Tehuti Networks Solutions

The Need for Host Hardware Offload

To meet the increasing needs of server, storage, and network appliance workloads, equipment manufacturers are looking for ways to accelerate packet processing to improve system performance.

Current Industry Needs

The Problem

IT personnel complain, “I've upgraded my system but my network applications are not responding any faster”. The root issue is the transportation of data from the network up to the applications. CPU performance and network bandwidth have improved, but not at an equal pace. Network bandwidth increases have outpaced CPU performance. Add to that, the method for moving TCP/IP data from the network into the applications has remained relatively unchanged.

Background of the problem

The Transmission Control Protocol/Internet Protocol (TCP/IP) is a suite of protocols that have been upgraded continuously since their conception between 1969 and 1973; however, the core of the protocol has remained essentially unchanged for more than a quarter century. As functional as the protocol has been, it was not designed to handle today's high-speed, high-performance networks.

The industry has transitioned to Gigabit Ethernet (GbE) speeds and is now shifting to 10 GbE. Increasingly more expensive and higher performance processors are required to keep pace with network throughput while due to design constraints the memory subsystem cannot physically keep up with the rest of the system. This means that latency is increasing as processors wait through hundreds of idle cycles while memory works at its slower speed to move data across the network to the data starved waiting applications, thereby requiring some sort of hardware assist of the processor to achieve near 10 GbE line rates.

Network bandwidth performance has been improving at a higher rate than Moore's Law

Servers and networks are the corner stone upon which today's industries operate. They perform functions as simple as email and web services to as complex as intrusion detection, sensor networks and grid computing. Since nearly all of the data required for those applications moves between systems in the form of packets with their inherent overhead, it is causing a system bottleneck that is fast approaching a critical level. The basic problem is that packet processing has to be accelerated, to prevent the system from being bogged down. The question now is what is the best approach to solve this fundamental problem?

This disparity, caused by the increasing arrival rate of packets from the network, is impacting the performance of today’s systems designed to handle the critical real-time transactions coursing through today's increased bandwidth networks. Due to increasing demands presented by powerful applications like e-Commerce, medical imaging and data warehousing, today's servers are designed to process larger files and move more data faster than ever before. For example, the 64-bit PCI-X bus architecture pumps out data in the multi-gigabit range and the PCI Express architecture extends this to handle 10 Gigabit Ethernet.

The solutions listed below have been brought to market to address the system bottlenecks.

What was the market's response?

The industry has attempted to address the system bottlenecks through a number of different approaches. Unfortunately each has significant limitations that inhibit wide-spread adoption in future system designs.

- Traditional approaches such as selective, albeit limited, offloading of specific tasks (for example, IP checksum) have been performed for several years. However, as traffic volume and network bandwidth increase, the performance of these specific tasks provides very little overall gain.
- Traditional TCP/IP offload engines (TOEs) attempt to address network I/O bottlenecks by assuming some of the burden of TCP/IP processing and freeing up CPU cycles for applications. TOEs typically require both extensive amounts of their own memory in addition to host memory to perform the offload functions. If the TOE relies on the host memory, it must be able to read or write to the host memory when required. This operation requires increased memory access and increased host bus interactions. The increased host bus interactions increases latency and decreases packet processing throughput. If the TOE uses its own memory, it must have enough memory to handle all the possible connections which results in a very large, expensive, and high power chip. In some instances, a large external memory is required to handle all the connections which further increases real estate requirements and cost.
- The RDMA protocol allows one computer to transfer data directly into a specific destination memory on another computer. It minimizes demands on bandwidth, reduces the overhead for processing packets, and reduces latency. The improvement though is limited to the improvements in the movement of data from the network into the host memory. There are no improvements to the processing of the TCP/IP overhead, and in many cases, RDMA requires modifications to the protocol stack.

Tehuti Networks Solution

Why hasn't the market accepted any of these responses as a way to increase performance? The main reasons are:

- Power budgets
- Numberous external chip requirements and footprint
- Cost
- Scalability

Therefore a different approach was needed to address the limitations associated with TCP/IP processing to meet the above customer's requirements. Tehuti Networks discovered that a series of relatively minor, but important, changes could be made to enhance the performance without adversely affecting the cost, ease of integration and the scalability.

The optimal response is an approach called host hardware offload.

Tehuti Networks' improves the flow of packets in enterprise applications to near 10 GbE line rates without increasing power or cost.

Implemented from wire to application, Tehuti Networks' host hardware offload technology dramatically speeds up enterprise platform to network communications by reducing latency, redistributing functions to components that will do each task better, and using additional enhancements as appropriate to improve server system throughput. With Tehuti Networks' technology, enterprise platforms can finally begin to take advantage of the high-speed, high-performance bandwidth that is available today.

About Tehuti Networks

Tehuti Networks is a fabless semiconductor company that develops innovative solutions-on-a-chip (SoC) for offloading TCP/IP processing in server and storage, and devices. The solutions improve processing performance and provide significant cost benefits to original equipment manufacturers and IT users.

More detailed information is available at:
www.tehutinetworks.net