kontron

VXFabric: PCI-Express Switch Fabric for HPEC

Poster B.7, Technologies and Systems

Robert Negre,
Business Technology Director
Kontron Modular Computers SAS
France

VXFabric: PCI-Express Switch Fabric for HPEC

- VXFabric is a new PCIe interconnect between multiple processor nodes
- VXFabric is based on Internet Protocol over PCI-Express
  - At the hardware level, it is built on PCIe Gen2/Gen3 links using non transparent PCIe bridges, allowing peer-to-peer DMAs. Implemented on Open VPX computers & PCIe switch boards
  - At the software level, VXFabric provides a socket API giving access to TCP/IP or UDP/IP protocols, Low latency, Light weight protocol, Plug & play thanks to IP sockets

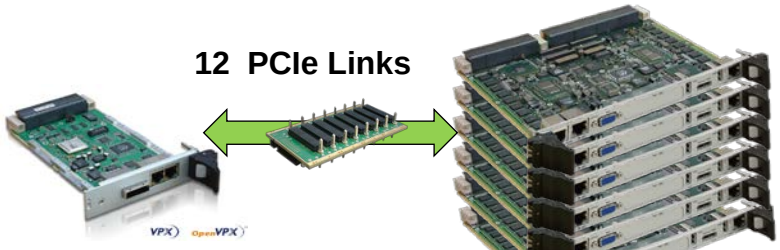

PCIe centralized switch

Six VX6060 = 12 processors = 24 cores

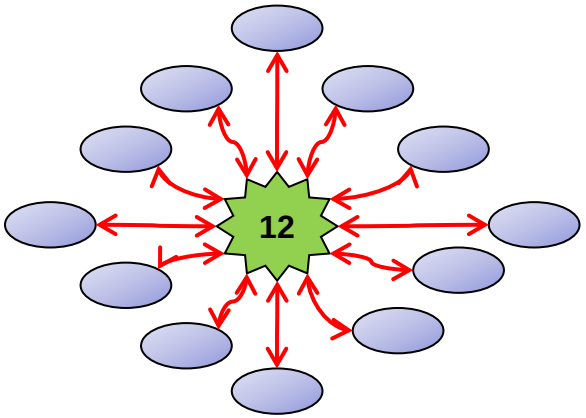

VXFabric Software API

- **VXFabric user programming model is based on a socket Internet Protocol allowing direct use of TCP/IP or UDP/IP**
  - **A low level protocol takes care of data transfers between processor nodes**

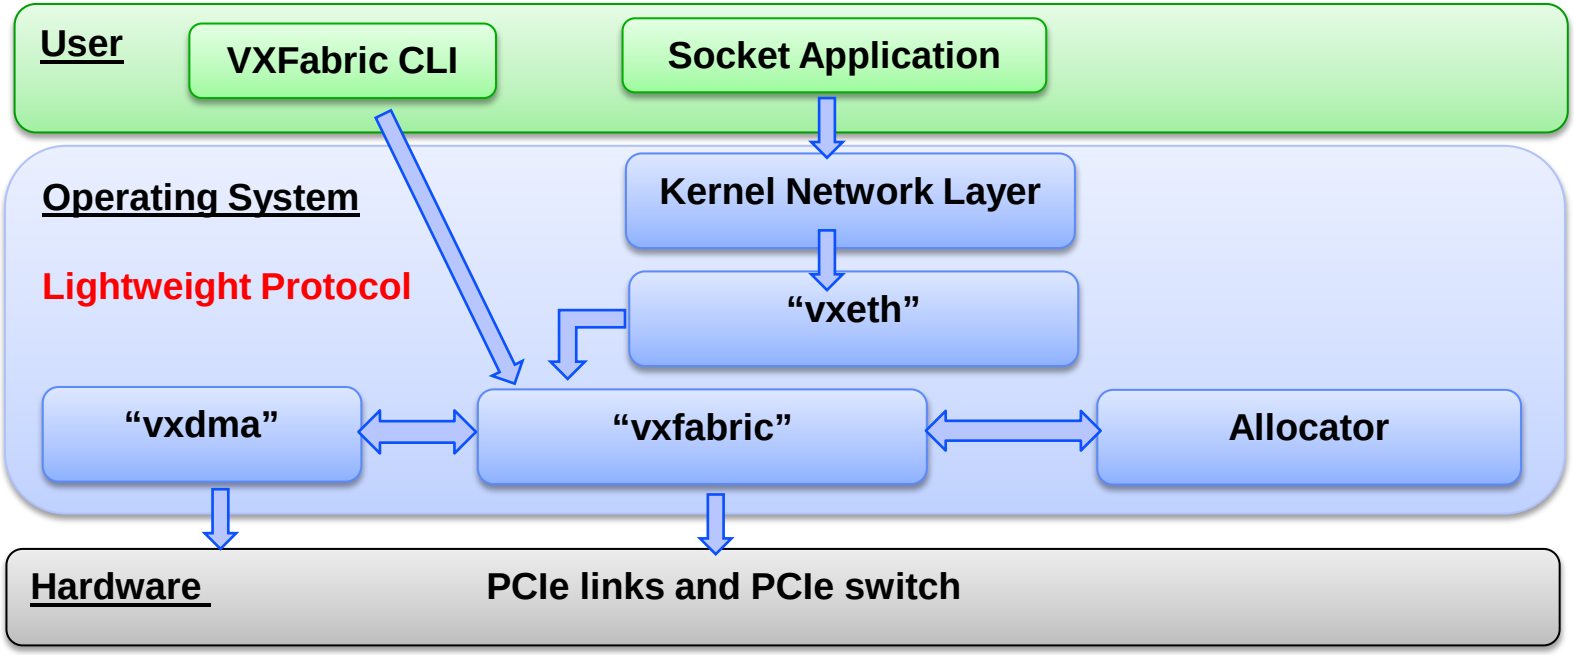

  - **"vxfabric" deals with the mappings, status and low level protocol of the fabric**
  - **"vxdma" implements the DMA API used by other modules to read/write data over PCIe**
  - **"vxeth" creates the Ethernet emulation of VXFabric**

VXFabric Performances under TCP/IP

- VXFabric measured sustained performance, per PCIe x2 Gen2 link

Iperf Linux benchmark	Server throughput	Client throughput	CPU-Load for VXFabric low level protocol Server	Client
1 client	4.28 Gb/s	4.28 Gb/s	4%	3%
2 clients	5.624 Gb/s	2.8 Gb/s	6%	3%

- The VXFabric software lightweight protocol frees up more of the CPUs bandwidth for processing data
  - Aggregate performance for a 12-node VXFabric exceeds 51 Gb/s for unidirectional transfers, and 102 Gb/s for bidirectional PCIe transfers
  - Latency for each packet is low: order of magnitude is a few µseconds

- VXFabric is scalable

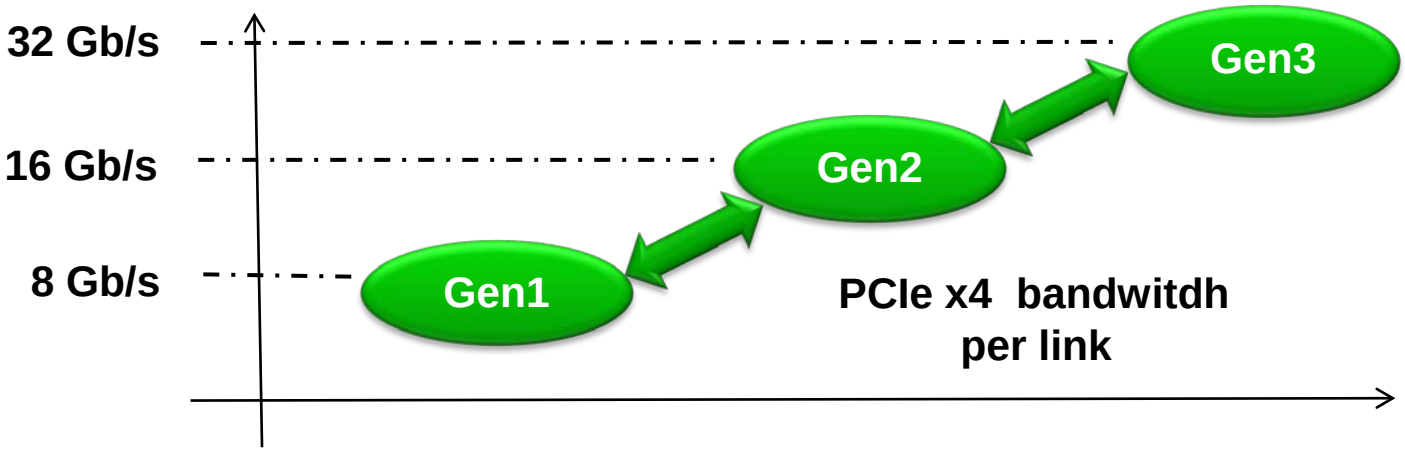

VXFabric Conclusion

- VXFabric high performance switch fabric running on the Kontron VPX Computers
- The VXFabric API, based on IP Socket, is plug & play for most applications, and does not require any long integration/troubleshooting phase: moving from an Ethernet switch topology to VXFabric is straight forward

• Have a look to our VXFabric poster and materials at Poster B.7, Demo B:Technologies and Systems

THANKS !