

VXFabric™: PCI-Express Switch Fabric for High Performance Embedded Computing

Robert Negre
robert.negre@kontron.com

Overview

The advent of fast communication connectors and high speed serial lines in computers opens a new era of rugged embedded computing. For example, VPX allows board computers to move away from decades of parallel bus architectures on the backplane and implements high speed serial link point-to-point connections between boards.

Multi-gigahertz signals like PCI Express (PCIe) are poised to become a real disruptor for many HPEC applications. High Performance Computing is the first domain to benefit from the tenfold increase in I/O bandwidth between computing boards, unleashing a new kind of unparalleled application for sensor data processing platforms used in radar, sonar, and general imaging.

VXFabric™ is a crossbar based on the native PCIe links connected to the computer nodes: it results in an efficient multi-node interconnect with high performance and low latency features. A key differentiator of VXFabric™ lies in the fact that it does not need any specific software tools, libraries or environment. The IP socket API makes VXFabric™ as easy as plug and play to use as an Ethernet network.

Communication over PCI Express

PCIe is a computer expansion serial link standard initially designed to replace the older PCI and PCI-X I/O peripheral interfaces. The PCIe topology is based on point-to-point serial links, rather than a shared parallel bus architecture. PCIe technology has evolved recently from Gen1 to Gen2, doubling the performance, and Gen3 is on the horizon which will again double throughput. Therefore PCIe is becoming a fast natural and native communication link between computing devices, not just I/Os. Built on bi-directional lanes used in aggregates from x1 to x32, PCIe provides huge bandwidth capabilities. For example, PCIe x4 Gen2 supports 4GB/s. (e.g. 32Gb/s)

Interconnecting CPU nodes means memories interconnect (through DMA or PIO) of each node, providing address translations and protection, DMAs, Doorbell interruptions, scratchpad registers and flow control, handled by PCIe switches in Non-Transparent modes.

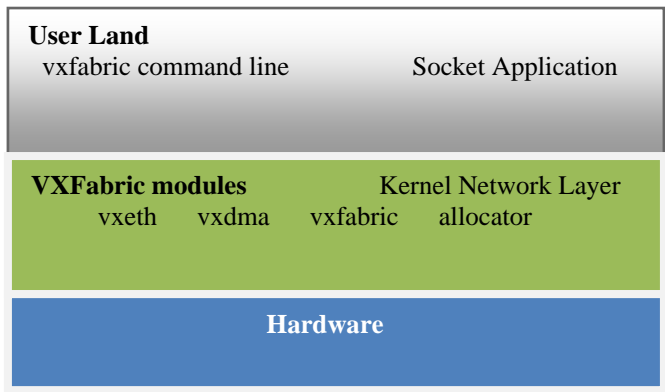
VXFabric™ technology (currently deployed with Kontron VPX boards ecosystem) gives a description of the various data communication APIs proposed by Kontron as well as real life performance figures. Those APIs insulates applications from the low level details of the current generation of PCIe silicon management. The use of the IP

socket API is protecting application software investment. Therefore, VXFabric™ facilitates migration path from 1Gb Eth, 10Gb Eth and 40Gb Eth communication on the backplane.

Current VXFabric™ implementation can interconnect simultaneously up to 12 nodes. The physical interconnect through the backplane can be made in various implementations:

- A centralized topology, using a PCIe switch.
- A distributed PCIe backplane can be used without the need of any additional PCIe switch

Software wise, the VXFabric™ is equivalent to an Ethernet network infrastructure mapped over a switched PCIe express fabric. VXFabric™ implements the layers allowing the user to handle the communication with a IP socket programmatic interface. From this API, all classical protocols like TCP or UDP are directly useable on VXFabric™, and require no modification of existing applications.



Conclusion

In conclusion, VXFabric™ implements a brand new switch fabric based on the PCI-Express using the plug and play capability of this fast link able to move data at ultra-high speed. PCIe presents the key advantage of being a native data bus in all modern processor chipsets. As a consequence, the software environment and ecosystem based on PCIe are huge and allow a large basic interconnect with peripherals. VXFabric™, taking advantage of the PCIe express Non Transparent switching capabilities, works out a peer-to-peer interconnect, allowing all processors of a VXFabric™ system to communicate in the crossbar.