Case Study of PCIe® as an On-Board System bus

Eric Esteve

Marketing Manager

PLD Applications

Agenda

- Introduction
- Case Descriptions
- High-speed link selection process
- Architecture
- Conclusion

Introduction

- Audience: we target embedded system designers.
- Their design will have to interface with, or be integrated into, other systems.
- Need for high bandwidth capability for external data transfer.
- Need for high throughput on-board system bus.
- Selected protocol has to support concept of switch fabric.
- FPGA are used in both cases because of:
  - ✓ TTM, Flexibility, Application specific needs

Introduction

- Case 1: I/O module
  - ✓ Limited computing resource
  - ✓ Support data interconnect between:
    - Sensor array or equivalent,
    - Various computing systems
  - ✓ High Bandwidth data link (10 Gbps) supported by a standard protocol
- Case 2: Intelligent adapter card
  - ✓ On board high performance CPU
  - ✓ Support various high-speed protocols
  - ✓ On-board switch for switching large amount of data
  - ✓ Need for one high-speed protocol to be used as a local system bus

Introduction

- Many high-speed protocols available:
  - ✓ RapidIO™, HyperTransport™, Ethernet, XAUI, SPI, Advanced Switching, PCI Express®, PCI/PCI-X®, Aurora, SerialLite…
- So, why select PCI Express?
  - ✓ Performance, scalability, evolution
  - ✓ Expectations from a protocol
  - ✓ Market acceptance
    - Interfacing with various systems
    - Technology well supported (ASIC, FPGA), Commercially Of The Shelf (COTS) product available
  - ✓ Total cost of ownership

Agenda

- Introduction
- Case Descriptions
- High-speed link selection process
- Architecture
- Conclusion

Case 1: I/O module

- Designed for point-to-point or broadcast communication between remote sensor arrays and processing systems
- Four 2.5 Gbps FPDP links (Front Panel Data Port: protocol used in data acquisition)
- Need for aggregate bandwidth of 10 Gbps
- High Bandwidth data link to be supported by a standard protocol
  - ✓ Widely used to be compatible with maximum number of processing systems
  - ✓ Connectors available to support CPU plug in

I/O module: block diagram

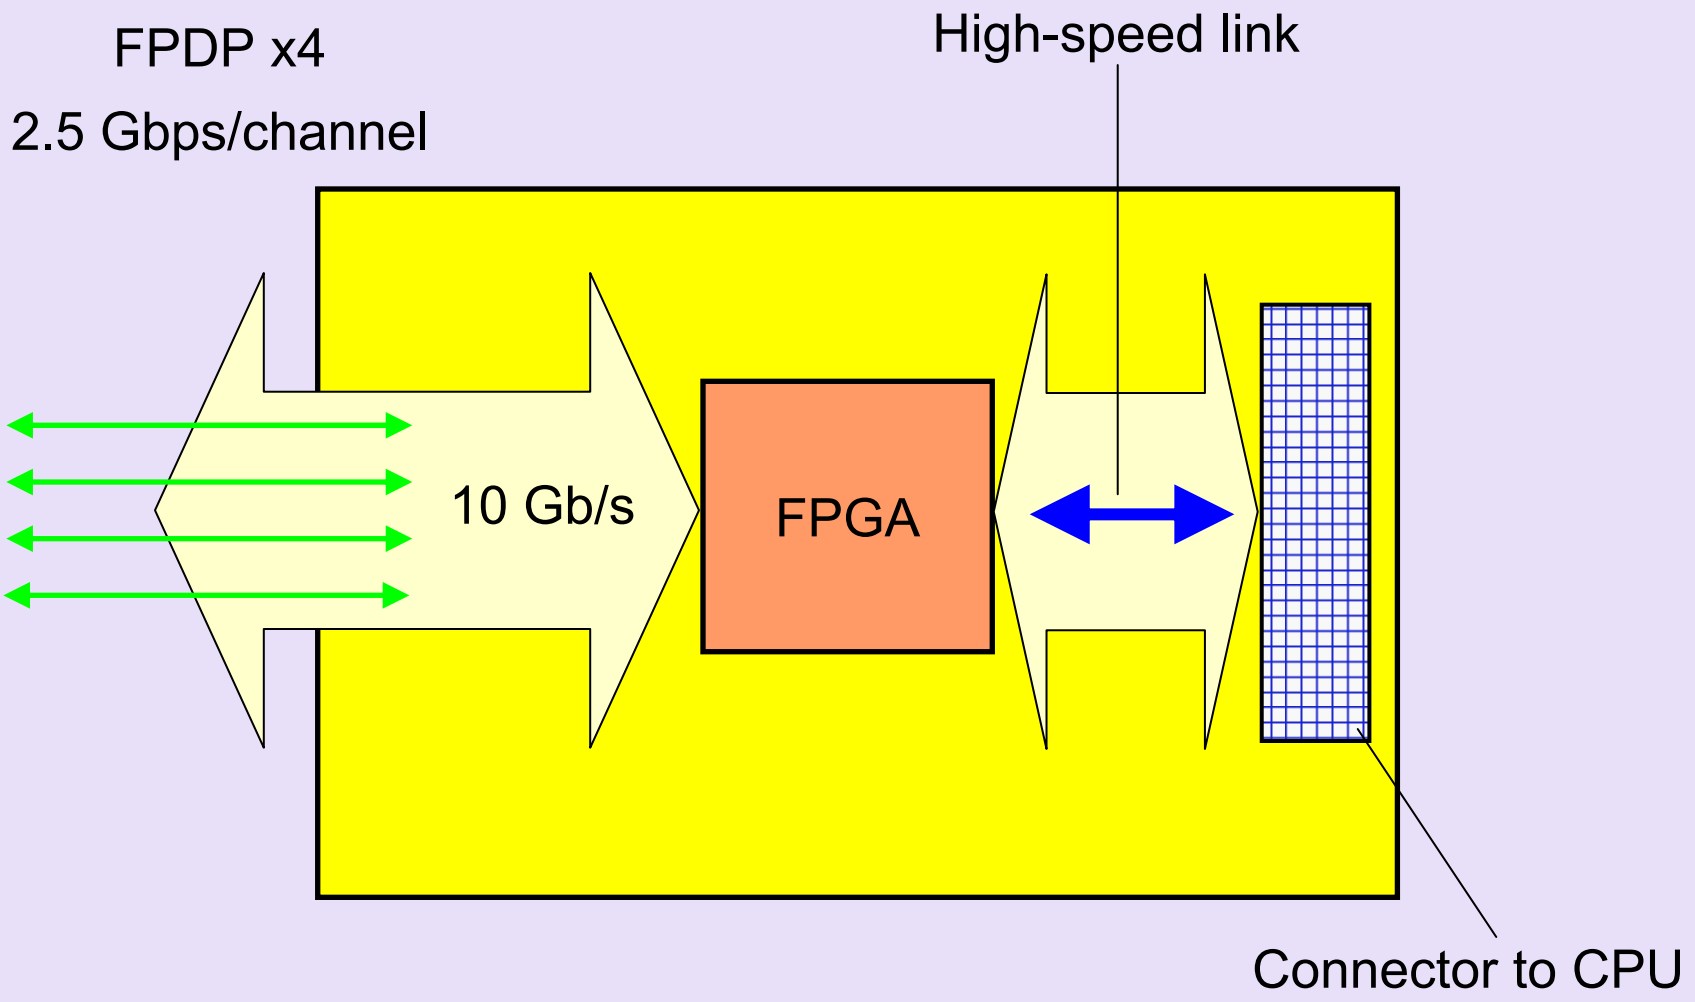

Case 2: Intelligent Adapter card

- On board high performance CPU
- Many high-speed I/O protocols on board
  - ✓ Backplane: two x4 VXS links
  - ✓ 2eSST VME interface
  - ✓ Dual 4 Gb Fibre Channel
  - ✓ Gigabit Ethernet
- Connectivity:
  - ✓ Connector must be standard to support plug-in of Off The Shelf modules
  - ✓ System bus selected has to be compatible with the connector

Intelligent Adapter card

- Functional requirement: switch data flow
  - ✓ Two x4 VXS links to backplane: 1.25 GB/s/direction each
  - ✓ 2eSST VME interface: 320 MB/s
  - ✓ Dual 4 Gb Fibre Channel: 8 Gb/s
  - ✓ Gigabit Ethernet: 125 MB/s
  - ✓ System bus: > 2 GB/s/direction
- Need to build a switch on board
- Need to select a system bus supporting
  - ✓ Switch
  - ✓ More than 1.25 GB/s/direction

Block Diagram

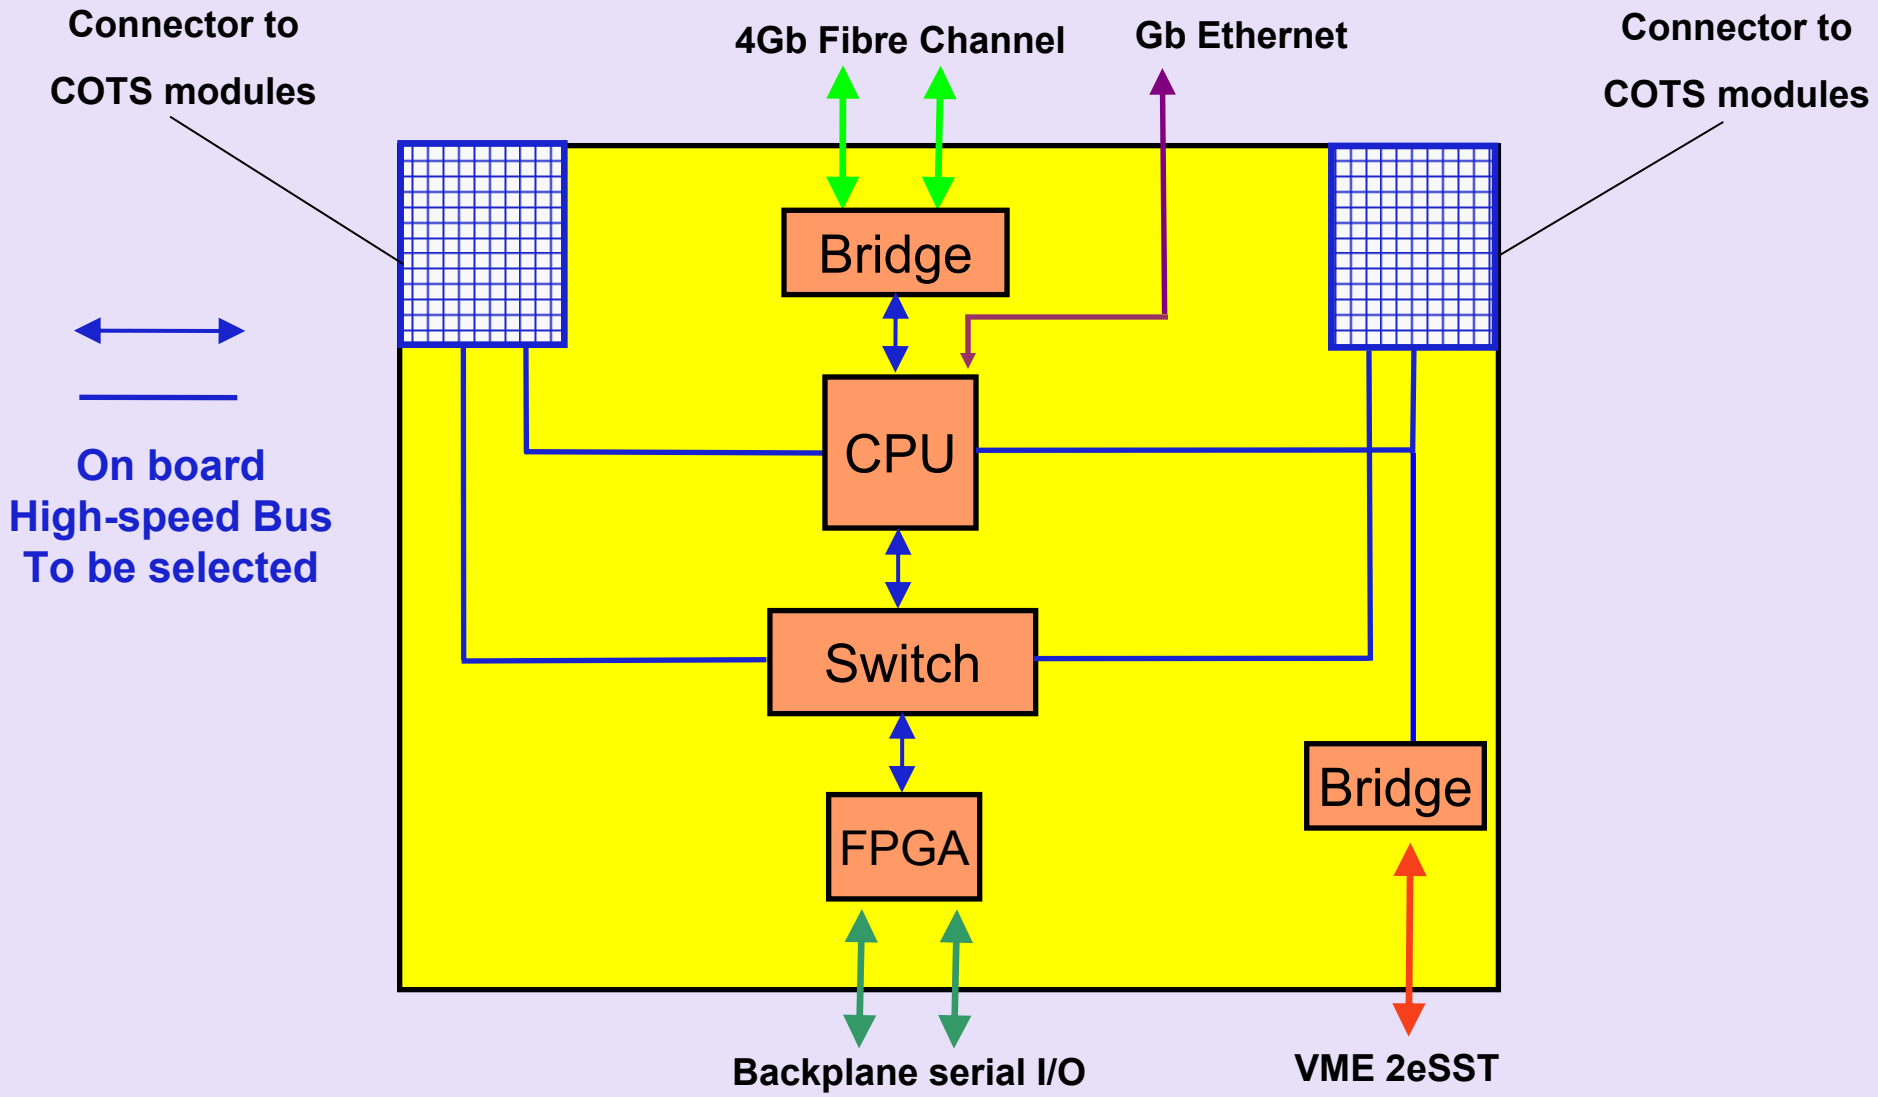

Agenda

- Introduction
- Case Descriptions
- High-speed link selection process
- Architecture
- Conclusion

High-speed link selection process

- Let's look at high-speed protocols available:
  - ✓ RapidIO, Hyper Transport, 10 Gb Ethernet (XAUI), Fibre channel, SPI 4.2, ASI (Advance Switching Interconnect), PCI Express, Aurora, SerialLite II…
- And try to differentiate these protocols in term of:
  - ✓ Does it fit my needs
    - Bandwidth
    - Used as a local system bus
  - ✓ Performance, scalability, evolution
  - ✓ Does it target a specific market segment?
  - ✓ Market acceptance
    - Interfacing with various systems
    - Technology well supported (ASIC, FPGA), COTS available

High-speed link selection process

- High-speed serial proprietary protocols:
  - ✓ Aurora is owned by Xilinx
  - ✓ SerialLite II is owned by Altera
- High-speed protocols
  - ✓ RapidIO, Hyper Transport
- Native high-speed serial protocols
  - ✓ PCI Express 1.1
  - ✓ Fibre Channel: Storage
  - ✓ 10 Gb Ethernet (XAUI): Networking
  - ✓ SPI 4.2: Networking and Telecom: ATM, SONET/SDH,…
  - ✓ ASI: switched fabric architecture (High End Communication, Computer, Storage products)

High-speed link selection process

- Shortlist (non proprietary, non market specific):
  - ✓ Serial RapidIO
  - ✓ Hyper Transport 3.0
  - ✓ PCI Express 1.1
- Select the most appropriate for our embedded system application…
- Switch fabric capability

High-speed link selection process

Protocol	Architecture & Market
Topology	Flow Control
First Spec	Peer to Peer
Bus Width	Hot Plug
Data Rate	Evolution
Max Bandwidth	Protocol agent
Protocol	Switch Agent
Priority Levels	Target market
Max Payload	Market acceptance
Latency	IP available
Virtual Channels	COT availability

HT 3.0: protocol

Topology	Daisy-chained
First Spec	3Q01
Bus Width	2, 4, 8, 16 bits
Data Rate	Up to 2.6 Gps DDR
Max Bandwidth	10.4GB/s (16 bits)
Protocol	Layered packets
Priority Levels	2 priorities
Max Payload	64 bytes
Latency	Low
Virtual Channels	3+16 VCs

HT 3.0: architecture & target market

Flow Control	Yes
Hot Plug	Yes (3.0 only)
Evolution	Bus: 32 bits
Protocol agent	Host – Tunnel – Slave Bridges to PCI/PCI-X/PCIe
Switch Agent	Yes
Target market	PC Chipset – Graphic – Gaming – Servers (PCIe is used complementary to HT for peripheral)
Market acceptance	Networking – Storage – Blade servers (3.0)
IP available	2 FPGA vendors (400 MHz DDR only) – Very few IP vendors
COT availability	HT 3.0: First CPU or ASSP available in 1Q06

Hyper Transport: topology

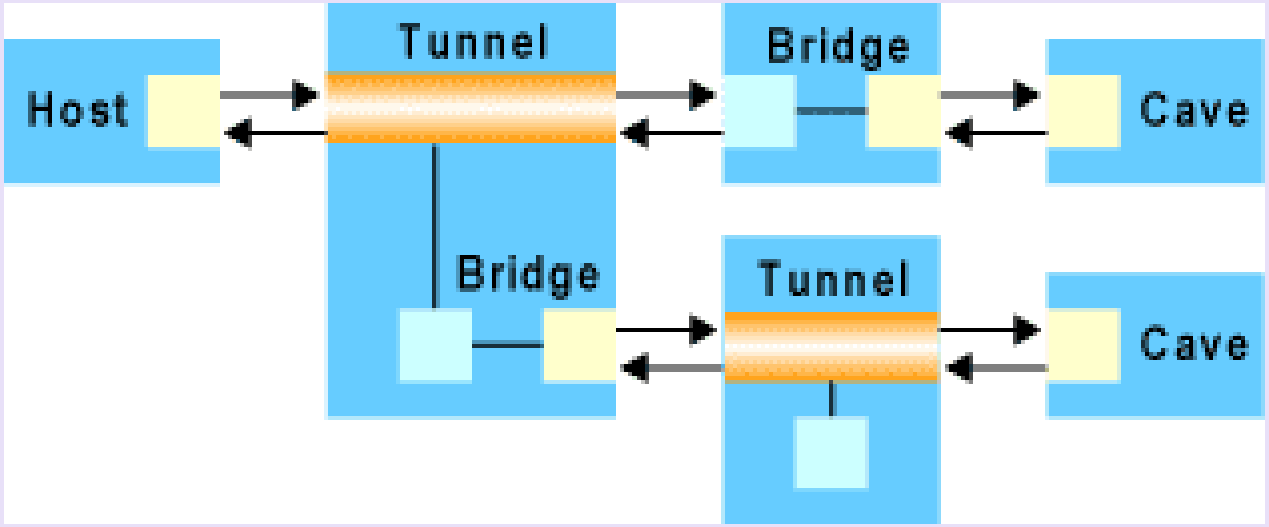

Serial RapidIO: protocol

Topology	switched
First Spec	4Q01
Bus Width	1, 4 Lanes
Data Rate	1.25/2.5/3.125Gbps
Max Bandwidth	1.25GB/s (4 lanes)
Protocol	Layered packets
Priority Levels	4 priorities
Max Payload	256 bytes
Latency	Medium
Virtual Channels	2 VC

sRIO: Architecture & target Market

Flow Control	Point-to-point + end-to-end
Hot Plug	No
Evolution	5 Gbps, 6.25 Gbps
Protocol agent	I/O Master – I/O slave Bridge to Ethernet, PCI, PCI-X
Switch Agent	Yes
Target market	Embedded – Wireless Communication Multi-DSP or PowerPC interconnection
Market acceptance	Not embedded in support function chips (Communication controller or security uP)
IP available	3 FPGA vendors – Several IP vendors
COT availability	First DSP, CPU and Switches available

sRIO: topology

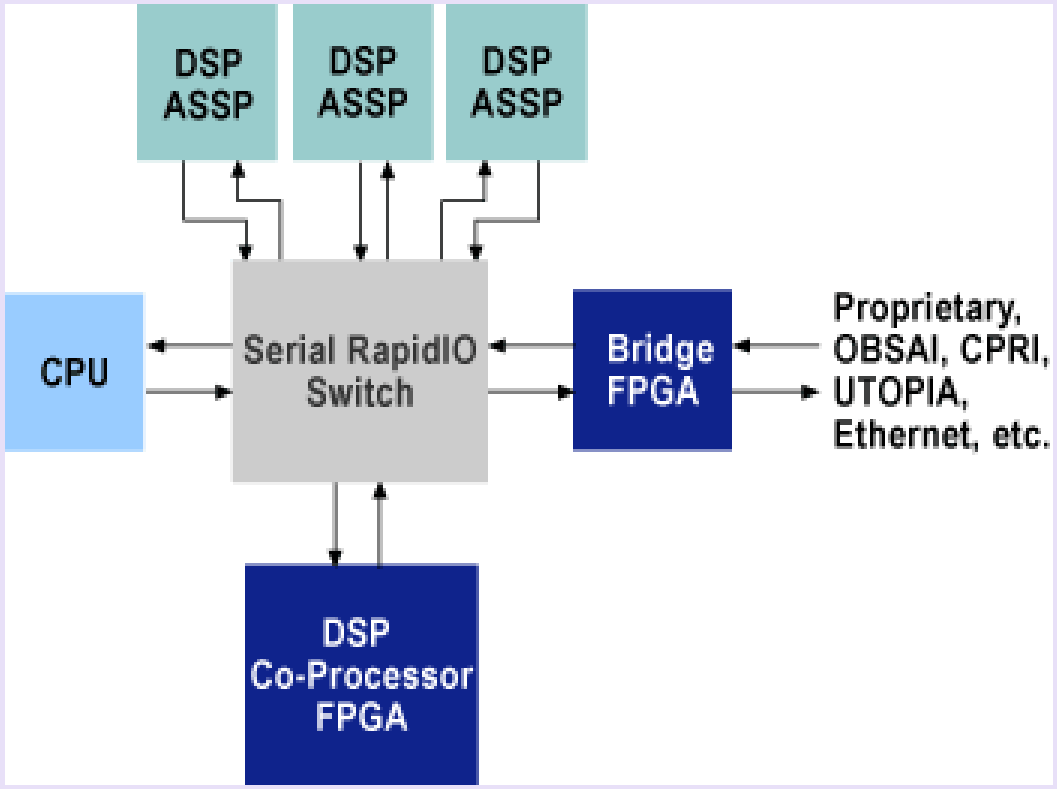

PCIe 1.1: protocol

Topology	Tree
First Spec	3Q02
Bus Width	Up to 32 lanes
Data Rate	2.5Gbps/direction
Max Bandwidth	8GB/s (32 lanes)
Protocol	Layered packets
Priority Levels	8 priorities
Max Payload	128 to 4,096 bytes
Latency	Medium
Virtual Channels	8 VC

PCIe 1.1: architecture & target market

Flow Control	Yes
Hot Plug	Yes
Evolution	-Data rate: 5.0 Gb/s /direction (PCIe 2.0) -IOV
Protocol agent	Root Port – End Point – Shared (EP/RP) Bridges to PCI, PCI-X
Switch Agent	Yes
Target market	PC – Graphic - Servers - Embedded – Communication – Storage – Security processors
Market acceptance	Printers – Networking – Switching – Industrial
IP available	3 FPGA vendors – Several IP vendors
COT availability	-Several vendors of dedicated ASSP (Switch, Bridges) -Integrated in 100's of ASIC, ASSP and CPU

PCIe: topology

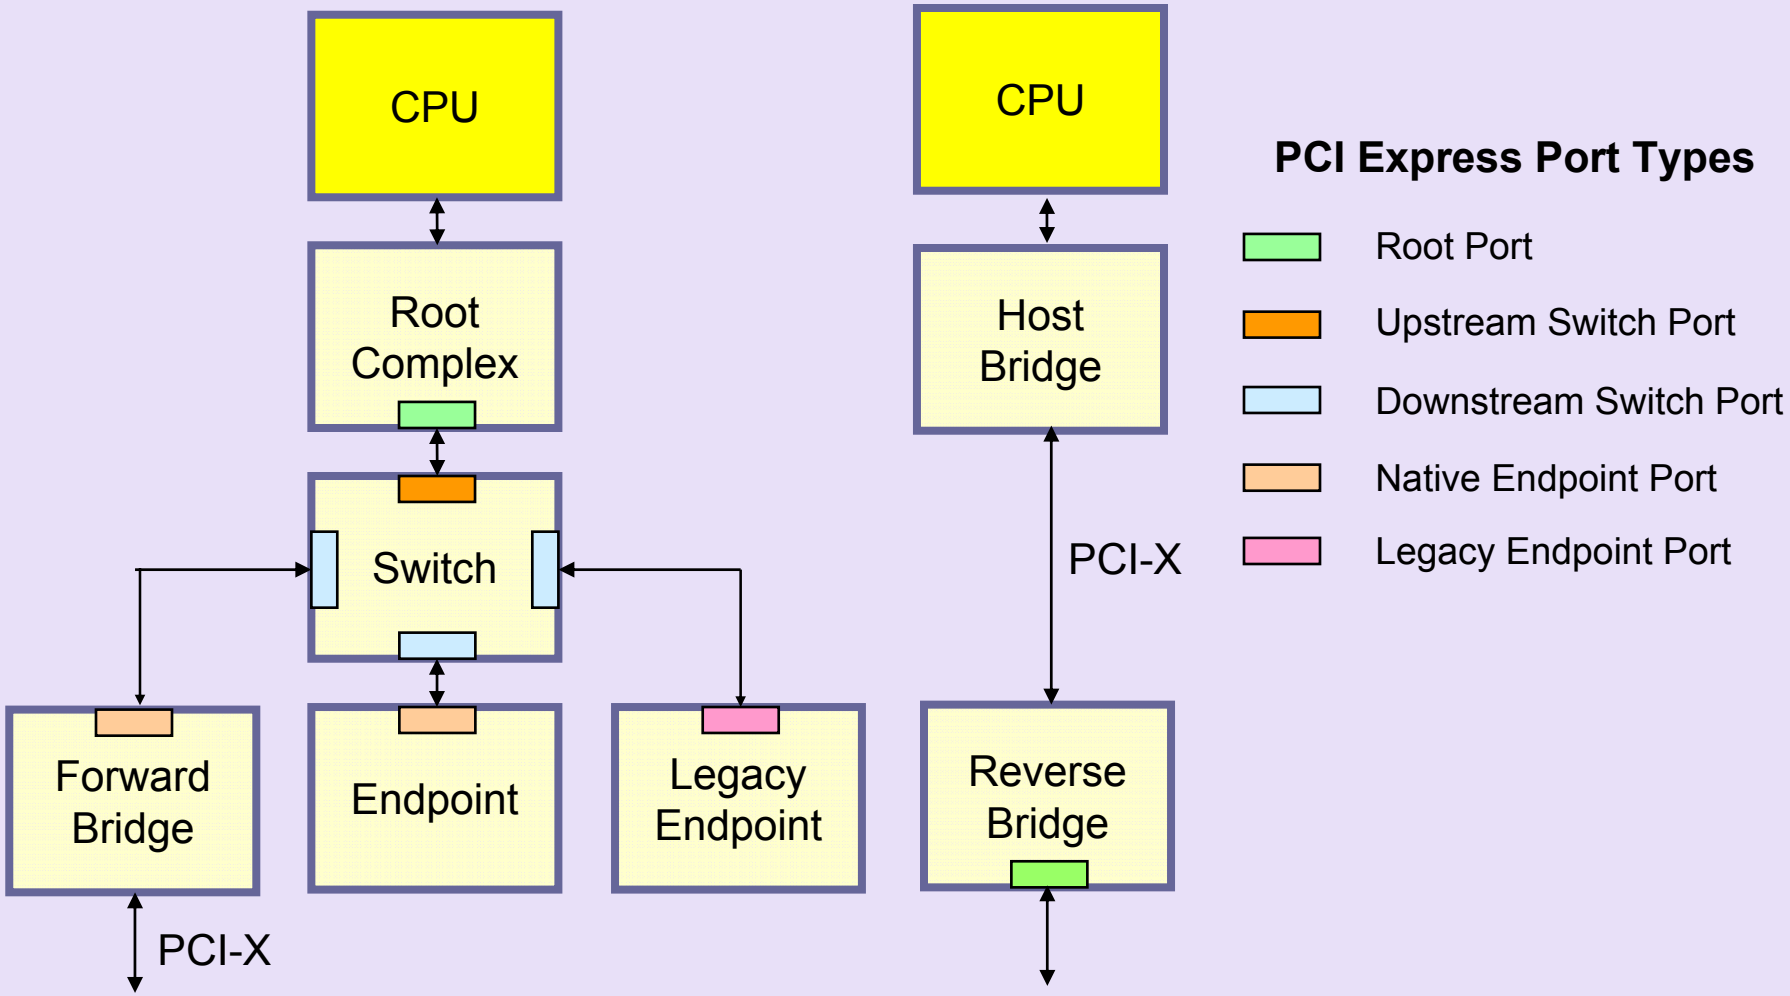

Comparison

	HT 3.0	sRIO	PCIe 1.1
Topology	Daisy-chained	Switched	Tree
First Spec	3Q01	4Q01	3Q02
Bus Width	2, 4, 8, 16 bits	1, 4 Lanes	1, 2, 4, 8, 16, 32 lanes
Data Rate	Up to 2.6 GHz DDR (800 Mbps FPGA)	1.25/2.5/3.125Gbps	2.5Gbps/direction
Max Bandwidth	10.4GB/s (16 bits) (1.6GB/s FPGA)	1.25GB/s (4 lanes)	8GB/s (32 lanes)
Protocol	Layered packets	Layered packets	Layered packets
Priority Levels	2 priorities	4 priorities	8 priorities
Max Payload	64 bytes	256 bytes	4,096 bytes
Latency	Low	Medium	Medium
Evolution	Bus 32 bit	5 Gbps, 6.25 Gbps	5.0 Gbps, IOV
Virtual Channels	3+16 VC	2 VC	8 VC

PCI Express selection

- Selection based protocol features is not effective: all three are close
- Target markets:
  - ✓ Hyper Transport: PC Chipset - Graphic - Gaming - Servers
  - ✓ sRIO: Embedded - Wireless Communication - (DSP and PowerPC)
  - ✓ PCIe: PC Chipset - Graphic - Peripheral - Servers - Embedded - Security processors
- Market penetration (beside target market) is higher for PCIe
  - ✓ Industrial – Medical – Networking - Communication – Storage
  - ✓ Because lower cost of ownership (Price, PCI/PCI-X S/W compatibility)
- PCI Express is a general-purpose bus
- Widely used in the industry:
  - ✓ A lot more COTS available
  - ✓ A single product can get better market penetration

PCI Express selection

- PCI Express can be used as for:
  - ✓ peripheral device interconnect,
  - ✓ chip-to-chip interconnect,
  - ✓ bridging to other interconnects:1394b, serial RapidIO, USB2.0, InfiniBand™, Ethernet, PCI, PCI-X, …
  - ✓ Design an application specific switch on board
- Can be used as a local system bus
- Backward compatible with existing PCI software model

PCI family

	Width	Transfer rate	Bandwidth
PCI 1.0	32 bits	33 MHz	133MB/s
PCI 2.x	64 bits	33-66 MHz	266-533MB/s
PCI-X 1.x	64 bits	133 MHz	1 066MB/s
PCI-X 2.0 *	16, 64 bits	266, 533 MHz	Up to 4GB/s
PCIe	1-32 lanes	2.5 GHz	Up to 8GB/s /direction

* **Not supported by FPGA technologies**

Agenda

- Introduction
- Case Descriptions
- High-speed link selection process
- Architecture
- Conclusion

Case 1: I/O module

- Four 2.5 Gbps FPDP
- Need for aggregate bandwidth of 10 Gbps
- PCIe x4 solution:
  - ✓ Offer a bandwidth of 10 Gbps/direction
  - ✓ Interface directly with XMC (Switched Mezzanine Card) connectors
  - ✓ Backward S/W compatibility also permits us to use PCI/PCI-X
- PCI-X 64/133 to target a wider market
  - ✓ Offer a similar bandwidth: 1.067 GB/s
  - ✓ Interface with PMC (Peripheral Mezzanine Card) connectors

I/O module: PCIe architecture

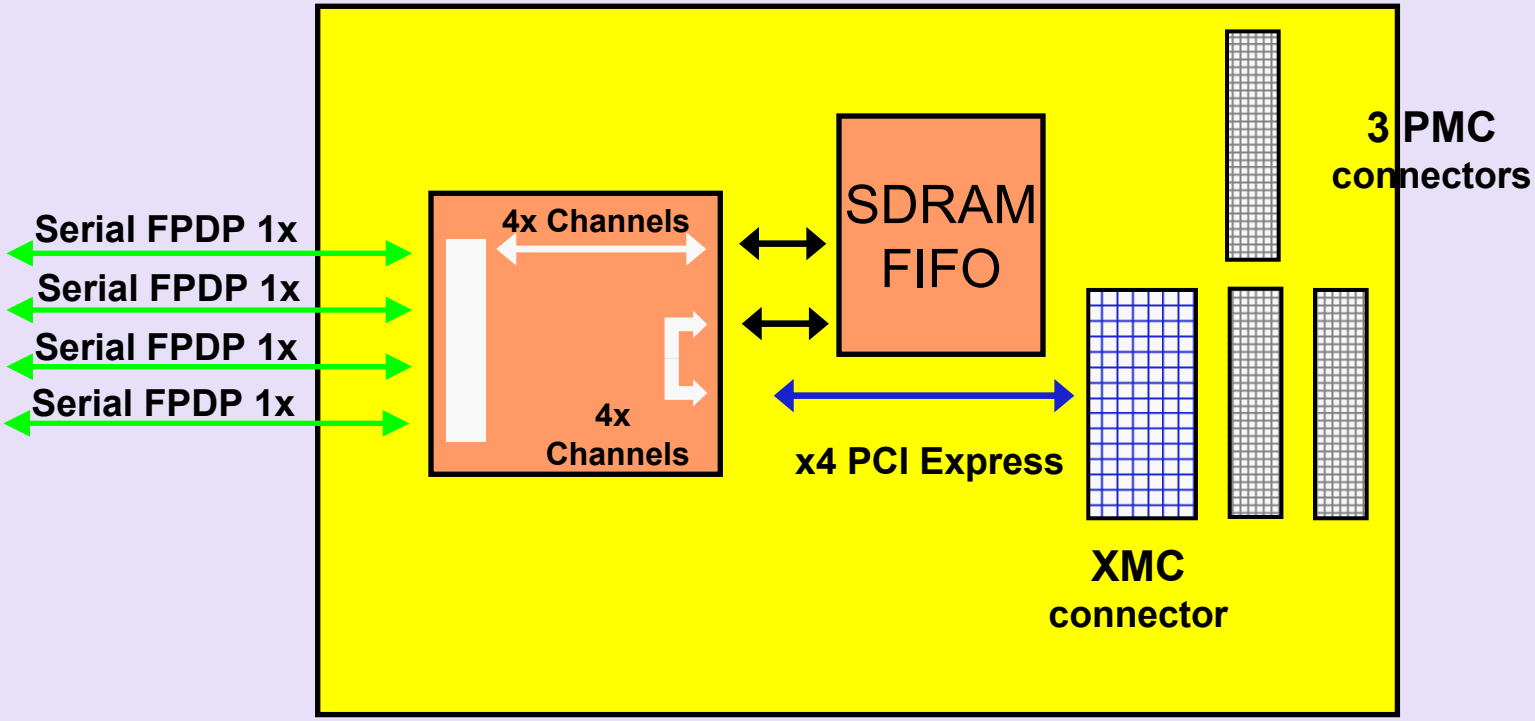

I/O module: PCI-X architecture

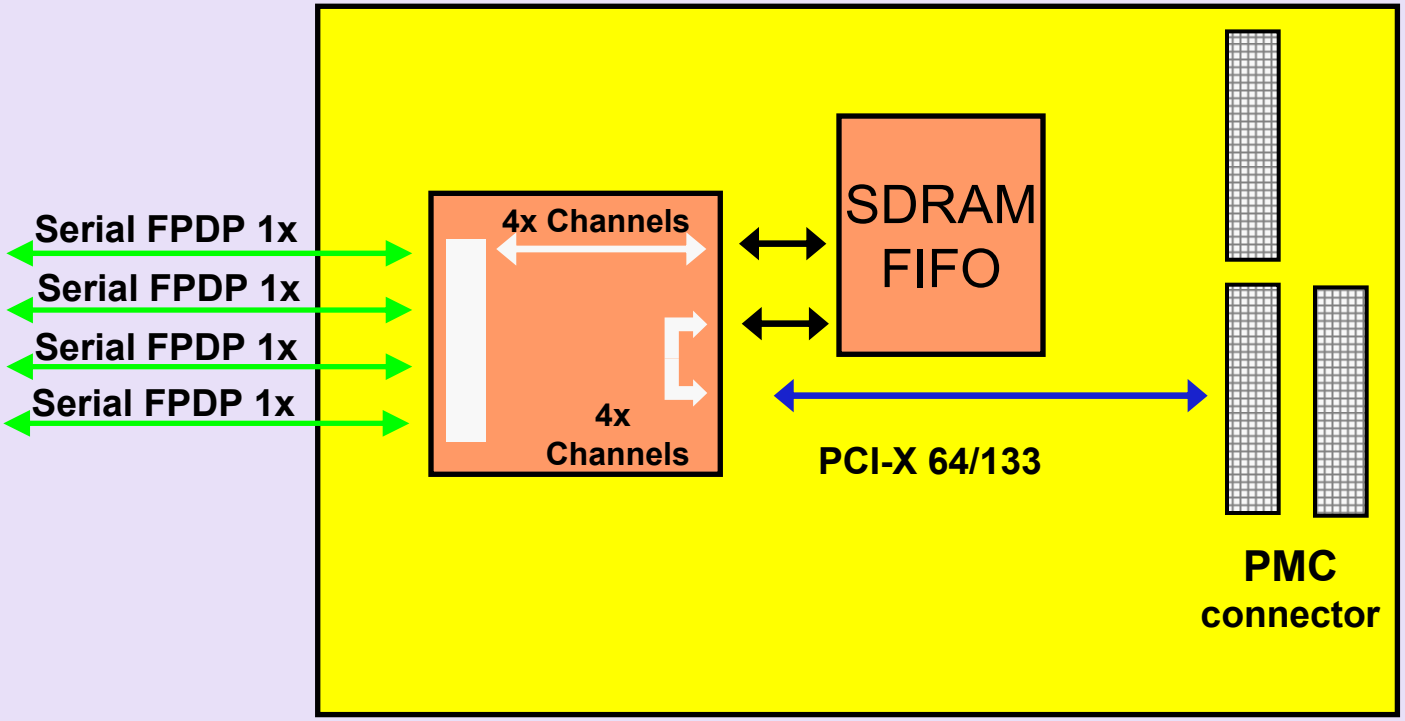

SFM Serial FPDP

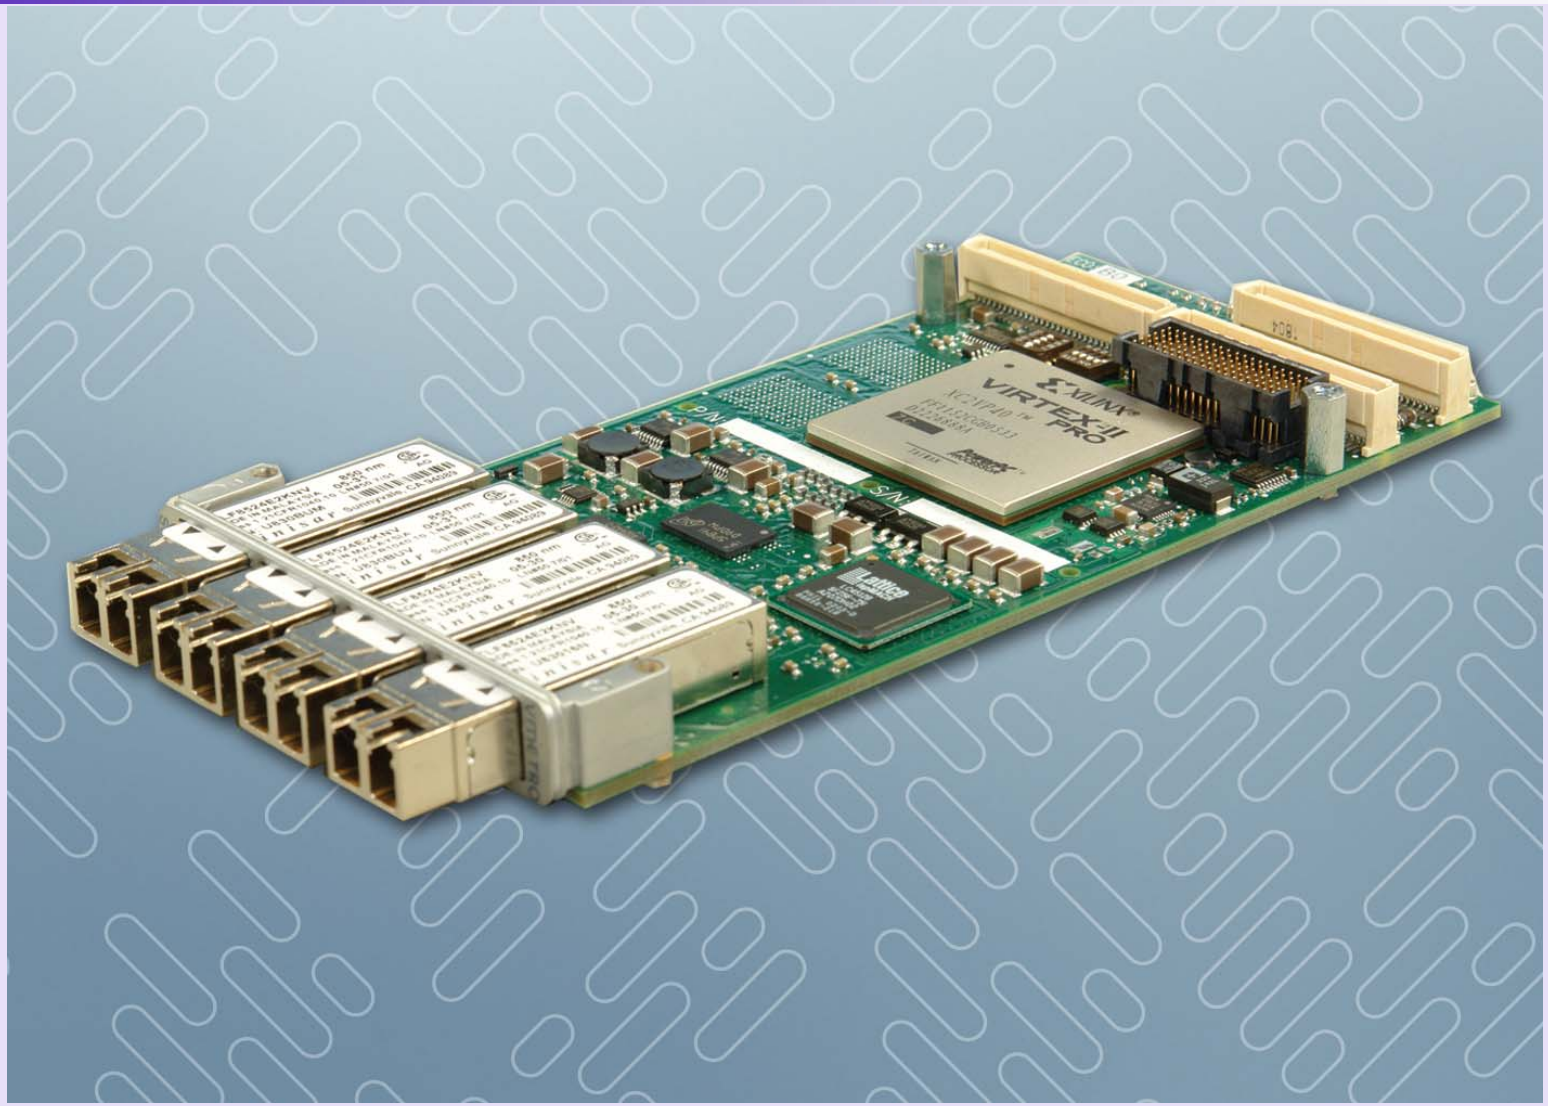

Case 2: Intelligent Adapter card

- On board CPU
- High-speed protocols:
  - ✓ Two x4 VXS links to backplane
  - ✓ 2eSST VME interface
  - ✓ Dual 4 Gb Fibre Channel
  - ✓ Gigabit Ethernet
- High speed link/On board system bus:
  - ✓ Two x8 PCI Express…
  - ✓ … also two PCI-X 64-bit/133 MHz
- Connectivity:
  - ✓ Dual XMC (switched mezzanine card) sites
  - ✓ Dual PMC (PCI mezzanine card) sites

block diagram (1)

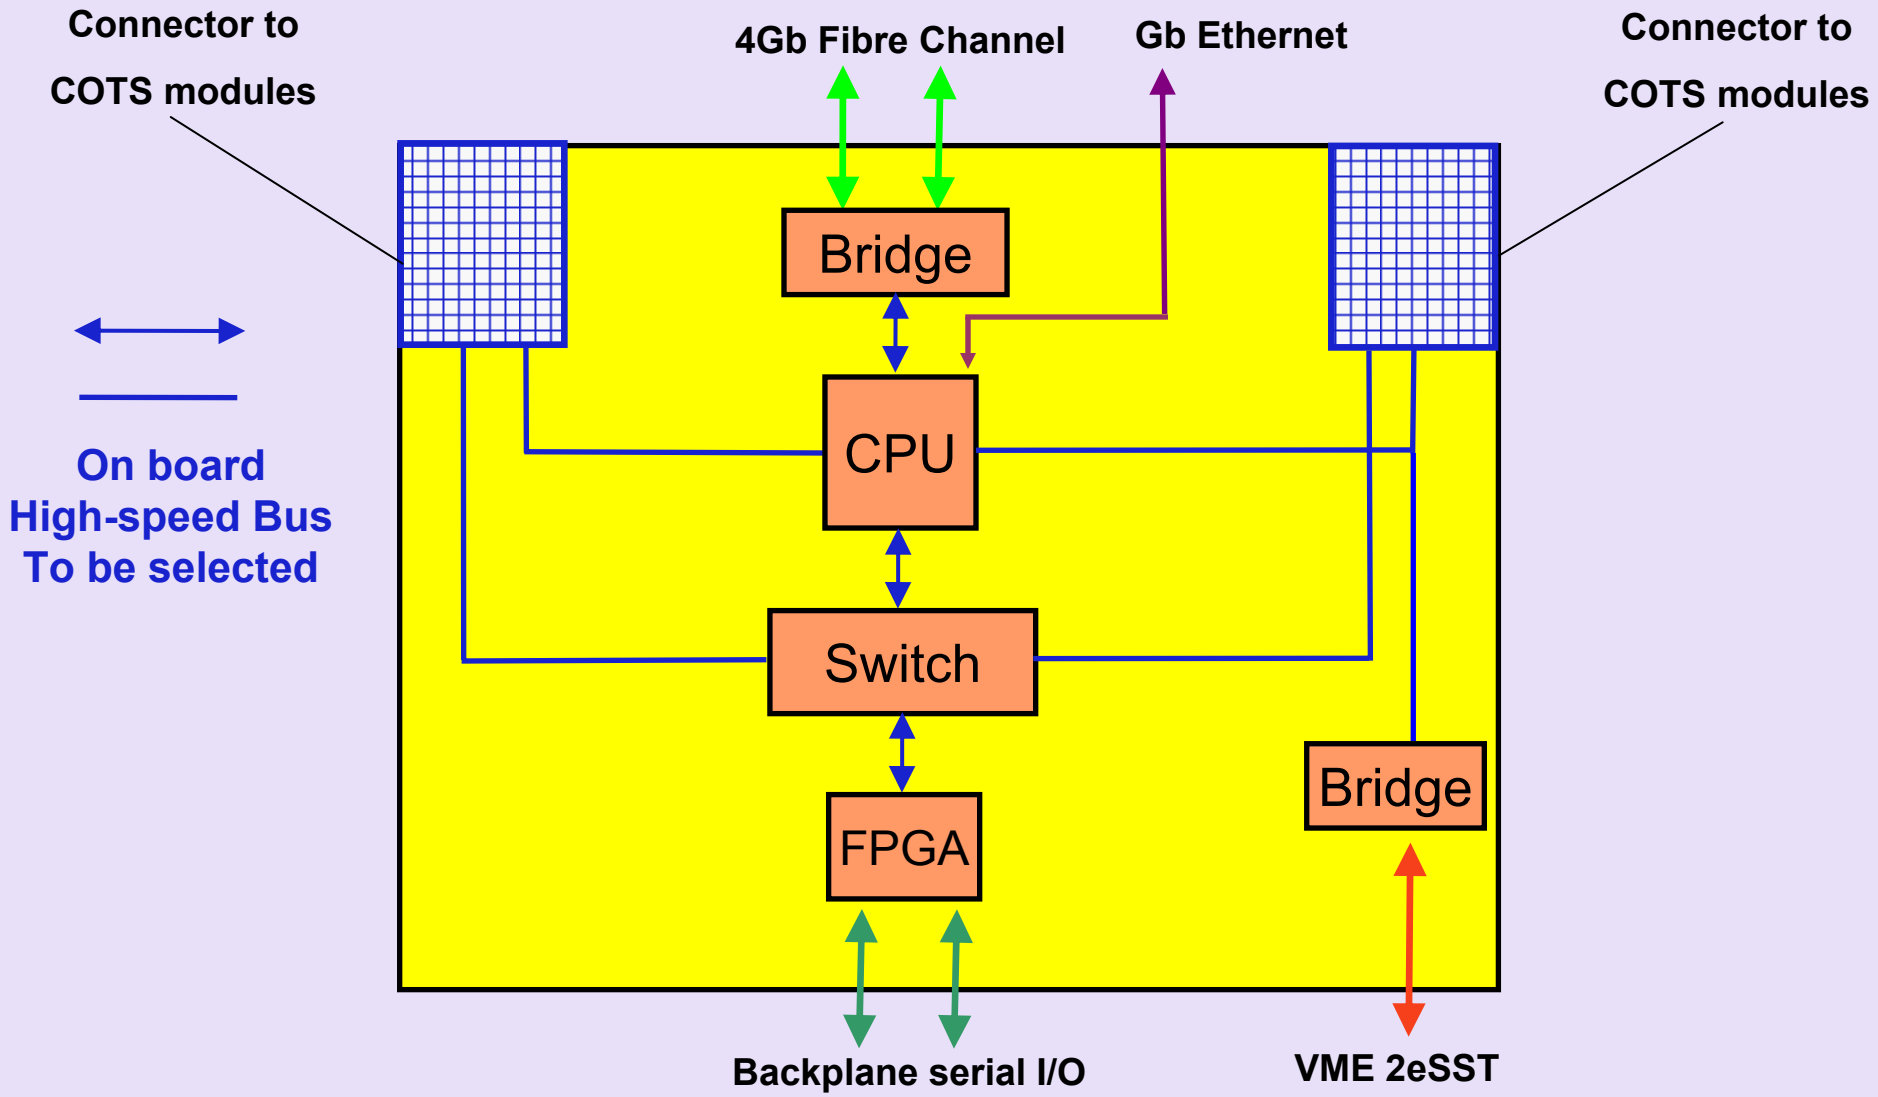

block diagram (2)

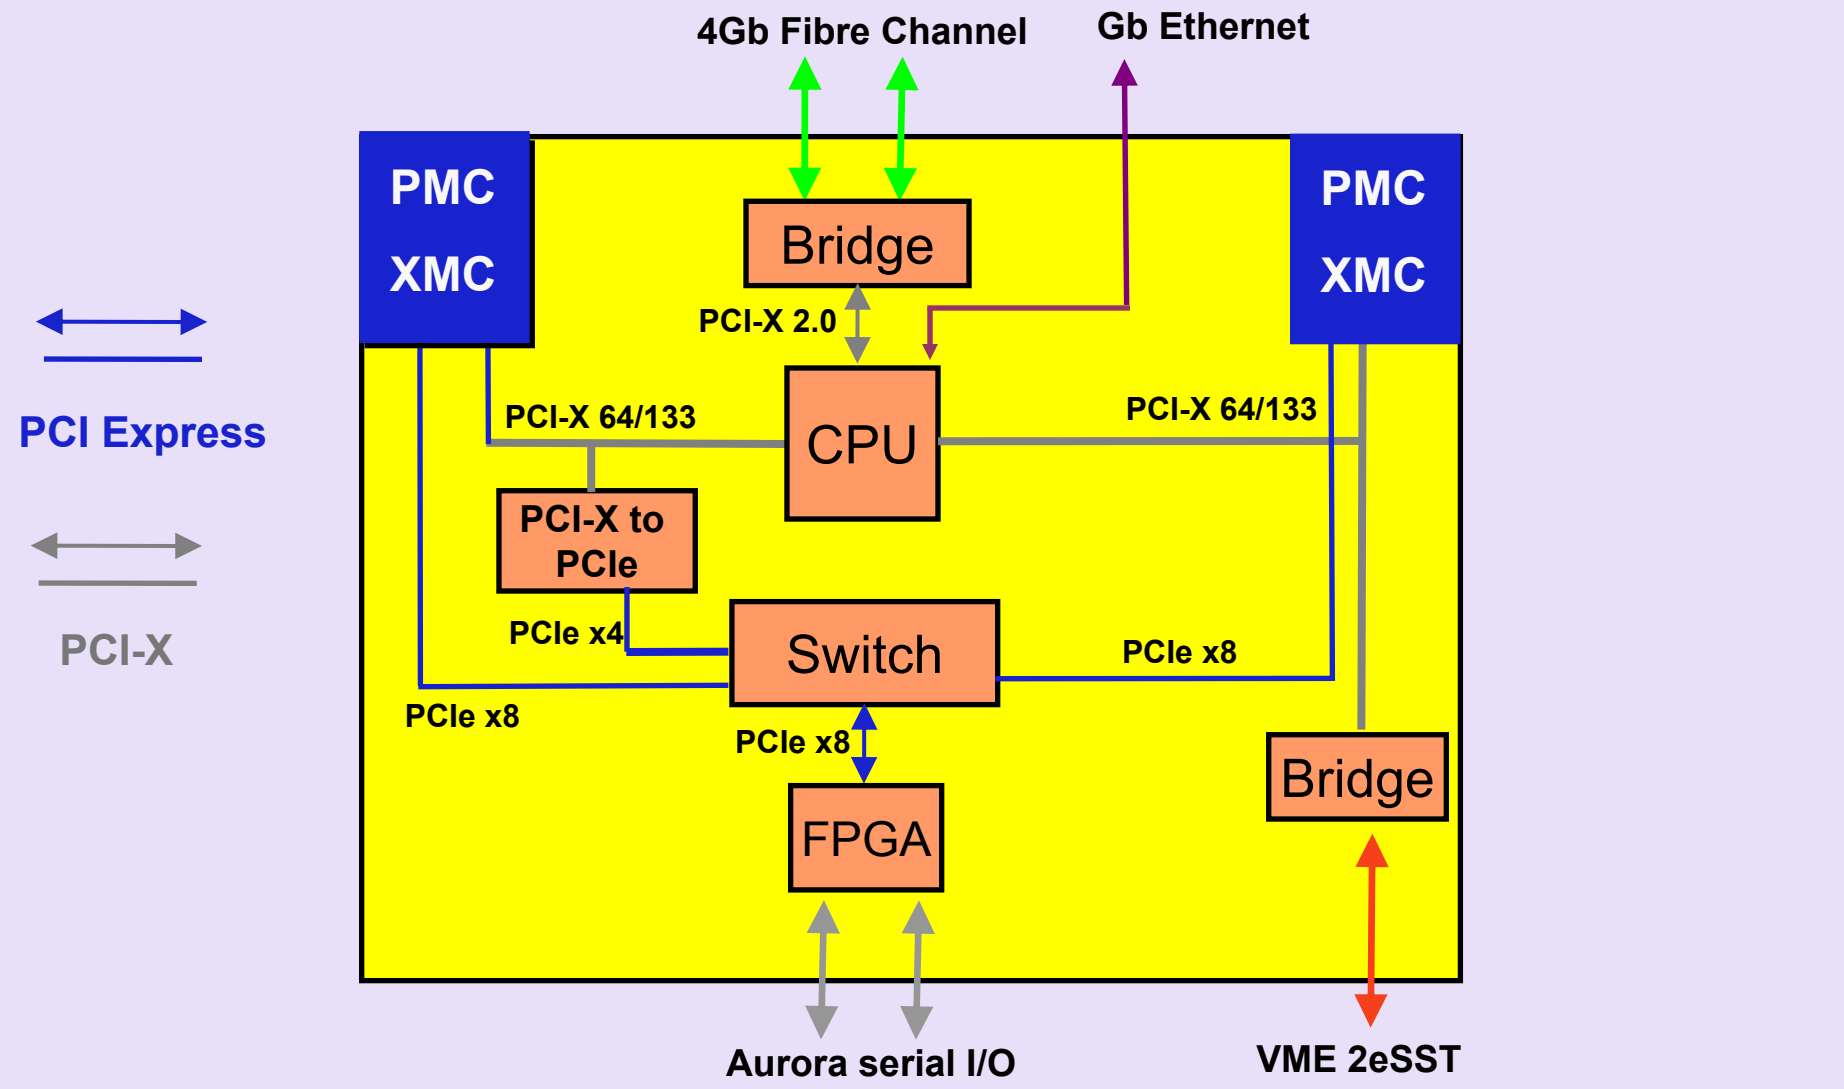

PCIe topology: using COTS only

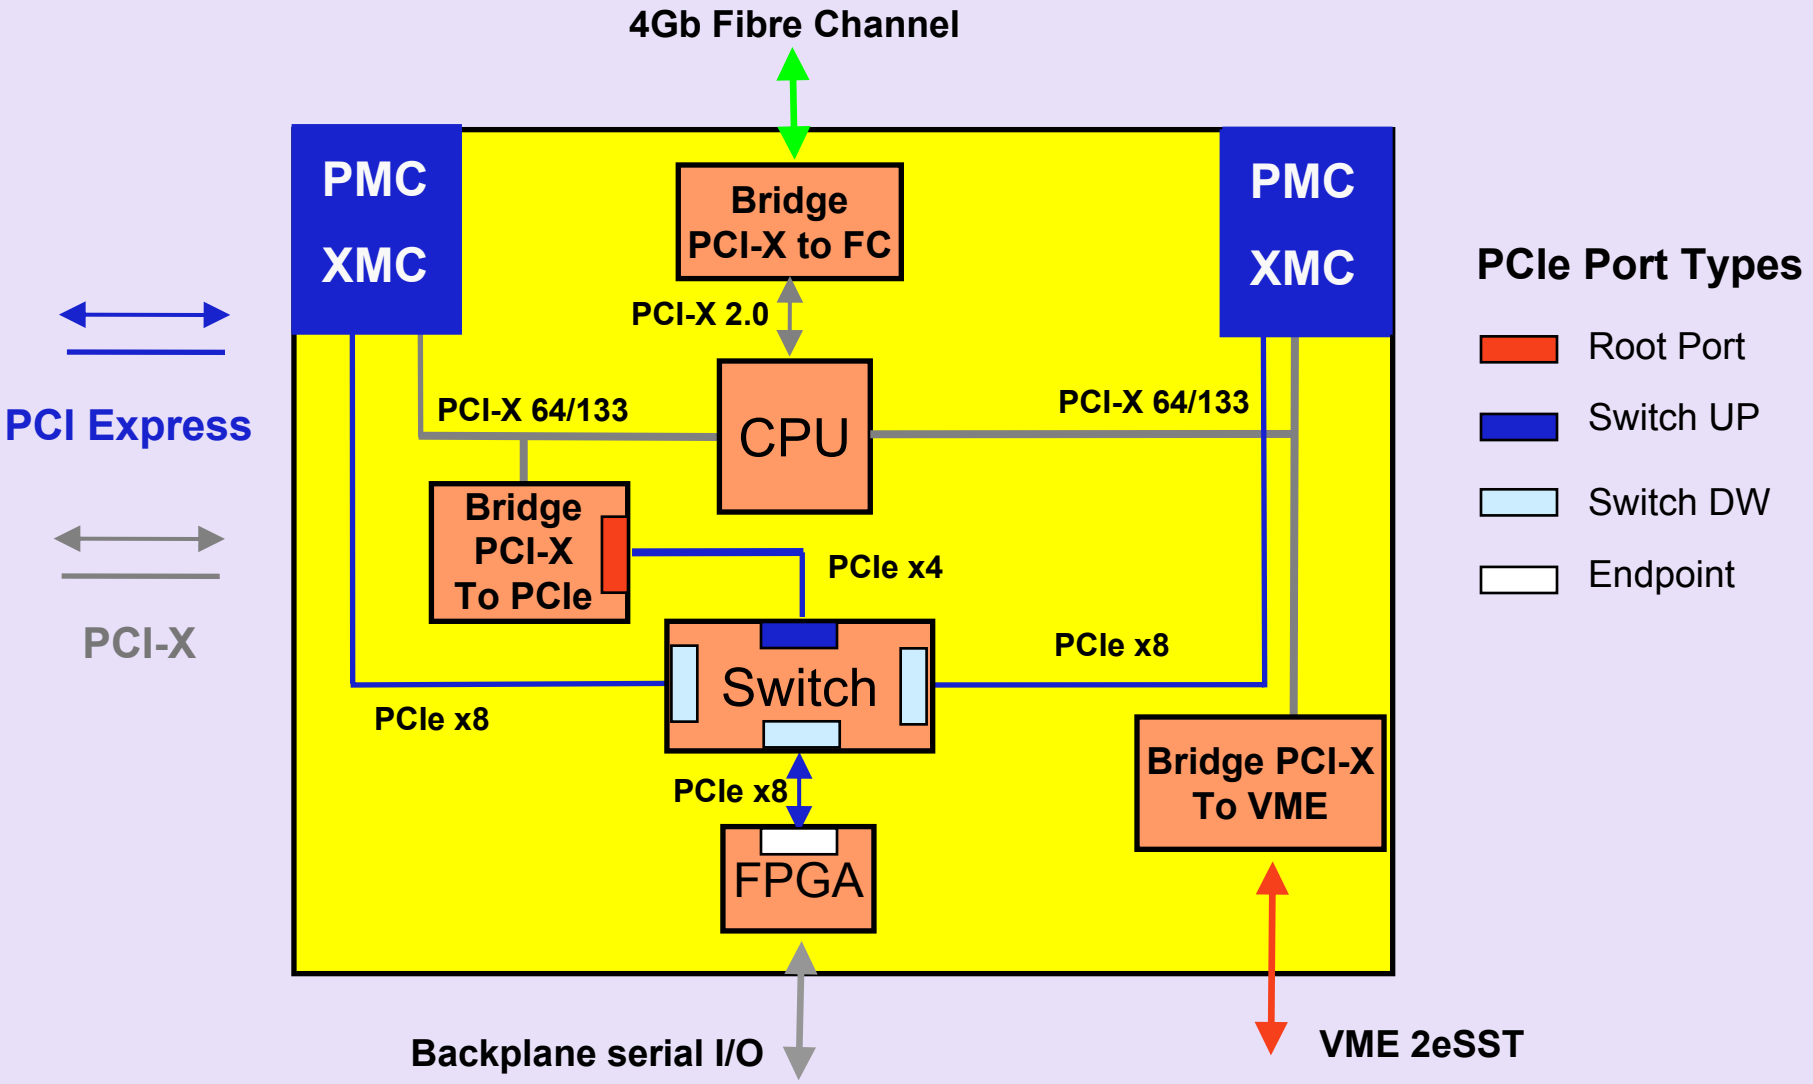

M6000 (Intelligent Adapter Card)

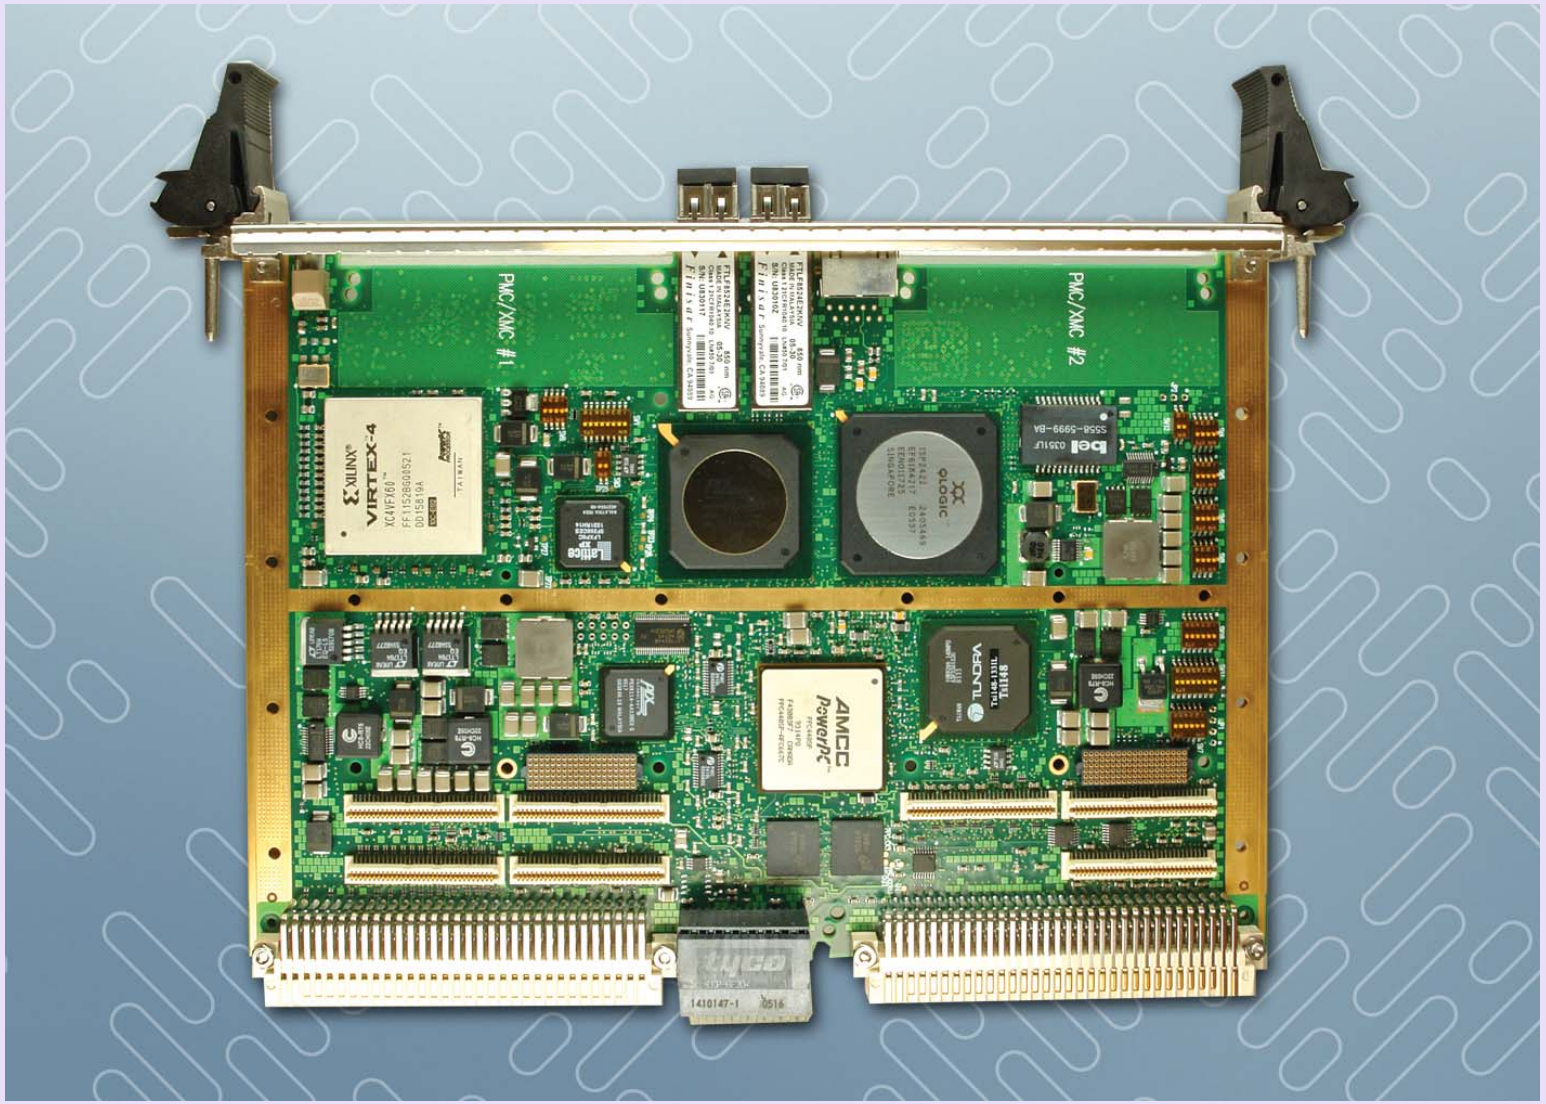

Agenda

- Introduction
- Case Descriptions
- High-speed link selection process
- Architecture
- Conclusion

Conclusion

- Many high-speed interconnect bus choices available
- Some can be the best solution for niche market segments
- PCI Express protocol was defined from the beginning
  - ✓ for serial high-speed I/O
  - ✓ To support switches and bridges
- PCI Express is the most widely used high-speed serial protocol
  - ✓ Allow to integrate a sub-system into various systems
- PCI Express can be used as a local system bus
  - ✓ Allow to build a complete sub-system
- Best use of PCI Express agents
  - ✓ End Point, Root Port, Switch and Bridge
  - ✓ Allow to easily build a communication sub-system

Thanks to:

- **VMETRO** for permitting us to highlight commercially available products that use:
  - ✓ PLDA's PCI-X 64/133 IP controller
  - ✓ PLDA's PCIe x4 and x8 IP controller

- **The Linley Group,Inc** as useful information has been extracted from:
  - ✓ *A Guide to High-Speed Interconnect* – Apr 2006

Thank you for attending the
PCI-SIG Developers Conference 2006.

For more information please go to
www.pcisig.com

Case Study of PCIe as an On-board System bus

Eric Esteve

Marketing Manager

PLD Applications