

Enabling Multi-Port 1 Gbps and 10 Gbps TCP/iSCSI Protocol Offload Solutions

The Virtex-5 LXT platform enables low-footprint, system-level, multi-port 1 Gbps and 10 Gbps TOE solutions.

by Sriram R. Chelluri
Senior Manager, Storage and Servers
Xilinx, Inc.
sriram.chelluri@xilinx.com

As the data center network infrastructure migrates to 10 Gbps, moving data traffic to an Ethernet-based solution becomes economically viable without sacrificing performance and latency. Hardware-based host interfaces like PCI Express and multi-Gigabit Ethernet (GbE) support open up design possibilities for low-cost, high-performance products in the computer and data-processing markets. The Xilinx® Virtex™-5 family of FPGAs sets the stage for designing system-on-chip (SoC) solutions with higher functionality and low power.

The Virtex-5 architecture brings to market critical features that make SoC designs easy to implement for TCP and iSCSI offload engines:

- Built-in PCI Express (PCIe) block – An integrated standards-compliant PCIe endpoint for supporting one to eight

lanes, capable of providing as much as 32 Gbps full-duplex performance

- Built-in Gigabit Ethernet MAC (GEMAC) – four hardcore GEMACs enable multi-port gigabit solutions, reducing total real estate requirements for SoC designs
- Real six-input LUT (6-LUT) technology – improves slice utilization and reduces routing latency for high performance
- 36-Kb dual-port block RAM – higher memory density with error-correction circuitry enables support for reliable computational logic structures and increased on-chip TCP sessions for simultaneous transmit and receive operations
- DSP48E slices – enable massively parallel computations for image processing and multimedia applications

Because the Virtex family is a programmable platform, you can adapt your designs to changing standards and market requirements. Leveraging the resources available in

the Virtex-5 family, you can design cost-effective TCP and iSCSI offload solutions for the server, storage, multi-protocol switch, and wireless base station markets with extended product life cycles.

TCP Offload Engine (TOE) Overview

Current TCP offload solutions rely on a complete software stack or on special network interface cards (NICs) based on ASICs for handling TCP/IP processing. An all-software solution is acceptable for low-bandwidth applications, but high-performance applications would consume all of the CPU resources, creating a system bottleneck for critical applications.

ASIC-based solutions are primarily from start-ups looking to capitalize on the high-performance 10 Gbps market. These solutions are still expensive and prone to vendor lock-in with an uncertain financial future.

Xilinx and its third-party IP partners provide fully standards-compliant TCP/iSCSI offload solutions that you can implement as is or customize for functionality, size, speed, or the target application.

FPGA-Based TCP/iSCSI Engine

With standards-compliant built-in GEMACs, a PCIe core, and increased block RAM, the Virtex-5 LXT device is a programmable platform chip that system architects can exploit for TCP and iSCSI protocol processing without worrying about serial connectivity issues on the network or host interface side. Some of the protocol offload design challenges are:

- The number of TCP connections to support
- TCP reassembly/reorder
- IP fragmentation and reassembly
- Latency
- On-chip versus off-chip TCP session management

These issues can be mitigated with the unique features of Virtex-5 devices and available IP cores. With built-in GEMAC and PCIe interfaces, you can implement direct memory access solutions with minimal FPGA resources, reducing memory transfer latencies and enabling TCP reassembly without using temporary memory. Virtex-5 FPGAs also feature a 36-Kb dual-port block RAM, allowing you to support twice as many TCP connections than previous generations. With the Xilinx LogiCORE™ high-speed memory controller, you can use external DDR2 memory to scale TCP session management. Let's look at the resources you could save in an FPGA-based network interface solution.

1 Gbps and 10 Gbps NIC Solution

An integrated multi-port 1 Gbps and 10 Gbps TCP offload NIC for IP storage and bladed servers enables companies to leverage network infrastructure for storage traffic. Figure 1 shows a typical FPGA-based NIC design.

Depending on the IP cores used, this design can take as many as 20,000 slices to implement. The Virtex-5 LXT platform can reduce resource utilization by 50%, enabling you to develop lower cost solutions without sacrificing performance. Besides hardware efficiency, system archi-

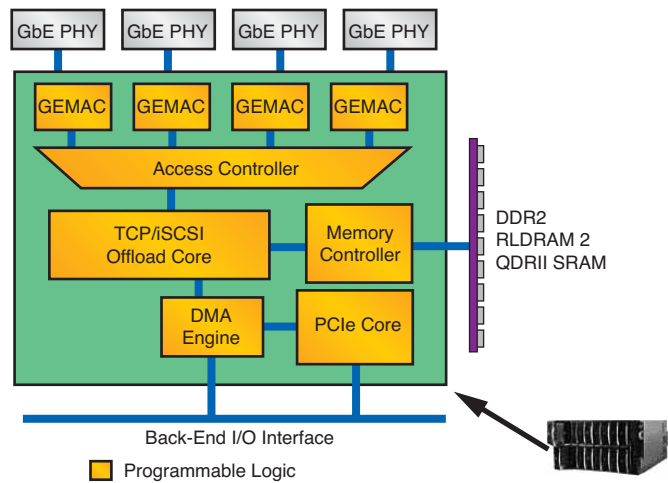


Figure 1 – Designing a TCP offload solution with traditional FPGAs

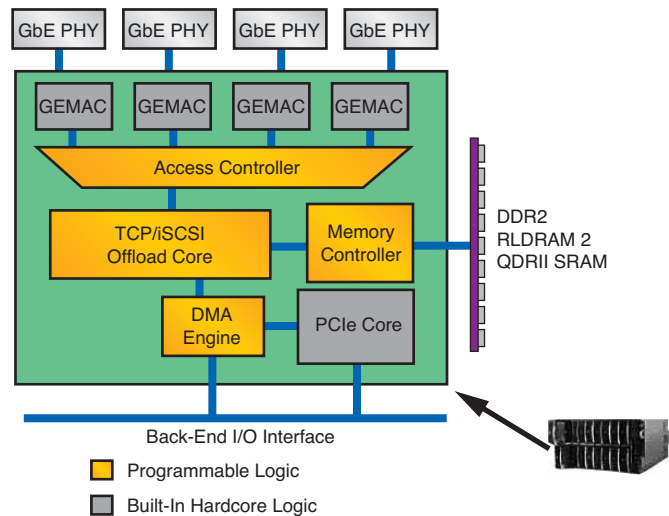


Figure 2 – Designing a TCP offload solution with Virtex-5 LXT FPGAs

ects can also reduce NRE costs because they are not required to implement high-speed I/O interfaces for GbE and PCIe. Figure 2 shows a redesign of the TCP offload NIC leveraging the built-in resources of the Virtex-5 family.

Conclusion

With standards-compliant TCP and iSCSI offload IP cores from third-party vendors implemented on Xilinx FPGAs, you can now design a drop-in or custom SoC at a much lower total cost of development. The

Virtex-5 LXT platform – with hardened GEMACs and PCIe Endpoint blocks, larger block RAMs, and 6-LUTs – uses fewer FPGA resources to implement complex solutions for the server, storage, multi-protocol switch, and wireless base station markets.

To learn more about Virtex-5 LXT FPGAs, visit www.xilinx.com/virtex5. To learn more about protocol offload solutions, visit www.xilinx.com/esp/storage/. And to learn how Xilinx FPGAs can help you in other applications, visit www.xilinx.com/esp.