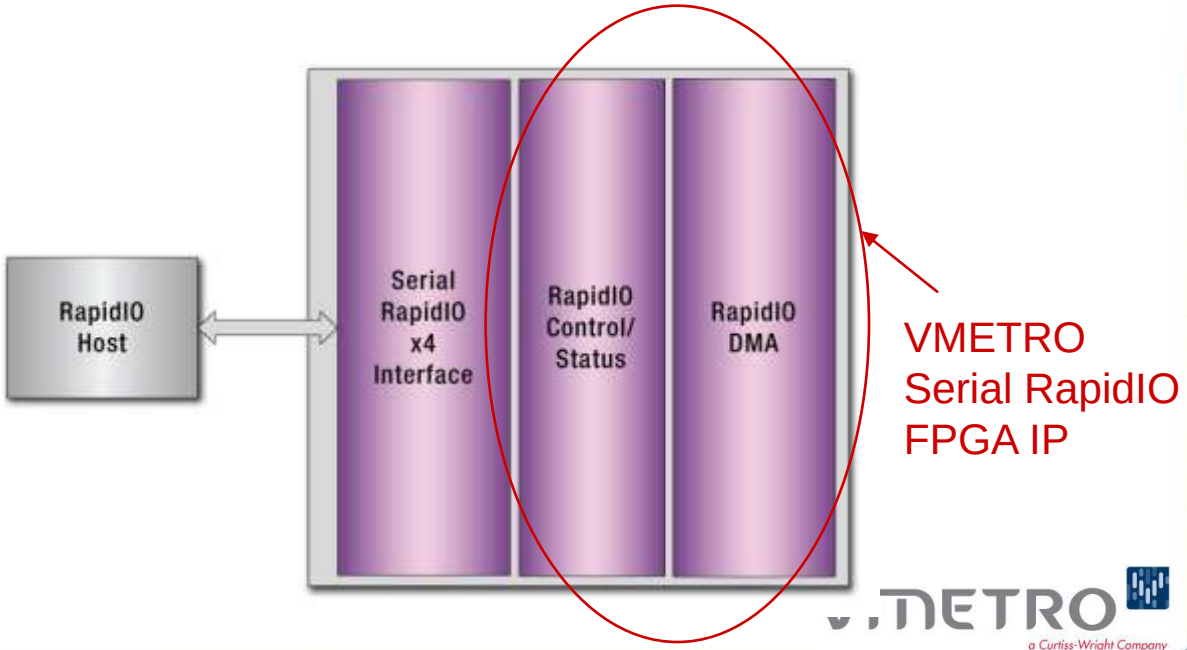

VMETRO-CW's FusionIPC / Continuum
Inter-Processor Communication Software Suite

- Targets Real Time, Distributed Multi-Processing (DMP) Applications
- Designed for low overhead and close control of hardware resources
- Based on scalable fabric endpoints, Shared Memory Buffers (SMBs), POSIX or Linux standards, and open platforms
- Utilizes RDMA, Zero Copy implementation – avoids excessive data copying
- Dynamic and Abstract Configuration Management as opposed to Static and Geographical
- Well defined User API provides Distributed DMA and Distributed Name Service

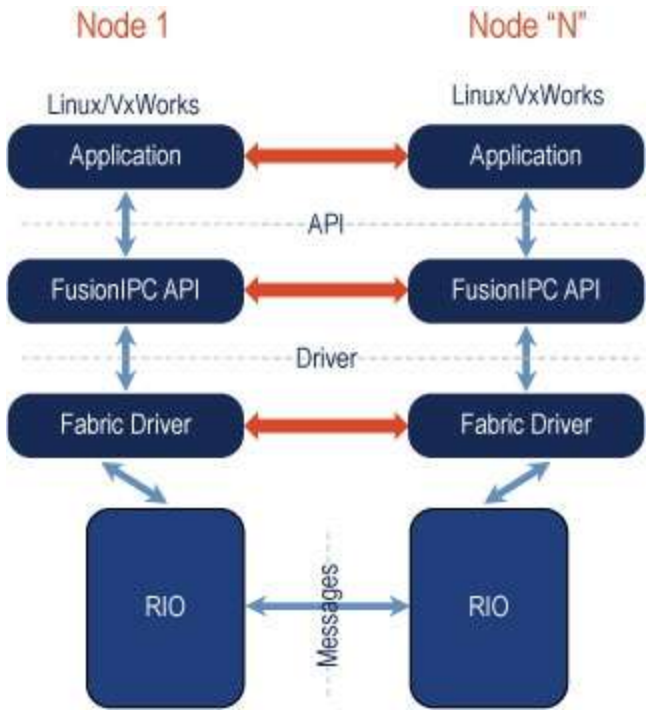

VMETRO
a Curtiss-Wright Company

Interoperability

Extensive Validation Testing of VMETRO's Serial RapidIO FPGA IP with other industry offerings

Full memory to memory transfers and comprehensive system diagnostics

- Components/IP
  - Freescale 8548 PowerPC
  - Freescale 8641D
  - Tundra Serial RapidIO Crossbar Switch
  - Mercury Computers' pRIO-to-sRIO bridge and Crossbar Switch
- Boards
  - Extreme Engineering XES-5203 XMC
  - Mercury Computers' Ensemble ATCA carrier with GDA Freescale RMCs
  - Curtiss-Wright VPX-185
  - Curtiss-Wright AV6
- Chassis
  - Hybricon VXS system (two sRIO x4 on VXS P0)
  - Hybricon VPX system (five slot full sRIO x4 mesh on P1)
  - Mercury Ensemble ATCA chassis
- Software IP
  - VMETRO FusionIPC sRIO Libraries
  - Curtiss-Wright sRIO ContinuimIPC
  - Mercury Computers' sRIO API framework

VMETRO
a Curtiss-Wright Company

VPX Products
with Serial RapidIO

NEW!

RapidIO VPX REDI

MPE730
Quad-core PPC Multi Processing Engine

- Two Dual core 8641D processors up to 1.5GHz
- sRIO and x8 PCIe fabric switches
- Two PMC / x8 PCIe XMC.3 sites
- Slot-1 capable, incl. VME I/F
- Rugged Air & conduction-cooled versions

SBC731
Single Board Computer

- One Dual core 8641D processor up to 1.5GHz
- sRIO and x8 PCIe fabric switches
- Two PMC / x8 PCIe XMC.3 sites
- Slot-1 capable, incl. VME I/F
- Rugged Air & conduction-cooled versions

VMETRO
a Curtiss-Wright Company

VPX Products
with Serial RapidIO

NEW!

RapidIO

VPX REDI

VIRTEX 5

HPE720
PowerPC & FPGA Hybrid Processing Engine

- Dual-core Freescale MPC8641D processor up to 1.33GHz
- Two Xilinx Virtex-5 FF1738 FPGAs
- LX110/220/330T, SX240T or FX200T build options
- Two FMC sites, alt. one FMC and one XMC/PMC
- sRIO and x8 PCIe fabric switches
- Rugged Air & conduction-cooled versions

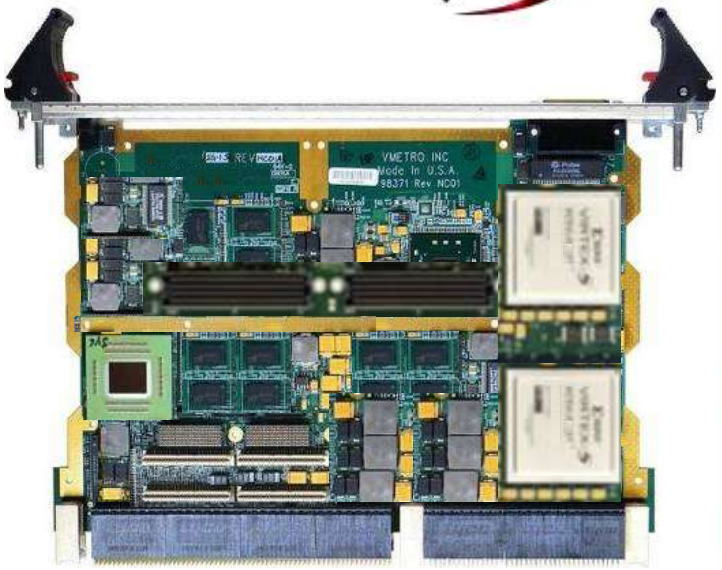

MM-6171
Buffer Memory XMC Module

- 4GB of SDRAM Memory
- x4 sRIO fabric port
- Rugged Air & conduction-cooled versions

VMETRO
a Curtiss-Wright Company

VPX Products
with Serial RapidIO

NEW!

RapidIO

VPX REDI

VIRTEX 5

CHAMP-AV6 VPX Quad 8641D

- Quad 8641/8641D processors @ 1GHz
- 64GFLOPS peak performance
- 512 Mbytes – 1Gbyte DDR2 SDRAM per processor
- 256MB Flash
- On-board SRIO fabric with 5GB/s bandwidth
- Four off-board Serial RapidIO ports, up to 10Gbytes/s
- One PCI/PCIe XMC site
- Gigabit Ethernet per processor, on-board switch
- Four EIA-232, two EIA-422 ports

VPX6-185 SBC Single 8641(D)

- single 8641 device
- up to 2GB DDR II SDRAM, up to 512MB Flash
- High-Performance I/O Complement
- three GbE Ethernet ports (+ IPM)
- two PMC/XMC with PCI-X and PCI-E
- optional VME64 interface
- 2 x SIO, 2 x USB 2.0
- std IPM options (1553, serial, SCSI, SATA, TTL & diff'l discretes)
- 4-lane fabric ports, each selectable for SRIO / PCI-E

VMETRO
a Curtiss-Wright Company

Thank you

CURTISS WRIGHT Controls
Embedded Computing

ABOVE & BEYOND

www.cwcembedded.com